

HP OpenView Performance Insight

Administration Guide

Version: 5.0

HP-UX, Solaris, Windows®



Manufacturing Part Number: J5223-90026

April 2004

© Copyright 2004 Hewlett-Packard Development Company, L.P.

Legal Notices

Warranty

Hewlett-Packard makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

A copy of the specific warranty terms applicable to your Hewlett-Packard product can be obtained from your local Sales and Service Office.

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause in DFARS 252.227-7013.

Hewlett-Packard Company
United States of America

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Copyright Notices

© Copyright 1992-2004 Hewlett-Packard Development Company, L.P.

No part of this document may be copied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard Company. The information contained in this material is subject to change without notice.

Trademark Notices

HP; Hewlett-Packard Company; and the HP logo are registered trademarks or trademarks of Hewlett-Packard Company. Microsoft®, Windows®, MS Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation. Oracle® is a registered US trademark of Oracle Corporation, Redwood City, California. All other product names are the property of their respective trademark or service mark holders and are hereby acknowledged.

Related Documentation

This document is part of the OVPI documentation set, which also includes the following guides:

- *Performance Insight Guide to Building and Viewing Reports*
- *Performance Insight Installation Guide*
- *Performance Insight Release Notes*
- *Performance Insight Reference Guide*

In addition to customizing existing reports and learning how to create your own reports, you also might want to explore advanced features that allow you to integrate OVPI with other systems and import data from other sources. Procedures for using these more advanced features are found in the following guide: *Performance Insight TEEL Reference Guide*.

You can obtain this guide by attending OVPI customer training (go to <http://www.hp.com/education/courses/u1616s.html> for information) or by calling your HP sales representative.

The OVPI guides are bundled with the OVPI software on the product CD. These guides are also posted to the following Web site:

<http://support.openview.hp.com>

Because updated guides are posted to the web site on a regular basis, the PDF available from the Web site is likely to be more current than the PDF included on the CD. **Always** check for updates on the support web site before using any PDF that was shipped with on the CD.

To locate guides, select **product manuals** from the **using HP software** section; this displays the product manuals search page. Select **performance insight** from the select product box, and then select a version number from the select version box. To see the user guides for report packs, datapipes, and preprocessors, select **reporting and network solutions**, and then select a version number from the select version box.

Every title on the Product Manuals Search page indicates the date of publication. This date comes from the guide's title page.

Online Manuals

The documentation shipped with OVPI is available in Portable Document Format (PDF), which you can view using Adobe Acrobat Reader. You can find the PDF documents in the following locations:

- The OVPI Installation CD—The documentation set and the Adobe Acrobat Reader software are available *before* you install OVPI in the `Documentation` folder located on the OVPI CD-ROM.
- Your system—The documentation set is available *after* you install OVPI in the `docs` directory under the OVPI installation directory (for example, `installation_directory\docs` on a Windows system).

Technical Support and Training

You can find technical support and training information on the HP OpenView World Wide Web site at:

<http://openview.hp.com/>

There you will find contact information and details about the products, services, and support that HP OpenView offers.

You can go directly to the HP OpenView eCare web site at:

<http://support.openview.hp.com/>

The eCare support site includes:

- Downloadable documentation
- Troubleshooting information
- Patches and updates
- Problem reporting
- Training information
- Support program information

contents

Chapter 1	Performance Insight Overview	17
	Understanding How OVPI Works	17
	Collect Data Function	18
	Determining the Types of Data to Collect	19
	Determining the Sources of Data to Collect	20
	Provisioning Reports	21
	Polling Parameters	21
	Process Data Function	21
	Report Data Function	22
	Using OVPI Report Packs	24
	Creating Custom Reports	24
Chapter 2	Configuring OVPI	27
	Post-Installation Configuration Tasks	27
	Additional Configuration Tasks	35
	Configuring User Access	35
	Creating User Accounts	36
	Creating a User Group with a Filter	38
	Configuring a Distributed Installation	41
	Configuring Polling Policies for Remote Pollers	41
	Changing the Polling System	41
	Creating New Policies for Multiple Pollers	42
Chapter 3	Maintaining Your OVPI Installation	47
	Post-Installation Administration Tasks	47

Verifying that OVPI is Running	47
Scheduling Reports	49
Modifying User Accounts	49
Modifying a User Group	50
Modifying a User Group Filter	51
Deleting a User Group Filter	51
Daily Administration Tasks	52
Checking the Status of Data	52
Weekly Administration Tasks	52
Backing Up Your Oracle Database	53
Backing Up Your Sybase Database	53
Backing Up a Sybase SQL Installation Server	54
Backing Up the Sybase SQL Server	58
Backing Up the Sybase SQL Server Using a Script	64
Infrequent Administration Tasks	68
Backing Up Your System	68
Starting and Stopping OVPI Processes	69
Starting and Stopping Processes on Windows	69
Starting a Service	70
Starting and Stopping Processes on UNIX	70
Relocating the OVPI Sybase Database	71
Verifying the Relocation of the OVPI Database	73
Setting Up a Proxy Server	74
Chapter 4 Using the Management Console	75
Starting the Management Console	76
UNIX System	76
Windows System	77
Logging On	77
The Management Console Interface	78
Menu Bar	79
File Menu	80
Edit Menu	80
Tools Menu	81
Window Menu	82
Help Menu	82

Toolbar	83
Changing Database Servers	83
Saving User Preferences	84
Displaying Online Documentation	85
Chapter 5 Discovering the Network	87
Using the Discovery Process	88
Using the SNMP Discovery Wizard	88
Using the Type Discovery Process	96
Starting the Type Discovery Wizard	96
Chapter 6 Installing Report Packs	99
Package Overview	99
Report Packs	100
Datapipes	100
Report Pack and Datapipe Dependencies	100
Starting Package Manager	101
Windows Desktop	101
OVPI Management Console (UNIX and Windows)	101
Command Line (UNIX and Windows)	102
Installing Packages	102
Uninstalling Packages	111
Chapter 7 Managing Polling Policies	113
Polling Strategies	114
Polling Groups	114
Polling Frequency	115
Starting Polling Policy Manager	116
Viewing Functions	118
Sorting Columns	118
Changing the Size of a Column	118
Refreshing a Policy List	118
Selecting Items	119
Viewing Polling Policies	119
Viewing the Current Polling Policies	120

Creating Polling Policies	120
Editing Polling Policies.....	123
Deleting a Polling Policy.....	125
Importing Nodes.....	126
Creating the Import File.....	126
Importing the File.....	129
Managing Polling Groups.....	129
Creating a Polling Group	131
Creating a Single Node Group.....	131
Creating All Nodes of the Same Type Group	132
Creating All Nodes in Same View Group	133
Create Combination of Type and View Group	134
Creating a Specific Instance Group.....	136
Editing a Polling Group	139
Managing Nodes.....	140
Finding a Node	141
Creating Nodes	141
Editing Nodes	143
Deleting Nodes.....	143
Managing Community String Profiles	144
Creating a Community String Profile	145
Editing a Community String Profile	145
Deleting a Community String Profile.....	146
Managing SNMP Profiles.....	147
Creating an SNMP Profile	147
Editing an SNMP Profile	148
Deleting an SNMP Profile.....	149
Managing Datapipe Installations	149
Creating a Datapipe Installation	150
Editing a Datapipe Installation	150
Deleting a Datapipe Installation	151
Modifying Property Table Values.....	152
Chapter 8 Managing Managed Objects	155
Viewing Managed Objects	156

Changing the Objects View	158
Finding a Specific Object	159
Viewing Log Information	160
Modifying the Objects Tree Display	160
Importing Nodes	163
Creating the Import File	163
Use a comma as a placeholder for any parameter that you do not include; otherwise, the node record will be skipped. Table 7 on page 126 describes the parameters for the import file.	163
Importing the File	163
Exporting Nodes	164
Using Forms	164
Modifying the Database	164
Creating New Managed Objects	165
Chapter 9 Grouping Managed Objects	167
Grouping Types	167
Grouping Methods	168
Grouping Nodes	169
Grouping Other Managed Objects	169
Group Categories	170
Viewing Managed Object Groups	170
Changing the Groups View	172
Finding a Specific Group	173
Creating Groups	174
Creating an Enumerated List Group	174
Creating a Derived Group	176
Creating a Property Group	178
Creating a Folder Group	180
Importing Groups	182
Exporting Groups	182
Managing Groups	183
Viewing a Group	184
Editing a Group	184
Deleting a Group	185
Viewing Log Information	186

Setting the Display Options	187
Managing Group Polling Policies	188
Creating a Group Polling Policy	188
Modifying a Group Polling Policy	191
Deleting a Polling Policy	192
Changing the Polling Policy View	192
Using Groups in Reports	193
Chapter 10 Managing Your System	195
Starting the System and Network Administration Application	195
The System/Network Administration Window	196
Displaying System Information	197
Changing Web Access Server Appearance in the Viewing Pane	199
Changing the Details Display	199
Managing Your Web Access Server Systems	200
Adding a Web Access Server	201
Viewing and Modifying Web Access Server Properties	201
Synchronizing Web Access Servers	204
Changing Your Web Access Server	205
Deleting a Web Access Server	206
Managing Your Databases	206
Adding a Database	206
Viewing Database Properties	208
Changing Database Initialization Parameters	210
Increasing Database Size	211
Database Component Sizing Guidelines	211
Increasing the Size of the Database	212
Changing the Database Server	215
Chapter 11 Managing Tables	217
Starting Table Manager	218
Viewing Database Information and Statistics	219
Viewing Database Information	219
Viewing Database Statistics	220
Monitoring Database Usage	220

Viewing Table Information	221
Viewing Functions	221
Sorting Information	221
Changing Column Size	222
Refreshing Information	222
Finding Alias Names	222
Viewing Data Tables by Type or Category	223
Viewing Property Tables	225
Changing the Table Column Display	225
Viewing Data and Internal Table Details	227
Viewing Data Table Column Information	228
Setting Data Table Storage Times	228
Adding a Data Table Description	229
Viewing Data Table Column Information	229
Viewing Data Table Index Information	230
Viewing Property Table Details	231
Viewing Property Table Information	232
Adding a Property Table Description	232
Viewing Property Table Column Information	233
Viewing Property Table Index Information	234
Displaying Table Data	235
Displaying Reports	235
Using the Menus	236
Managing Data	238
Printing Table Data	238
Saving Table Data	240
Exporting Table Data	241
Formatting Data	241
Formatting Statistics	242
Formatting Cells	243
Formatting Columns	244
Formatting Rows	247
Changing Table Queries	248
Selecting Statistics	248
Sorting the Data in the Report	248
Changing the Time Period	250

Adding Constraints	252
Changing Constraint Values	255
Modifying Constraints	256
Deleting Constraints and Clauses.	256
Selecting Nodes and Interfaces	256
Changing the Maximum Number of Rows Returned	258
Creating Tables.	259
Creating Data Tables	260
Creating Property Tables	275
Setting Table Defaults	277
Setting Data Table Defaults	278
Time-Period Table Default	279
Setting the Display Options	280
Chapter 12 Using the Web Access Server	283
Accessing the Web Access Server.	284
Web Access Server Overview	286
Administration Link	287
User Accounts.	287
Service Configuration	288
Log Files	290
Java Settings	290
Report Settings.	291
Catalog Link	293
Maintaining User Accounts	295
Setting Password Validation Rules	295
Creating a User Account.	297
Modifying a User Account.	301
Deleting a User Account	302
Sending Email to a User Account.	302
Maintaining User Groups	303
Understanding How User Groups Work	304
Creating a User Group	304
Modifying a User Group	308
Deleting a User Group	308
Creating a User Group Filter	309

Modifying a User Group Filter	309
Deleting a User Group Filter	310
Changing Log File Settings	311
Viewing Log Files	312
Changing Java Settings	314
Changing Report Settings	315
Creating Views	318
Creating a View	318
Modifying a View	320
Changing the Appearance of the Web Access Server	321
Changing Site Options	321
Changing Link Bar Options	322
Changing Link Bar Aliases	324
Changing Navigation Frame Options	325
Changing Results Frame Options	326
Deploying Reports	327
Changing Settings for the Services	329
Changing the Web Access Server Settings	329
Resetting Administrator Privileges	332
Changing the SSL Service Settings	333
Importing Your Digital Certificate	335
Generating a Certificate Signing Request (CSR)	335
Importing a Signed Digital Certificate	336
Changing System Manager Settings	337
Displaying the Performance Insight Systems Page	337
Choosing the Default System	338
Adding a System	338
Modifying a System's Settings	341
Deleting a System	345
Changing Mail Server Settings	345
Modifying Graph X-Axis Time Formats	347
Chapter 13 Tuning the System	351
Hardware Considerations	351
Processor Type and Clock Speed	352

Number of Processors	352
Memory	352
Shared Memory / Swap File Size	353
Disk Drive Interfaces	353
Number of Physical Disk Devices	354
Configuring Operating System Parameters	355
Configuring Oracle Parameters	355
Configuring Sybase Parameters	356
Using the Database Tuner	356
About Sybase Parameters	361
Using the sp_configure Command	367
Operating System - Asynchronous I/O Processing	368
Transaction Log Cache	368
Transaction Log I/O Size	370
Traceflag Parameters	371
Other Configuration Parameters	373
Binding the Property Tables	374
Chapter 14 Managing MIBs	379
Starting the MIB Browser	381
The MIB Browser Window	382
Navigation Pane	383
Viewing Pane	383
Menu Bar and Toolbar	385
File Menu	385
View Menu	385
Actions Menu	386
Help Menu	386
Loading MIBs	386
Issues to Consider	387
Including the Top of the MIB Tree	387
Validating Object Types	387
Loading a MIB	389
Unloading a MIB	391
Querying SNMP Data	392

Setting the Node	393
Setting the Community String	393
Finding a Group or Object in the MIB Map	393
Viewing the Data	394
Setting SNMP Values	397
Creating Collection Tables	397
Collection Table Features	397
64-Bit Values in SNMPv1	398
Collection Table Requirements	399
Creating a Collection Table	400
Chapter 15 Interpreting Log Files	413
The OVPI Log Files	413
Viewing Log Files	414
Log File Message Format	415
timestamp	415
time_zone	416
GMT_offset	416
appl_name	416
module_name	416
severity	417
process_ID	417
parent_proc_ID	417
message_ID	417
message	418
Example Log File Messages	418
WARNING and INFO Messages	418
ERROR Messages	418
The audit.log File	419
Log File Format	419
The metrics.log File	421
Log File Format	422
mw_collect Messages	423
bcp_gateway Messages	424
trendcopy messages	425
The builder.log and viewer.log Files	426

Appendix A Troubleshooting	429
OVPI Client Applications	429
Oracle Unable to Resolve a Service Name	429
Displaying the Navigation Frame for the Web Access Server	431
Java Stack Trace Entries in the website.log File	432
Web Access Server Help Does Not Display Properly	432
Troubleshooting Sybase Space Allocation Problems	432
Appendix B Rendering Graphics Using Hardware	435
Configuring X	436
Index	439

Performance Insight Overview

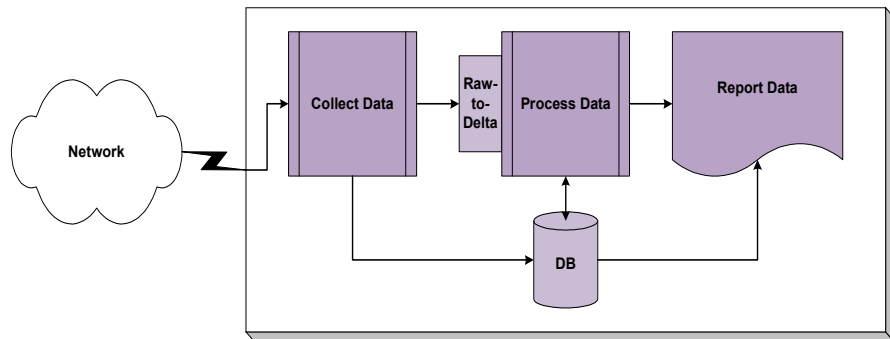
This chapter describes how HP OpenView Performance Insight (OVPI) works.

Understanding How OVPI Works

OVPI is a network management system that performs the following functions:

- Collect Data
- Process Data
- Report Data

The following figure shows an overview of these functions.



These functions work as follows:

- **Collect Data**

This function identifies the nodes (network devices or systems) in your network, and then obtains data from those nodes by polling them directly or by importing their flat file extractions. A *flat file* is a self-contained set of historical information for nodes produced by network device utilities on a periodic basis. In either case, the system stores the collected data in a database. See [Collect Data Function on page 18](#) for details.

- **Process Data**

This function converts raw data into rate data, performs summarization routines on data, and ages data out of the database. The summarization routines convert rate data into hourly, weekly, monthly, quarterly, and yearly data. They also produce forecasting data, which allows you to see what could happen based on past performance. See [Process Data Function on page 21](#) for details.

- **Report Data**

This function produces reports from the summarized data. It automatically generates the reports based on a reporting solution, known as a report pack. A *report pack* provides the report templates and all the supporting files and scripts for a particular type of solution.

You can view these reports by using a web-based interface (Web Access Server) or by using a client application (Report Viewer). It is also possible to create your own customized reports with another client application (Report Builder); however, you should contact your HP representative if you want to use this feature. See [Report Data Function on page 22](#) for details.

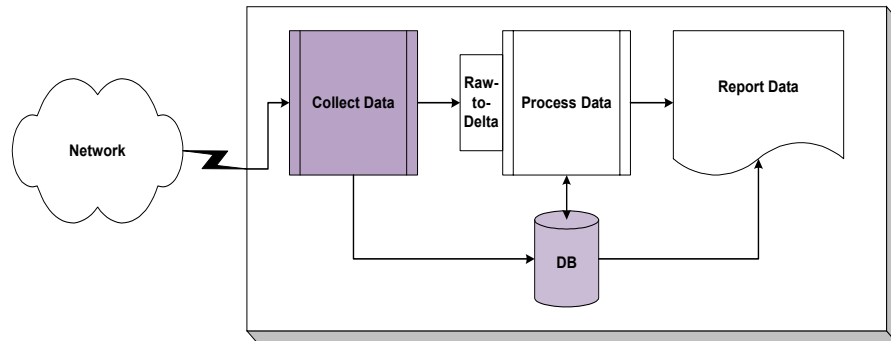
OVPI's trendtimer program controls when all these functions occur. This program runs as a service (Windows) or as a daemon (UNIX). It triggers the various processes that are part of the OVPI system. See [Chapter 16, Managing Services](#) for details about the trendtimer program.

Collect Data Function

When OVPI collects data, it must determine what type of data to collect as well as where it comes from and when to collect it. OVPI can collect data from following sources:

- Simple Network Management Protocol (SNMP)-type devices
- OpenView performance agents (OpenView Performance Agent and OpenView Operations Agent)
- Non-SNMP devices

In this context, SNMP-type devices include Remote Monitoring (RMON) devices. The following figure shows the Collect Data function in the overall OVPI system.



Determining the Types of Data to Collect

OVPI uses a list of object identifiers (OIDs) from SNMP-type devices to determine what data to collect. The OIDs are in management information base (MIB) tables, which are sets of defined objects. OVPI automatically provides these tables when you install a report pack.

You can set up your own collection using MIB Browser or a data definition language called TEEL with `datapipe_manager` to define the table.

When OVPI collects data from a non-SNMP device, it uses TEEL files to import and store the data from the device. OVPI automatically provides the definition file for the corresponding tables and data file when you install a report pack. If you want to define your own collection, you can use Table Manager or `datapipe_manager` to define the table, and then import the data, using the appropriate collection tool (for example, `ee_collect` or `pa_collect`). Contact your HP representative if you want to develop your own solution.

Determining the Sources of Data to Collect

For SNMP-type devices, the data comes from the nodes. You can use any of the following methods to let the system know which nodes to poll:

- **SNMP Discovery.** Lets you use a wizard to ping a range of network addresses to locate the nodes that respond. See [Using the Discovery Process on page 88](#) for details.
- **Polling Policy Manager.** Lets you enter nodes one at a time. See [Chapter 7, Managing Polling Policies](#) for details.
- **node_manager.** This utility lets you import an existing list of nodes into the system. In this case, you set up an ASCII file with the required information and then import the nodes using the utility. For more information about the `node_manager` command, refer to the *Performance Insight Reference Guide*.

After the nodes have been specified, you can use the Type Discovery process to identify and assign the nodes into groups. Typically, nodes go into a group specific to data or device of a report pack, which is the *type* list. However, if you want to poll only a specific set of nodes, you can create your own groupings, which are *view* lists. See [Creating a Polling Group on page 131](#) for more information about defining a node list. When the poller polls the nodes in these lists, it places the data that it acquires into data tables.

To enable data collection from OpenView performance agents that have SNMP enabled, use SNMP Discovery to discover these nodes. For non-SNMP-enabled systems with OpenView performance agents, use `node_manager` to import a list of nodes.

For non-SNMP devices, data comes from a flat file. The method to obtain the data from these flat files requires setting up a TEEL file to provide the instructions for importing the data. The TEEL files also define the tables to store the data from these files. You can use TEEL files to create SNMP and non-SNMP tables. You must place non-SNMP data on the server that contains the Performance Manager component in order for the data to be collected properly; SNMP data is collected directly from the source node. Once the data is located in the database, all data is processed the same regardless of the way in which it was collected.

Provisioning Reports

Report packs may provide a property import datapipe that allows you to *provision* your report pack with information specific to your network. This process involves importing static data specific to your network using any of the following methods:

- Create your own flat file from a spreadsheet program or from a network provisioning database that automatically generates a file.
- Export existing data from the database.
- Modify selected fields interactively using Polling Policy Manager.

The best time to modify an existing file is after the type discovery process runs the first time and before any other processing executes. At this point, you create or modify the file, and then import it back into the proper location. The specific report pack documentation includes the instructions for modifying and importing these files.

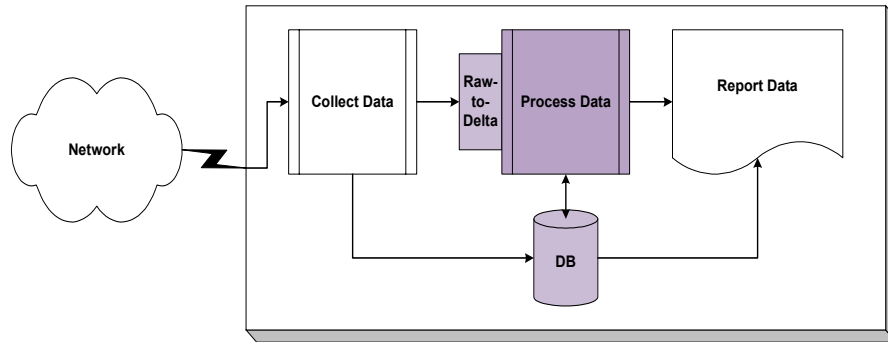
Polling Parameters

After the system locates the nodes to poll, it needs know what data to collect and how frequently. The system obtains these instructions from a *polling policy*. Report packs contain the polling policy definitions; however, if you set up a distributed environment you may need to modify an existing polling policy or create a new one. See [Chapter 7, Managing Polling Policies](#) for more information about polling policies.

Process Data Function

OVPI primarily manages time-series data. Accordingly, each record stored in the database must have an associated timestamp. In addition to a timestamp, each record has an indicator that identifies the *managed object* associated

with the data. Some typical examples are router, interfaces, customer, and location. The following figure shows the Process Data function in the overall OVPI system.



Once the data is collected, the system assigns timestamps to the records and stores the data in *data tables*. The system will then convert the raw data into *delta* (or rate) data. Delta data is the difference between two values over time.

Once there is delta data, the aggregation process continues. During the aggregation process, the system reads data from a *source* (input) *table*, then aggregates it, and stores it in the *destination* (output) *table*. The aggregation process creates different destination tables for the different types of statistics. Statistics include hourly, weekly, monthly, quarterly, and yearly summaries; 30-, 60-, and 90-day forecasts; percentiles; and averages. Report packs contain all the necessary instructions to perform the multiple aggregations.

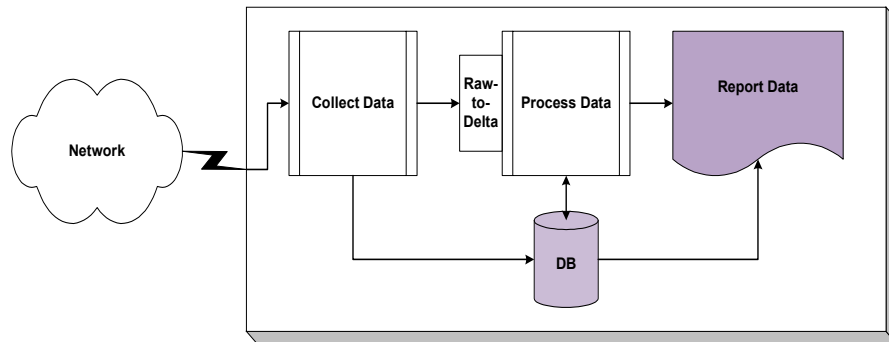
To prevent the accumulation of excessive data in the tables, OVPI ages the data. Aging involves the concept of *retention time*, which is the number of days to keep the data in the table. The aging process deletes the data when it has been in the table longer than the specified retention time for that table. Each table has a default retention time; you can change the retention time using Table Manager. See [Setting Table Defaults on page 277](#) for more information about changing the retention time for a table.

Report Data Function

After the data is processed, it is available for producing reports. You can produce reports using either of the following methods:

- **Automatic.** This method uses OVPI report packs to create, deploy, schedule, and generate the reports. See [Using OVPI Report Packs on page 24](#).
- **Custom.** This method lets you create custom reports using Report Builder. See [Creating Custom Reports on page 24](#).

Both methods create the report template, which is a set of instructions that describe the report. The following figure shows the Report Data function in the overall OVPI system.



Reports are available in the following formats:

- A *report definition* file (.rep file extension) contains multiple tables, graphs, graphics, and text and retrieves its data from the database.
- A *dataset* file (.srep file extension) is similar to a report, except when you view a dataset report, its data is retrieved from the report *not* the database. A dataset report is a snapshot of the data taken at the time the report is generated. This is performed by a batch generation process, which is typically set up by the Administrator. Like a report definition file, you can change many of the display attributes of a dataset report.
- A *form* (.frep file extension) is a special kind of report. A form lets you change the data contained in a single database table. All elements that are available for reports are available for forms. In addition, you add form elements, which allow you to change the data in the database table. A form is an easy way to modify or create new data in the database. You might want to create a form that will let you update a customer in the database or add a new customer to the database for a specific report pack.

Using OVPI Report Packs

By installing a report pack for a particular technology, your system is automatically set up to create, schedule, and generate reports.

Report packs provide the following:

- Table definitions to store the data for the solution
- Collection information (including scripts, polling policies, and commands) to collect the data for the solution
- Aggregation instructions to generate the summary data for the solution
- Report templates to generate the reports for the solution
- Deployment instructions to identify the location for the report templates
- Instructions and commands for automatic scheduling and generation of reports

Creating Custom Reports

You can use Report Builder to create a report template. When you create the report template you define how to display the data, what information you want to use from the database, and how to link the data to the report. To display the data in the report template, you can use a graph or table. You can also include graphics and text. Thus, a completed report can comprise graphs, tables, graphics, and text.

After you create a report, you will want to decide which application to use to view it, and then set up a schedule to generate the report. You can view the report using either of the following applications: Web Access Server or Report Viewer.



To view a report using the Web Access Server, you must deploy the report to the Web Access Server before you can view it. You can deploy a report using the Web Access Server, the Deployment Wizard (an option available from the Report Builder and Report Viewer applications), or the Deployment tool, a command-line interface.

When you are ready to view a report, you view it using the Web Access Server application, where the report will retrieve the latest data for the report. Additionally, you can schedule a report for generation if you want the data it contains to reflect a static point in time or you want to produce the report in

.pdf or .csv formats. You can generate a report according to a generation time specified by a schedule or you can generate a report independently of its scheduled generation time. You can schedule a report for generation using the Web Access Server or the Schedule tool, which is a command-line interface. The generation process combines a defined report template with data from the database to create the final report.

Refer to the *Performance Insight Guide to Building and Viewing Reports* for more information about Report Builder, Report Viewer, and Web Access Server applications. Contact your HP representative if you want to create custom reports.

Configuring OVPI

This chapter describes how to configure OVPI after you have successfully installed it in a standalone or distributed environment.

Post-Installation Configuration Tasks

This section describes the configuration tasks that you should perform after you have successfully installed OVPI in a standalone or a distributed environment.

The tasks for configuring OVPI are as follows:

- Tune the database.
- Apply OVPI Service Packs, if necessary.
- Discover your network or import nodes.
- Install report packs.
- Discover device types.
- Change polling policies.
- Verify collections.
- Set up the backup server
- Back up the database

Task 1: Apply OVPI Service Packs

You should apply any Service Packs for the release. To locate any Service Packs for this release, go to the following Web site:

http://support.openview.hp.com/cpe/ovpi/patch_ovpi.jsp

Task 2: Tune the Database

Complete the tuning recommendations for your database (Sybase or Oracle).

Sybase

Run the Database Tuner. See [Using the Database Tuner on page 356](#) for details.



The Database Tuner is only available for Sybase databases.

Oracle

Because Oracle 9i includes the ability to optimize itself automatically, OVPI only needs to set a few parameters in order for the database to reach full optimization.

During the OVPI installation, there was opportunity to change the default values of these parameters. (See “Changing Oracle Configuration Parameters” in Chapter 4 of the *Performance Insight Installation Guide* for details.)

If you did not change the values of these parameters during the installation but would like to do so now, you can change the parameters using Oracle’s Enterprise Manager Console. Refer to your Oracle documentation for additional tuning recommendations.

Task 3: Discover Your Network

Before OVPI can collect data, OVPI must know which devices you want to manage. To obtain this information, you need to discover your network. You can discover your network by using either of the following methods:

- Discover nodes.

This method is useful in environments where you do not have a list of managed nodes provisioned or where NNM or other fault management platforms are not present. For large environments, discovering nodes can be a long process; you may want to consider importing nodes first.

OVPI will discover the nodes that are available to it when you use the Discovery wizard, which is available from the Management Console. It finds the nodes in your network and determines whether or not they are SNMP manageable.

▶ For networks that contain a large number of devices, this may not be a good method to use. Import your nodes instead.

- Import nodes.

This method is useful when you have a list of managed objects that you want to monitor (for example, NNM or other network provisioning software such as Cisco Works). For example, if you only have five routers, you will most likely know the information for them; thus, you can target the discovery towards only the objects that you want to poll.

You can import nodes from other sources (for example, a provisioning database).

▶ Report packs may require some configuration before you discover the network; therefore, refer to your report pack documentation before you begin the discovery process. To locate reporting solutions documentation, see [Related Documentation on page 3](#).

Discovering Nodes

- 1 From the Management Console's Tools menu, choose **SNMP Discovery**. The Discovery wizard opens.
- 2 Type the minimum and maximum values for the IP Address and the subnet mask, and then click **Next**.

▶ Subnet mask settings are only for determining network and broadcast addresses. If the mask is wrong, the discovery process will still occur but nodes may not be added correctly.
- 3 Type the community read string or strings for your SNMP devices, and then click **Next**.
- 4 Enter the SNMP port if it is different than 161, and then click **Next**.
- 5 Click **Discover**.

The Discovery wizard will verify the existence of a device at each IP address and forward a GET request to each address. If the request is answered, the device can be managed by SNMP.

See [Using the Discovery Process on page 88](#) for complete details. Alternatively, you can use the `trend_discover` command. Refer to the *Performance Insight Reference Guide* for details.

Importing Nodes

There are two steps involved when importing nodes. You must first create an ASCII file that contains the nodes you want to import, and then you must import that file. See [Table 7 on page 126](#) for details about how to format the ASCII file.

To import a list of nodes:

- 1 From the File menu, choose **Import Nodes**.
The Import from dialog box opens.
- 2 Navigate to the file that contains the nodes you want to import.
- 3 Double-click the filename, or select the filename and click **Open**.

Alternatively, you can use the `node_manager` command to import nodes using the command line. For more information about the `node_manager` command, refer to the *Performance Insight Reference Guide*.

Task 4: Install Report Packs

This procedure will show you how to install report packs using the RNS CD.



Before you continue with this procedure, make sure that you have obtained and **read** the user guides for the report packs you want to install. These guides contain important configuration information (for example, installing in distributed environments) about the report packs that you should read before you install them. To obtain these guides, see [Related Documentation on page 3](#).

- 1 Log on to the local system.
 - *Windows:*
Log on as a user with administrative privileges.
 - *UNIX:*
Log on as `root`.
- 2 Insert the RNS CD in the CD-ROM drive. If you are a UNIX user, you must also mount the CD.

The following table provides instructions for mounting the CD on each of the UNIX platforms:

Operating System	Command
Sun Solaris	<p>The CD should mount automatically when you insert it in the CD-ROM drive. If the CD-ROM mounts automatically, the File Manager window opens. If not, as <code>root</code>, type the following:</p> <pre>mkdir /cdrom mount -r /dev/dsk/<i>cd_devicename</i> /cdrom</pre> <p>In this instance, <i>cd_devicename</i> is the name of the CD-ROM device.</p>
HP-UX	<p>You can use the SAM utility to mount the CD-ROM, or you can type the following commands:</p> <pre>mkdir /cdrom mount /dev/dsk/<i>cd_devicename</i> /cdrom</pre> <p>In this instance, <i>cd_devicename</i> is the name of the CD-ROM device.</p>

- 3 *Solaris systems only:* Do the following:
 - a To display the `cd_label`, type the following command:


```
ls -l /cdrom
```
 - b To change directory, type the following command:


```
cd /cdrom/cd_label
```

In this instance, *cd_label* identifies the CD name.

- 4 The installation program should start automatically. If it does, go to [Step 6 on page 32](#); otherwise, go to [Step 5](#).
- 5 Do *one* of the following to start the setup program:
 - *Windows*:
From the top level folder on the CD, double-click `setup.exe`.
 - *UNIX*:
Locate the setup program, and start it by typing the following command:

```
./setup
```

A menu opens.
- 6 Select **OVPI solution components**. Once the installation script has copied every package to the Packages directory on your system, the installation script starts Package Manager.
- 7 Click **Next** to continue.
- 8 Click **Install**, specify the directory in which you want to install the report packs if you are not using the default shown, and then click **Next**.
- 9 Do one of the following:
 - Select the **Deploy Reports** box. If you select this option, the reports are published to the Web Access Server, so the reports can be viewed there.
Enter an OVPI administrator username and password, and click **Next**. (You can use the default OVPI administrator account, `trendadm`, which was created during the installation.)
 - Click to clear the **Deploy Reports** box.
There may be instances when you do not want to deploy reports to the Web Access Server but you still want to collect data for those report packs. You can deploy reports later, and all the data that has been collected will be displayed.
- 10 Select report packs and datapipes you want to install, and click **Next**.
Dependent packages will automatically be selected if they are located in the same directory as the report packs. If not, the installation will not continue. For information about necessary datapipes for the report packs you are installing, see the report pack's user guide.

- 11 Select or click to clear **Run OVPI Discover**, and then click **Next**.

OVPI Discover determines which devices in your network can be polled for the report packs you are installing.

It is recommended that you run type discovery for all the report packs that you are installing.

Some report packs add in types that may be of no interest to you. If you are an advanced user and know which types you have in your environment, you may want to click to clear this option, and run the Type Discovery wizard at a later time.

- 12 Verify the information that is displayed on the Selection Summary window, and then click **Install**.

The system should now be configured to collect data for the report pack or report packs you just installed.

To collect data, you need to make sure that you perform a type discovery using the Type Discovery wizard if you did not perform a type discovery in [Step 11 on page 33](#).

Task 5: Discover Device Types.

This discovery determines the device type of a discovered node and then records this information in the database. This information lets OVPI target a specific type of device for data collection.

Do one of the following:

- If you did not perform a type discovery when you installed your report packs (see [Step 11 on page 33](#)), then you can perform a manual type discovery using the Type Discovery wizard, which is available from the Management Console.
- If you performed a type discovery, proceed to [Task 6](#).

To perform a type discovery:

- 1 From the Management Console's Tools menu, select **Type Discovery**. The Type Discovery wizard opens.
- 2 Select the types you want to discover. To select types that are not consecutive, press and hold down the **CTRL** key, and then click each type.
- 3 Select **Discover all known types** if you want to discover all types or if you are unsure, and then click **Next**.

4 Click **Discover**.

Alternatively, you can use the `trend_discover -t` command. Refer to the *Performance Insight Reference Guide* for details.

Task 6: Change Polling Policies

A polling policy defines the frequency by which the defined type of data will be collected. It is also used to specify whether the local system will poll devices for data or whether the polling will be assigned to the remote poller



Be careful when you modify polling policies. Usually, you should change only the policy's polling frequency. If you change tables or groups, it may have adverse effects on your installation.

In a standalone environment, you may want to adjust the frequency of a polling policy. For example, suppose you know your network will be changing frequently one day due to the addition of several new devices. If your current polling policy polls the network once a day, you may want to change the polling frequency of the policy to once every hour in order to poll the new devices as they become available. After the network stabilizes, you would return the polling frequency to once a day.



If you change your policy so that data is collected too frequently, it may severely affect OVPI and your network performance.

For information about configuring polling policies for remote pollers, see [Configuring Polling Policies for Remote Pollers on page 41](#).

To change polling policies:

- 1** Start the Management Console, and then click the **Polling Policies** icon in the navigation pane.
- 2** Select the **All** page in the viewing pane to display all of the current polling policies.
- 3** Locate the policy you want to change, and double-click it. The Edit Polling Policy dialog box opens.
- 4** To change the polling interval, select an interval from the Poll Interval list.
- 5** Click **OK**.

6 Repeat [Step 3](#) through [Step 5](#) for each polling policy.

See [Chapter 7, Managing Polling Policies](#) for more details about polling policies.

Task 7: Back up your system.

It is recommended that you back up your system after you install and configure OVPI. Use the procedure you normally use to back up your file system.



It is recommended that you stop all OVPI services and the database before you back up your system. For information about stopping the OVPI services, see [Starting and Stopping OVPI Processes on page 69](#). For information about stopping Sybase, refer to the *Performance Insight Installation Guide*. For information about stopping Oracle, refer to your Oracle documentation.

Task 8: Back up the database.

It is recommended that you back up your database after you install and configure OVPI and, thereafter, back up the database weekly. See [Backing Up Your Oracle Database on page 53](#) or [Backing Up Your Sybase Database on page 53](#) for details.

Additional Configuration Tasks

This section describes the additional tasks that you can perform to configure OVPI.

Configuring User Access

All OVPI client applications (Report Builder, Report Viewer, Web Access Server, and the Management Console) require users to log on. You configure user access to these applications using the Web Access Server.

By default, OVPI creates one user account when you install it. This is the `trendadm` account; you should create a user account for each person that will use the OVPI client applications.

You can also create groups to organize users, for example, by functional areas or by company. You can then take the groups and apply filters to them; using filters, you can limit user access to information.



When you create users and groups for a distributed installation, they are created on the Web Access Server that you are currently logged on to; this Web Access Server will provide the log on authentication for the Web Access Server application and the Management Console. If you have more than one Web Access Server, then you must make sure that you monitor the users and the Web Access Servers they need to access. For example, if there are users that need to access two different Web Access Servers, you must create an account for those users on each Web Access Server they need to access.

If a user only needs to view reports with data from the different database servers, then you can create a view. A *view* allows users to see a certain list of deployed reports. The view applies to any database that the user may access. See [Creating Views on page 318](#) for more information about creating views.

Creating User Accounts

To create a user account:

- 1 Access the Web Access Server. See [Accessing the Web Access Server on page 284](#) for details.
- 2 From the Links bar, select **Administration**.
- 3 Expand **User Accounts** in the Administration tree, and click **Users**. The User Administration page opens; it lists the current user accounts.
- 4 Click **Add New User**. The Add a New User page opens.
- 5 Enter the requested information into all boxes from Username to E-mail Address, noting the constraints that follow.

Ensure that you do **not** use any of the following characters in any of the fields on this page: ampersand (&), backslashes (\), comma (,), double quotation mark (“), forward slash (/), question mark (?), percent (%), and single quotation mark (‘).

The valid entries are as follows:

- All alphabetic characters (uppercase and lowercase), all numeric symbols, and all special characters except those in the previous list.

- The Username must be a minimum of one (1) character, can be a maximum of 255 characters, and can only contain alphanumeric characters, spaces, underscores (_), and hyphens (-).
 - The user's real name and company/department have no minimum length (that is, they can be blank) and can be a maximum of 255 characters.
 - The password can be a minimum of four (4) characters, a maximum of 255 characters, and cannot contain a space (blank) character.
 - The telephone number can be in any format but cannot exceed 255 characters.
 - The email address can be in any format but cannot exceed 255 characters.
- 6 Select the **Is Administrator** check box only if the user will have an Administrator role. If you do not select this option, the account is automatically assigned a user role.
 - 7 In the Database list, click the database that the user will access. For information about how to add another database to the current list, see [Changing System Manager Settings on page 337](#).
 - 8 In the View list, click the view that the user will use. For information about how to create a view, see [Creating a View on page 318](#).
 - 9 To choose the password expiration setting for the user account, do *one* of the following:
 - Click the **Never Expires** option, which will keep the password for the account indefinitely.
 - Click the **Password Expires** option, and then enter a number of days to set an expiration period for the password.

When a password expires, a page will display informing the user that the password for the account has expired and will give the user the option to change the password for the account.

In addition to this option, you can do the following:

- Set up password validation rules. For example, you can set a grace period for a password. This option lets the user to log on to an account a specific number of times after the password for the account has expired. When this period expires, the account is locked. See [Setting Password Validation Rules on page 295](#) for details.

- Specify the number of unsuccessful log on attempts that a user can have before an account is locked. See [Changing the Web Access Server Settings on page 329](#) for details.

10 Click **Apply**.

11 Click **OK** to confirm the new user account.

After you confirm the new user account, the User Administration page lists the new user account information. You can also view the user account in the navigation tree.

Alternatively, there is a command to create users (**userimport**), and a command that you can use to manage users (**userctl**). Refer to the *Performance Insight Reference Guide* for more information about these commands.

Creating a User Group with a Filter

A *group* is:

- A collection of user accounts

For example, you can create a group called *Thunderbolt* for all users in a company called Thunderbolt, Inc. This is useful if you are a service provider who wants to limit the data that certain companies can view.

- A collection of user accounts and groups

For example, you can create a group called *All*, which contains some user accounts and four groups called *North*, *East*, *South*, and *West*.

The constraint placed on the users in the *North*, *East*, *South*, and *West* groups lets those users view only the data for their specific region.

The users in the *All* group can view all the data that the users in the *North*, *East*, *South*, and *West* groups can view because the top-level group inherits the constraints of the groups below it in the hierarchy.

When you create groups, it is recommended that you try to create effective constraints for the groups to limit the number of groups you might have. For example, instead of using `interface` as a constraint and having 1000 or more groups for each interface, use `cust-id` as a constraint instead. A large number of groups can affect the efficiency and size of the query sent to the database.



If there is a report that only one user should see, deploy that report to the user's report directory under the users folder rather than using a view. See [Deploying Reports on page 327](#) for details.

To create a group with a filter:

- 1 Start the Web Access Server. See [Accessing the Web Access Server on page 284](#) for details.
- 2 From the Web Access Server's Links bar, select **Administration**.
- 3 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens.
- 4 Click **Add New Group**. The Add New Group page opens.
- 5 To create a new group, do the following:
 - a In the Group box, enter the name of the group you want to create.



Do *not* enter a group name that already exists or use the name "Group".

The group name must adhere to the following constraints:


Ensure that you do **not** use any of the following characters in the group name:

ampersand (&), asterisk (*), backslash (\), colon (:), comma (,), double quotation mark ("), exclamation point (!), forward slash (/), greater than (>), left bracket ([), less than (<), period (.), pound (#), question mark (?), right bracket (]), single quotation mark ('), percent (%), and a blank space.

The valid entries for the group name are as follows:

- The name can contain all alphabetic characters (upper case and lower case), all numeric symbols, and all special characters except those cited above.

- Minimum of one (1) character and a maximum of 255 characters.
- b** Select the user or users and the group or groups from the Available Users/Groups list that will belong to this new group, and then click **Add**. The user names, group names, or both are moved to the Members list.
- c** Click **Next**. The Group Filters page opens.

 A *filter* is a comparison that involves a table column name, a numerical operator, and a value for the statistic.

6 To create a filter:

- a** In the Column field, enter the table column name to filter on.

To locate the table column, use Table Manager which is available from the Management Console. Note that only table columns from property tables are valid values in this field. See [Viewing Data Table Column Information on page 228](#) for details.

- b** Click the **Operator** arrow, and then select the operator.
- c** In the Value field, enter the value.

For example, you are a service provider who has the ABC Company as a client. Create the ABC Sales group with a filter that allows this group only to see ABC's reports. To do this, you would create a filter similar to the following:

```
cust_id = 05
```

`cust_id` is the table column name.

`=` is the numerical operator.

`05` is ABC's customer identification number. You would have created this value when you installed a reportpack. See your report pack documentation for details.

- d** Click **Add Filter**.

7 Click **Finish**.

The Contents of Groups page lists the new group; the group also appears in the navigation frame under the Groups folder.

Alternatively, there is a command that creates groups (**groupimport**), and a command that you can use to manage groups (**groupctl**). Refer to the *Performance Insight Reference Guide* for more information about these commands.

Configuring a Distributed Installation

This section describes the tasks that you may want to perform after you install OVPI to configure a distributed installation. In a distributed installation, the OVPI components are installed on multiple physically separate systems and operate as a single, distributed system. Refer to the *Performance Insight Installation Guide* for more information about distributed installations.

Configuring Polling Policies for Remote Pollers

A *remote poller* is a system that collects data from network nodes and stores the data in an OVPI database. This database is not present on the remote poller system. The remote poller collects data from the nodes in the network and then stores the data on a server running the OVPI database.

After you install a remote poller, you must configure the polling policies for your report packs. You must set up the polling policies, which are located on the database server, to identify the remote poller from which they will collect data. The database server will transfer the data from the remote poller after the collection is completed for that period.

This section describes the following ways in which you might modify polling policies in for a distributed configuration:

- Change the system which will execute the polling policy.
- Create new policies for multiple pollers.

Changing the Polling System

When you install report packs on a Satellite Server, the polling policies for the report packs are initially set up to poll from the system on which they are installed. Thus, you need to edit the polling policies for the report packs to poll from the Remote Poller.

The following procedure shows how to change the polling system for a polling policy from the Satellite Server (referred to as `neast1`) to the Remote Poller (referred to as `neast2`):

- 1 Start the Management Console, and then click the **Polling Policies** icon in the navigation pane.
- 2 Select the **All** page in the viewing pane to display all of the current polling policies.
- 3 Locate the policy you want to change, and double-click it. The Edit Polling Policy dialog box opens. The value for the Polling Assigned to box will be `neast1`.
- 4 To add a remote poller, do the following:
 - a Click **Create** next to the Polling Assigned to box. The Create Datapipe Installation dialog box opens.
 - b Type the name of the system (in this example, `neast2`) that will be the remote poller in the Hostname box.
 - c Click **OK**.
- 5 To change the polling system, select `neast2` from the Polling Assigned to list.
- 6 Click **OK**.
- 7 Repeat [Step 3](#) and [Step 5](#) for every polling policy associated with each report pack.

See [Chapter 7, Managing Polling Policies](#) for more details about polling policies.

Creating New Policies for Multiple Pollers

When you have multiple pollers polling different devices, you will need new create new policies for the pollers. For example, suppose you installed the Common Properties and Interface Reporting report packs on your Satellite Server and you had two pollers, Poller A and Poller B. You want Poller A to poll 10 devices and all of its interfaces and Poller B to poll the other 10 devices and all of its interfaces.



Before you create the polling policies, you must create polling groups for the two sets of devices.

Grouping Your Devices

To create the two groups necessary for this scenario:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#).

- 2 Click the **Groups** icon.

The Group Management window appears in the Management Console. The left side of the window displays the group management tree, which lists all the groups for the type of view that is selected.

You create a group by selecting one or more managed objects.

- 3 Expand the **Grouping** folder, and then expand the **Managed Objects** folder.

- 4 Right-click the folder that contains the managed objects that you want to group. A shortcut menu opens, which lets you create a new instance of any of the following group types: enumerated, derived, property, and folder.

For the example scenario, you need to create two groups, one for 10 devices and their interfaces and another for 10 devices and their interfaces.


- 5 Right-click the folder, and select **Create Enumerated** from the menu. The Create Group dialog box opens

- 6 Type the name for the group in the Group Name box (for example, DevicesA).

- 7 Type a description of the group in the Description box.

- 8 To add nodes to the group, do either of the following:

- Double-click a node in the Other Nodes box to move it to the Nodes in Group box.

- Select a node or an interface in the Other Nodes box, and click  to add the node to the Nodes in Group box.

- 9 To remove any node from the group before you create it, do either of the following:

- Double-click a node in the Nodes in Group box to move it to the Other Nodes box.

- Select a node in the Nodes in Group box, and then click  to return the node to the Other Nodes box.

- 10 Click **OK** when you are finished.
- 11 Repeat [Step 5](#) through [Step 10](#) of this procedure to create the second group (for example, DevicesB).

For more information about Group Manager, see [Chapter 9, Grouping Managed Objects](#).

Creating the First Polling Policy

This procedure shows you how to create the first polling policy that you will need. You will modify an existing polling policy to create the new polling policy.

To modify a polling policy:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#).
- 2 Click the **Polling Policies** icon.
- 3 Select the **All** page in the viewing pane to display all of the current polling policies.

You are going to create the two polling policies you need by modifying an existing one and then using it to create another one.

- 4 Double-click the ifAlias_Collection policy. The Edit Polling Policy dialog box opens.
- 5 Change the values for the following options:
 - a Type a distinct and unique name (Policy1 in this example scenario) for the polling policy in the Policy Name box.
 - b Select **Custom Groups** from the Collect Data From list.
 - c Select the name of the first group (**DevicesA** in the example scenario) from the Select Group to Poll From list.
 - d Select the name of the poller that will collect the data from the devices (Poller A in the example scenario) from the Polling Assigned to list.
 - e Type a description for this polling policy in the Description box.
 - f Click **OK**.

You have created the first polling policy called Policy 1.

Creating the Second Polling Policy

This procedure shows how to create the second polling policy by modifying the policy that you just created in the procedure described in [Creating the First Polling Policy on page 44](#).

To create the second polling policy:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#).
- 2 Click the **Polling Policies** icon.
- 3 Select the **All** page in the viewing pane to display all of the current polling policies.
- 4 Locate the name of the first policy that you created, and double-click it.
The Edit Polling Policy dialog box opens.
- 5 Change the values for the following options:
 - a Type a distinct and unique name (Policy2 in this example scenario) for the polling policy in the Policy Name box.
 - b Select the name of the second group (**DevicesB** in the example scenario) from the Select Group to Poll From list.
 - c Select the name of the poller that will collect the data from the devices (Poller B in the example scenario) from the Polling Assigned to list.
 - d Type a description for this polling policy in the Description box.
 - e Click **OK**.

You have created the second polling policy called Policy 2.

Maintaining Your OVPI Installation

This chapter describes the administration tasks that you should perform at specific intervals to maintain your OpenView Performance Insight (OVPI) installation.

Post-Installation Administration Tasks

After you successfully install OVPI, you should perform some subsequent administration tasks beginning the day after you install OVPI. Following are the administration tasks that you should perform the day after your initial installation of OVPI:

- Verify that OVPI is running.
- Schedule reports.
- Modifying users, groups, and group filters.

Verifying that OVPI is Running

To verify that OVPI is running properly, do the following:

- 1 Review the database's error log file for possible database corruption or possible hardware problems.

The Sybase error log file (`error.log`) is located in the directory in which you installed Sybase in the `install` folder.

Oracle's log files are usually found in `$ORACLE_BASE/admin/<SID>`.

- 2 Verify that data is being collected for the provisioned nodes by reviewing the `trend.log` file, making sure that there are no SNMP session errors and nodes that continuously cause SNMP timeout errors.

Locating the Log File

The `trend.log` file contains messages for all of OVPI's back-end processes (for example, `trendtimer`). Each message specifies the start and end times for the logged process.

If the default log file location is used when OVPI is installed, the log files are located in the `log` folder under the OVPI installation directory. (For example, `C:\OVPI\log` for a Windows system.)

To determine where the log files reside if they are not in the default location, do one of the following depending on your operating system:

Windows

- a On the Windows taskbar, click the Start button, point to Settings, and then select **Control Panel**. The Control Panel window opens.
- b Double-click **System**. The System Properties dialog box opens.
- c Click the **Advanced** tab, and then click the **Environment Variables** button.
- d Find the `TREND_LOG` variable in the System variables list. For example, it will appear similar to the following:

```
TREND_LOG          C:\OVPI\log
```

UNIX

As the operating system user (`trendadm`), type the following at the command line:

```
echo $TREND_LOG
```

This command will show the location of the log folder.

Resolving SNMP Errors

If you find SNMP errors, you should modify the **mw_collect** entries in the `trendtimer.sched` file. In particular, you should adjust the values supplied to the **mw_collect** command using the `-o` (timeout) and `-r` (retries) options. Refer to the *Performance Insight Reference Guide* for more information about the **mw_collect** command.

- 3 Use Table Manager, which is available from the Management Console, to check to see if you are collecting data. See [Viewing Property Tables on page 225](#) for details.

If you suspect your reports are not being populated correctly, you can place the Web Access Server in debug mode to view the SQL statements that are being used to generate the report. (For information about how to put the Web Access Server in debug mode, see [Changing the Web Access Server Settings on page 329](#).) When you are done reviewing the SQL statements, make sure that you turn off debug mode for the Web Access Server; otherwise, your system performance may degrade.

Scheduling Reports



Do **not** schedule all of the reports for all of your report packs.


It is recommended that you only schedule a few reports at a time and only in certain situations. For example, use a schedule when your manager wants a summary report in PDF format generated weekly. Use the Web Access Server to view reports interactively without having to schedule them. For more information about scheduling reports, refer to the *Performance Insight Guide to Building and Viewing Reports*.

Modifying User Accounts

You may need to occasionally modify a user account. For example, you might want to lock a user out of an account.

To modify a user account:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand `User Accounts`, and then click **Users**. The User Administration page opens, which lists the user accounts.

- 3 Do one of the following:
 - Click the user account on the User Administration page.
 - Click the **Edit** icon () next to the user account that you want to modify.

The *Edit Username* page appears.

- 4 Make your changes, and then click **Apply**.


For example, to lock a user account, select the **Do not allow user to login** check box.

- 5 Click **OK** to confirm your changes.

For more information about modifying user accounts, see [Maintaining User Accounts on page 295](#).

Modifying a User Group

To modify a user group:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Do one of the following:
 - Click the user group on the Contents of Groups page.
 - Click  (the Edit icon) next to the user group. If you choose this option, go to step 5 of this procedure.


The Contents of Group_Name page opens, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Group**. The Modify Group page appears, which is identical to the Add New Group page.
- 5 Select the user or user group from the Available Users/Groups list that you want to add to the user group, and then click **Add**. The user or user group is added to the Members list.

- 6 Select the user or the user group from the Members list that you want to remove from the user group, and then click **Remove**. The user or user group is removed from the Available Users/Groups list.
- 7 Click **Next**, and then click **Finish**.

Modifying a User Group Filter

To modify a user group filter:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand **User Accounts**, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Do one of the following:
 - Click the user group on the Contents of Groups page.
 - Click  (the Edit icon) next to the user group.


The Contents of Group_Name page appears, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Filter**. The Group Filters page opens.
- 5 Select the filter, and click **Edit**. The values for the filter appear in the Column, Operator, and Value boxes.
- 6 Edit the values you want, and then click **Add Filter**.
- 7 Click **Finish**.

Deleting a User Group Filter

To delete a user group filter:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand **User Accounts**, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Do one of the following:
 - Select the user group in the navigation frame.

- Click the user group on the Contents of Groups page.
- Click  (the Edit icon) next to the user group.

The Contents of Group_Name page appears, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Filter**. The Group Filters page opens.
- 5 In the Group Filters box, select the filter you want to delete, and then click **Delete**.
- 6 Click **Finish**.

Daily Administration Tasks

You should check the status of data by making sure processes and summarizations are completing properly.

Checking the Status of Data

You can check the status of processes and summarization by viewing the contents of the `audit.log`, `trend.log`, and `metrics.log` and by checking the operating systems process list. For more information about the information contained in the log files, see [Chapter 15, Interpreting Log Files](#).

Weekly Administration Tasks

Following is an administration task that you should perform on a weekly basis:

- Back up the database.

The more often you back up, the less potential there is for data loss in case of a fatal error or disk failure.

Backing Up Your Oracle Database

If you have an Oracle database, you can do either of the following:

- Offline backup. To perform this backup, you must shut down Oracle and OVPI.

To shut down OVPI, you will need to stop the OVPI HTTP Server and OVPI Timer services. See [Starting and Stopping OVPI Processes on page 69](#) for details.

- Online backup. To perform this backup, the database is available while you back up files.

For details about either of these types of backups, refer to your Oracle documentation.

Backing Up Your Sybase Database

Use the following guidelines to back up your Sybase database:

- If you can, dump everything to disk and then move disk contents to tape. If you have several systems, you might dump all databases to the backup server on one system and put all the files together on the same tape.
- Ensure that a backup server is running, and add that backup server to the interfaces file of all systems that access it. (If you create a new backup server, restart the Sybase SQL Server after you have created it.)
- The backup server program allows you to dump and load your database. This program runs on the same system as the Sybase SQL Server. You can perform backups over the network, using one backup server on a remote system and another on the local system.
- On UNIX systems, the backup server is automatically created during the Sybase SQL Server installation (`SYB_BACKUP`); however, if you want to create a backup server with a different name, use the `srvbuild` utility. See [Creating a Backup Server on a UNIX System on page 54](#) for details.



On HP-UX, you must install the Sybase EBF to correct a `srvbuild` problem **before** you run the `srvbuild` utility. The Sybase EBF is provided on the separate EBF CD in your 5.0 product media distribution.

- On Windows systems, the backup server is automatically created during the Sybase SQL Server installation; however, if you need to create another backup server with a different name, use the Server Config utility. See [Creating a Backup Server on a Windows System](#) on page 57 for details.
- Back up all the databases (`master`, `dpipe_db`, and `sybssystemprocs`) on the existing server, separately. See [Backing Up the Sybase SQL Server](#) on page 58 for details.

Backing Up a Sybase SQL Installation Server

Dumps and loads of the Sybase database are performed by the backup server program, which runs on the same system as the Sybase SQL Server. You can perform backups over the network, using one backup server on a remote system and another on the local system.

This section describe how to create a back up server on a Sybase 11.9.2.x system and backup up the Sybase SQL Server installation. For additional information about using the backup server, refer to the *SYBASE SQL Server System Administration Guide*.

Creating a Backup Server on a UNIX System

When you install Sybase, a backup server named `SYB_BACKUP` is created by default. If you want to create a backup server with a different name, use the `srvbuild` utility.

To create a new backup server:

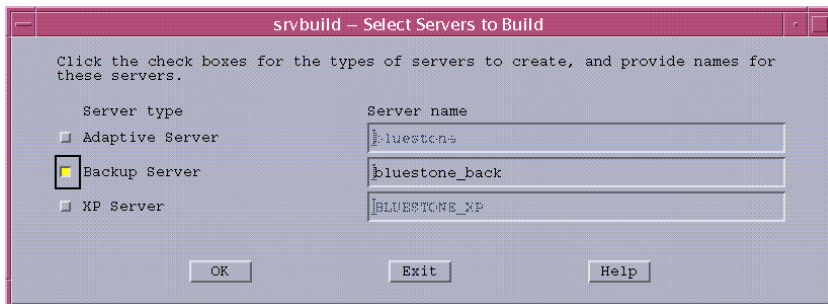
- 1 Log on to the system as `root`.
- 2 Open a terminal window.
- 3 Ensure that the `$SYBASE` variable points to the directory in which Sybase is installed.
- 4 Change to the directory in which Sybase is installed by entering the following command:

```
cd $SYBASE/bin
```

- 5 Run the `srvbuild` utility by entering the following command:

```
./srvbuild
```

The Select Servers to Build dialog box opens.

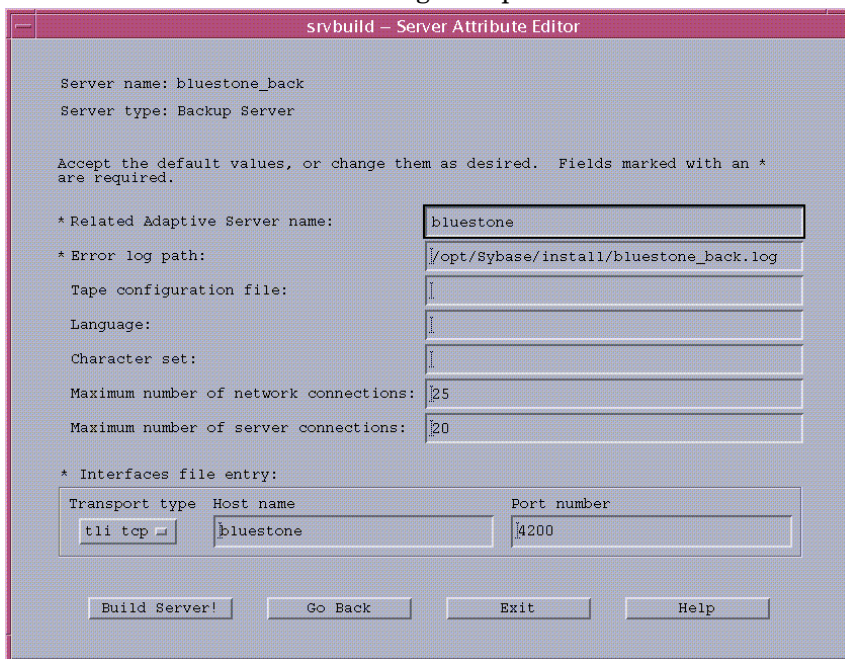


- 6 Under Server type, select **Backup Server**.

The default name of the backup server appears in the corresponding Server name box.

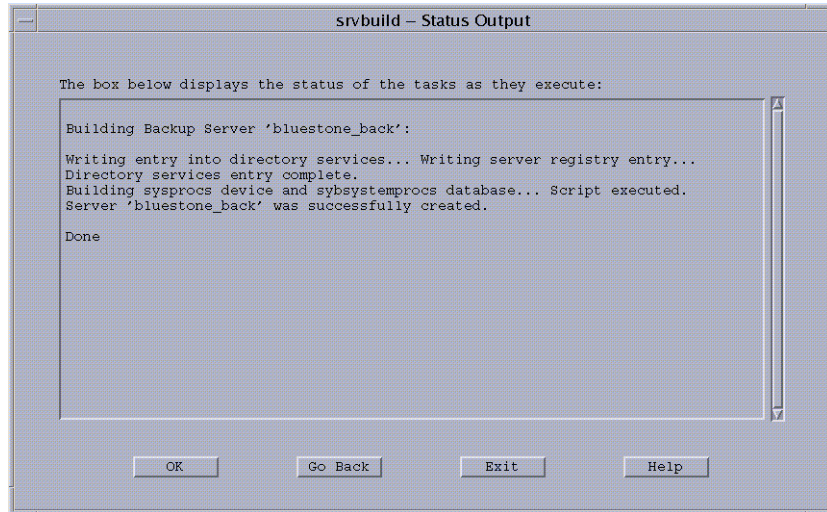
- 7 To accept the default name, click **OK**.

The Server Attribute Editor dialog box opens.



- 8 To accept the default values, click **Build Server!**.

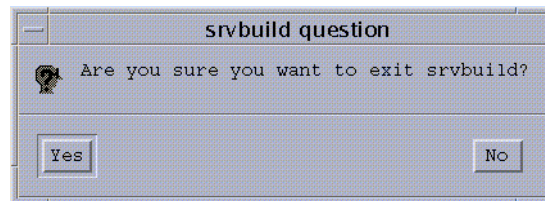
The Status Output window opens.



When the word `Done` appears in the Status Output window, the build is complete.

- 9 Click **Exit**.

The `srvbuild` question dialog box opens.



- 10 Click **Yes**.

After you create the new backup server, you need to drop the entry for the old backup server from the `sysservers` table and add an entry for the new backup server.

- 11 Start the `isql` utility by entering the following command:

```
isql -Usa -P
```

- 12 Drop the entry for the old backup server, which has the name `"SYB_BACKUP"` in the following example, by entering the following command:


```
1> sp_dropserver "SYB_BACKUP"
2> go
```

```
Server dropped.
(return status = 0)
```

- 13** Add an entry for the new backup server, which is `bluestone_back` in the following example, by entering the following command:

```
1> sp_addserver "SYB_BACKUP", null, "bluestone_back"
2> go
```

```
Adding server 'SYB_BACKUP', physical name 'bluestone_back'
Server added.
(return status = 0)
```

Creating a Backup Server on a Windows System

When you install Sybase, a backup server is created by default. If you want to create a backup server with a different name, use the Sybase Server Config utility.

To create a new backup server:

- 1 Log on to the system using an account with Administrator privileges.
- 2 On the taskbar, click the Start button, point to Programs, point to Sybase, and choose **Server Config**. The **Configure Sybase Servers** dialog box opens.
- 3 From the Products column, select **Backup Server**.
- 4 Click **Create Backup Server**. The Backup Server dialog box opens.
- 5 Click **Continue** to accept the default. The Backup Server Options dialog box opens.
- 6 Click **Continue** to accept the default values; however, if you need to add a connection, do the following:
 - a Click **Network Addresses** to display the Network Connections dialog box.
 - b Click **Add** to display the Enter Connection dialog box.
 - c Enter the appropriate information, and then click **OK**.
 - d Click **OK** in the Network Connections dialog box.
- 7 Click **Continue**. The New Backup Server Complete window opens and displays a **FINISHED!** message.

- 8 Click **Continue**.
- 9 Click **Exit** in the Configure Sybase Servers dialog box.

Backing Up the Sybase SQL Server

To back up the Sybase SQL Server, do the following:

- Start the backup server (see [Starting the Backup Server on a UNIX System on page 58](#) or [Starting the Backup Server on a Windows System on page 59](#)).
- Create a dump device (see [Creating a Dump Device on page 59](#)).
- Dump the database (see [Dumping the Database on page 61](#)).
- Load the database (see [Loading a Database on page 63](#)).
- Shut down the backup server (see [Shutting Down the Backup Server on page 64](#)).

You can back up the server manually, as described in this section, or you can create a script that will start the backup server, create a dump device, dump the database, and shut down the backup server. For more information about creating a backup script, see [Backing Up the Sybase SQL Server Using a Script on page 64](#) for details.

Starting the Backup Server on a UNIX System

To start a backup server:

- 1 Log on to the system as the Sybase Administrator.
- 2 Open a terminal window.
- 3 Ensure that the \$SYBASE variable points to the directory in which Sybase is installed.
- 4 Enter the following command to start the backup server:

```
startserver -f path/server_name
```

In this instance, *path* is the fully qualified SYBASE path to the location of the backup server; *server* is the name of the backup server. The SQL Server's default *server_name* is SYB_BACKUP.

Starting the Backup Server on a Windows System

To start a backup server:

- 1 Log on to the system using an account with Administrator privileges.
- 2 On the taskbar, click the Start button, point to Settings, and then choose **Control Panel**.
- 3 Double-click the **Administrative Tools** icon.
- 4 Double-click the **Services** icon.
- 5 From the Service list, select the Sybase backup server (Sybase BCKserver_ *hostname* _BS).

In this instance, *hostname* is the name of the system that is running the backup server.

- 6 Click **Start**.

A message appears, indicating that the system is attempting to start the backup server. After the server starts, *Started* appears next to the Sybase backup server in the Status column.

- 7 Close the Services and the Control Panel windows.

Creating a Dump Device

To create a dump device, use the **sp_addumpdevice** command as follows:

```
sp_addumpdevice "devicetype", devicename, "physicalname", size:
```

where: *devicetype* is the type of device you are dumping to enclosed in quotes. The valid values are **tape** and **disk**.

devicename is the name of the device.

physicalname is the physical name of the device, enclosed in quotes.

size specifies the size of the device.

To create a dump device to tape, the tape device must be a non-rewinding device, and you must use the size parameter with the **sp_addumpdevice** command.

To create a dump device to a disk, omit the *size* parameter with the **sp_addumpdevice** command; otherwise, a warning message will appear. The system will still create the device, however.

Alternatively, you can use Sybase Central to create a dump device to disk for Windows. To start Sybase Central, on the Windows taskbar, click the Start button, point to Programs, point to Sybase, and then choose **Sybase Central**.

Example: Creating a Dump Device to Tape

The examples in this section show how to create a dump device to tape for UNIX and Windows systems, respectively. Both examples specify a database size of 200 MB; the actual value should be just larger than your actual database size.

- *UNIX:*

```
isql -Usa -P
1> sp_addumpdevice "tape", tapedump, "/dev/rmt/0mn", 200
2> go
```

- *Windows:*

```
isql -Usa -P
1> sp_addumpdevice "tape", tapedump, "\\.\tape0", 200
2> go
```

Example: Creating a Dump Device to Disk

The following examples show how to create a dump device to disk for UNIX and Windows systems, respectively:

- *UNIX:*

```
isql -Usa -P
1> sp_addumpdevice "disk", diskdump,
"/export/home/diskdump.dat"
2> go
```

- *Windows:*

```
isql -Usa -P
1> sp_addumpdevice "disk", diskdump, "C:\HOME\diskdump.dat"
2> go
```

Dumping the Database

Use the **dump database** command to back up your databases. If you encounter any problems during the upgrade, you may need these dumps to recover.

The **dump database** command has two different formats, which let it:

- Dump data to a previously created dump device.
- Dump data to a file for which you supply the fully qualified path name in quotes.

► Ensure that you dump all the databases (`master`, `dpipe_db`, and `sybssystemprocs`) on the existing server, separately.

Dumping the Database to a Device

You can dump a database to a previously created dump device. (To create a dump device, see [Creating a Dump Device on page 59](#).)

Use the following commands to dump the database to a dump device on a UNIX or Windows system. The device appears in the location specified in the **SYBASE** (`%SYBASE%` on Windows) directory.

```
isql -Usa -P
1> dump database db_name to devicename with init
2> go
```

In this instance, use `master`, `dpipe_db`, and `sybssystemprocs` for `db_name` to dump each database.

► It is recommended that you use the `init` parameter, which causes the system to overwrite any existing data on the destination dump device (`devicename`).

The following examples show the commands you would type to dump a database to a predefined dump device (tape or disk):

- *Tape:*

```
isql -Usa -P
1> dump database dpipe_db to tapedump with init
2> go
```

- *Disk:*

```
1> dump database dpipe_db to diskdump with init
2> go
```

Dumping the Database to a File

You can dump a database to a file. Ensure that you include the quotes around the fully qualified file name, which is the directory path with the dumpfile name.

Do *one* of the following:

- *UNIX:*

```
isql -Usa -P
1> dump database db_name to "/path/dumpfile"
2> go
```

- *Windows:*

```
isql -Usa -P
1> dump database db_name to "\path\dumpfile"
2> go
```

The following examples show a command stream that backs up the **master**, **sybssystemprocs**, and **dpipe_db** databases on UNIX and Windows systems.

- *UNIX:*

```
isql -Usa -P
1> dump database master to
"/apps/syb_dumps/hostA_master_dump.dat"
2> go
```

```
1> dump database sybssystemprocs to
"/apps/syb_dumps/hostA_procs_dump.dat"
2> go
```

```
1> dump database dpipe_db to
"/apps/syb_dumps/hostA_dpipe_db_dump.dat"
2> go
```

- *Windows:*

```
isql -Usa -P
```

```
1> dump database master to "c:\sybdump\mastdmp.dat"
```

```
2> go
```

```
1> dump database sybsystemprocs to "c:\sybdump\procdmp.dat"
```

```
2> go
```

```
1> dump database dpipe_db to "c:\sybdump\dpipedmp.dat"
```

```
2> go
```

Loading a Database

This section provides examples that show how to load and restore a database from a predefined dump device (tape or disk).

The following examples shows how to load a database from tape or disk:

- *Tape:*

```
isql -Usa -P
```

```
1> load database dpipe_db from tapedump
```

```
2> go
```

In this instance, *tapedump* is a tape device.

- *Disk:*

```
1> load database dpipe_db from diskdump
```

```
2> go
```

In this instance, *diskdump* is a dump device.

The following example shows how to restore the database:

```
1> online database dpipe_db
```

```
2> go
```

Shutting Down the Backup Server

To shut down a backup server:

- 1 Log on to a Sybase SQL Server account with System Administrator privileges by entering the following command:

```
isql -Usa -Ppassword
```

- 2 Shut down the specified backup server by entering the following commands:

```
1> shutdown SYB_BACKUP  
2> go
```

When you use the **shutdown** command to stop a backup server, specify the **svrname** listed in the **syservers** table for the backup server. The SQL Server's default **svrname** is **SYB_BACKUP**.

As an alternative, you can do the following to shut down a backup server on a Windows system:

- 1 Log on to the system using an account with Administrator privileges.
- 2 On the taskbar, click the **Start** button, point to **Settings**, and then choose **Control Panel**.
- 3 Double-click the **Administrative Tools** icon.
- 4 Double-click the **Services** icon.
- 5 From the Service list, select the Sybase backup server (**Sybase BCKserver_hostname_BS**).
- 6 Click the **Stop** button.

Backing Up the Sybase SQL Server Using a Script

You can create a script to back up the database. You can schedule the script to run regularly using a scheduler program such as **cron** for UNIX or the **AT** command for Windows. (For more information about loading a database from a backup, see [Loading a Database on page 63](#).)



The commands and examples in this section apply to UNIX systems only.

Backup Script Prerequisites

Before you schedule the script, you must make the following decisions:

- 1 Decide if you want to back up smaller databases, such as **master** or **sybserverprocs**, to a filesystem rather than a tape device.

Perform either [Step a](#) or [Step b](#) for each database, depending on how you will define fixed disk files on a UNIX system:

- a If you are defining fixed disk files using the **sp_addumpdevice** command, do the following:
 - Add a new **setenv** line, which defines the variable name and the existing Sybase dump device that utilizes a disk file. For example:

```
setenv MASTER_DEVICE masterdump
```

- Change **\$(TAPE_DEVICE)** on the respective dump command to be the new variable name. For example:

```
dump database master to ${MASTER_DEVICE} file=
"master.dmp" with init
```

- b If you are defining disk files at run-time, do the following:
 - Change **\$(TAPE_DEVICE)** on the respective **dump** command to be the full path name. For example:

```
dump database master to '/Full_Pathname/OVPI_Backups/
master.dmp' with init
```

If you want to have multiple backup copies available at any instance, you can create these files with a date-designation by using the following syntax:

```
dump database master to '/Full_Pathname/OVPI_Backups/
master_`date +%y%m%d%H%M`.dmp' with init
```

- Remove the segment starting with **file=** and ending with the double quote (") from the respective **dump** command.



If you modify the **dump** command for the master database, it is important to keep the **with init** parameter.

- 2 Ensure that the proper access permissions exist on the **\$SYBASE/install** directory for the script to create the **error.log**.

- 3 Review the script's log file, **\$SYBASE/install/error.log**, on a regular basis to identify any error situations that may occur for a specific database.
- 4 Before you run the backup script, ensure the backup server exists and is running. See [Starting the Backup Server on a UNIX System on page 58](#) for details.

▶ If all the databases result in a tape volume change, the system will be waiting.

Creating the Backup Script

Create the backup script using a text editor.

▶ Replace the parameters identified by angle brackets (< >) in the following script with the requested information that appears in the angle brackets. Ensure that the arguments for the **bin/isql** line are all on the same line.

The following statements represent the contents of a script for a UNIX system:

```
#!/bin/csh -f
#
# script to backup (i.e. dump) OVPI and Sybase databases.
#

setenv DPIPE_HOME <OVPI-Directory>
setenv TAPE_DEVICE <Sybase-Name-of-Tape-Device>

# Setup correct OVPI environment.
#
source $DPIPE_HOME/data/Cshrc

# Log data to Sybase log file.
#
echo -n "Starting Sybase/OVPI backup on " | tee -a $SYBASE/install/
error.log
date | tee -a $SYBASE/install/error.log
```

```

# Start isql session to perform the backups
#
$SYBASE/bin/isql -Usa -P << EOFEOF | tee -a $SYBASE/install/
error.log
checkpoint
go
dump database master to ${TAPE_DEVICE} file="master.dmp" with init
go
dump database sybserverprocs to ${TAPE_DEVICE}
file="sybserverprocs.dmp"
go
dump database dpipe_db to ${TAPE_DEVICE} file="db_dpipe.dmp"
go
quit
EOFEOF

echo -n "Completed Sybase/OVPI backup on " | tee -a $SYBASE/install/
error.log
date | tee -a $SYBASE/install/error.log
echo " " | tee -a $SYBASE/install/error.log

exit

```

Running the Backup Script

Before you run the backup script, `BackupOVPIIDB.csh`, do the following:

- 1 Set the proper execute permissions. Typically, it is 755.
- 2 Place the script into the `$DPIPE_HOME/scripts` directory.
- 3 If you are using a scheduler program, such as **cron**, ensure that it has the environment it needs to know about `$DPIPE_HOME` and has permission to access this directory.

Infrequent Administration Tasks

Following are administration tasks that you should perform on an infrequent basis:

- Backing up the system and the database

You should back up both your OVPI installation and database when the following occur:

- After your initial OVPI installation

It is important that you perform a complete back up of your system and database after your initial installation of OVPI.

- Before and after you make any major software change to the installation

For example, if you decide to install additional report packs after your initial installation of OVPI and report packs, you should back up the system and the database **before** you install the report pack and **after**.

See [Backing Up Your System on page 68](#) and [Backing Up Your Oracle Database on page 53](#) or [Backing Up Your Sybase Database on page 53](#) for details.

- Starting and stopping OVPI processes
- Relocating the OVPI Sybase database
- Setting Up a Proxy Server

Backing Up Your System

Use the procedure you normally use to back up your file system.

Starting and Stopping OVPI Processes

OVPI has the following processes:

- OVPI HTTP Server

This process controls the OVPI client's access to the OVPI database. In order for the Web Access Server, Report Viewer, and Report Builder applications to operate properly, this process must always be running on the system.

If you installed the Web Access Server during the OVPI installation (that is, you selected the Web Access Server option), you can change the following settings for the OVPI HTTP Server process using the Web Access Server: port number, realm, log file path, logging, debugging, gateway patch features, time-out value, authentication method, and account ID. See [Changing the Web Access Server Settings on page 329](#) for details.

- OVPI Timer

This process starts other OVPI applications which need to be run at specific times or regular intervals by processing the contents of the `trendtimer.sched` file.

Refer to the *Performance Insight Reference Guide* for more information about the **trendtimer** command.

Starting and Stopping Processes on Windows

When you installed OVPI, OVPI HTTP Server and OVPI Timer are set to start automatically as services. If they do not started automatically or are stopped for some reason, you must start them manually.

Checking the Status of the Service

Before you attempt to start a service on Windows, check to see if it is running by doing the following:

To check the status of a service:


- 1 From the Windows Desktop, click the Start button, point to Settings, and select **Control Panel**. The Control Panel window opens.
- 2 Double-click **Administrative Tools**, and then double-click **Services**. The Services window opens.

- 3 Make sure the following services appear in the window with the status of Started and startup set to Automatic:

OVPI HTTP Server

OVPI Timer

For more information about stopping a service or setting up a service to start automatically, see [Starting a Service on page 70](#).

- 4 Click  in the upper right corner of the Services, Administrative Tools, and Control Panel windows to close them.

Starting a Service

Typically, in Windows, if services are installed, they automatically start each time the system is rebooted. Occasionally, you may need to manually start or stop a service.

You must start or stop the services from the system where they were installed. Using the Control Panel, you can start or stop a service. Additionally, you can also start the OVPI HTTP Server service from Windows Explorer or a command-prompt window.

Using the Control Panel

To start or stop a service using the Control Panel:

- 1 From the Windows Desktop, click the Start button, point to Settings, and select **Control Panel**. The Control Panel window opens.
- 2 Double-click **Administrative Tools**, and then double-click **Services**. The Services window opens.
- 3 Right-click the service (OVPI HTTP Server or OVPI Timer).
- 4 Select **Start** or **Stop** from the shortcut menu. The Status column indicates if the selected service was successfully started or stopped.

Starting and Stopping Processes on UNIX

When you installed OVPI, OVPI HTTP Server and OVPI Timer were set to start automatically as daemons by default; daemon services automatically start each time the system is rebooted. Occasionally, you may need to manually start or stop a service.

You must start the daemons from the system where they were installed.

To start the daemons:

- 1 Open a new shell on the system where the application is installed.
- 2 Make sure that you are `root`.
- 3 Start the OVPI HTTP Server process by entering one of the following commands:
 - HP-UX:


```
/sbin/init.d/ovpi_httpd start
```
 - Solaris:


```
/etc/init.d/ovpi_httpd start
```
- 4 Start the OVPI Timer process by entering one of the following commands:
 - HP-UX:


```
/sbin/init.d/ovpi_timer start
```
 - Solaris:


```
/etc/init.d/ovpi_timer start
```



You can manually stop the daemon by replacing `start` with `stop` in the command line.

Relocating the OVPI Sybase Database

If you need to move the OVPI Sybase database from one host to another, HP recommends creating and using Sybase backup servers on both Host A (the source host) and Host B (the target host).

To relocate the OVPI database:

- 1 Create a Sybase backup server on Host A.
 - In Windows, a backup server is automatically created during the Sybase installation.
 - In UNIX, you must use the `srvbuild` utility to install a backup server.
- 2 Ensure both the Sybase SQL Server and the backup Sybase server are running on host A.

3 Install OVPI on Host B.



The database size on Host B must be equal to or greater than the one on Host A for you to complete the remaining steps in this procedure.

4 On Host B, turn off trendtimer so that there is no automatic database activity.

5 Install a Sybase backup server on Host B. Ensure both the Sybase server and the backup Sybase server are running on Host B.

6 Turn off trendtimer on Host A.

7 Wait for any collections to complete.

8 Dump the database to an available disk file. Ensure you have the disk space available to create the dump file. For example, you can get latest backup by starting an **isql** session and entering the following command:

```
isql -Usa -P
1> dump database dpipe_db to 'path-to-diskpace'
2> go
```

9 Do one of the following:

- Mount Host A's *fullpath-to-dumpfile* file system for export to Host B.
- FTP the created file to Host B depending on whether you have available disk resources on Host B to handle the temporary storage of the dumped database file.

10 Load the backup on Host B by starting an **isql** session. For example:

```
isql -Usa -P
1> load database dpipe_db from 'fullpath-to-dumpfile'
2> go

1> online database dpipe_dp
2> go
```

11 After the backup loads, dismount Host A's *fullpath-to-dumpfile* file system on Host B or delete the *fullpath-to-dumpfile* on Host B, depending on the method you chose to get the file onto Host B.

After you have confirmed that the database on Host B is complete, you can turn on **trendtimer** on Host B.

Verifying the Relocation of the OVPI Database

After you have relocated the database, be sure to check the following items on Host B:

- Polling requests

Using Polling Policy Manager, you must modify each polling request so it will poll the new host for data. For more information about modifying which system should poll for data in a polling requests, see [Chapter 7, Managing Polling Policies](#). Once you change the system that is polled, you can verify it using the **ping** command.

- Environment variables for `ovpi` and all other users

Ensure that the following files have been changed (either by new installation or by editing) to have new settings specifically for **DISPLAY**, and **DSQUERY**:

```
$DPIPE_HOME/data/Cshrc
$DPIPE_HOME/data/Profile
```

- Database revisions

You must change the **dsquery** column in the **tl_groups** table (for all entries) to be the new Sybase environment variable value for Host B. You can do this with **isql** using the following commands:

```
isql -Udsi_dpiped -Pdsi_dpiped
1> update tl_groups set dsquery = Sybase-DSQUERY-value,
2> data_db = Sybase-DSQUERY-value
3> go
```

- Report Packs

Make sure you do the following for report packs:

- If you installed report packs with Package Manager or added reports, then copy the following files and directories from Host A to Host B:

```
$DPIPE_HOME/data/print.* files
$DPIPE_HOME/data/*.info files
$DPIPE_HOME/reports directory
$DPIPE_HOME/scripts directory
```

- Make sure that any custom report files (`.qss`, `.qgr`, or `.gos`), which can reside in any directory on Host A, are copied to Host B.

- Make sure that any custom scripts, which can reside in any directory on Host A, are copied to Host B.
- Edit the `trendtimer.sched` file on Host B to match the times and entries of the `trendtimer.sched` file on Host A, especially the lines regarding `trend_proc -f xxx.pro` for the packages and the lines regarding any custom reports or scripts.
- *Distributed installation only*: If the `trendcopy` is being used to copy data from a satellite server to a central server, the appropriate `.pro` files for each report pack must be modified to reflect the new hostname.

Setting Up a Proxy Server

To set up a proxy server, you will need to add some lines to the `httpd.conf` file. For example, if you are using an Apache proxy server, you will find the `httpd.conf` file in the `conf` subdirectory of the directory into which you installed Apache.

Add the following lines to the `httpd.conf` file:

```
ProxyPass /images http://<hostname>/images
ProxyPass /reports http://<hostname>/reports
ProxyPass /jsp http://<hostname>/jsp
ProxyPass /webview http://<hostname>/webview
ProxyPass /system http://<hostname>/system
ProxyPass /results http://<hostname>/results
ProxyPass /index.jsp http://<hostname>/index.jsp
ProxyPass /help.html http://<hostname>/help.html
ProxyPassReverse /images http://<hostname>/images
ProxyPassReverse /webview http://<hostname>/webview
ProxyPassReverse /system http://<hostname>/system
ProxyPassReverse /results http://<hostname>/results
ProxyPassReverse /index.jsp http://<hostname>/index.jsp
ProxyPassReverse /help.html http://<hostname>/help.html
ProxyPassReverse /reports http://<hostname>/reports
ProxyPassReverse /jsp http://<hostname>/jsp
```

In this instance, `<hostname>` is the name of the system that is running the Web Access Server service.

Using the Management Console

The Management Console lets you manage most of the features of your OVPI system. OVPI also provides another application for managing various features of the product called the Web Access Server. For more information about the Web Access Server, see [Chapter 12, Using the Web Access Server](#).

Using the Management Console, you can access the following applications:

- **MIB Browser.** This application lets you query or modify the status and configuration information on a node by making requests to the agent running on it. Additionally, if you want to select a different set of objects to collect from the node, you can create a custom collection table. See [Chapter 14, Managing MIBs](#) for details.
- **SNMP Discovery.** This application lets you find the nodes in your network and determine whether or not they are SNMP manageable. See [Using the Discovery Process on page 88](#) for details.
- **Type Discovery.** This application lets you determine the device type of a node (for example, router, hub, or switch). See [Using the Type Discovery Process on page 96](#) for details.
- **Package Manager.** This application lets you install and uninstall packages and deploy and undeploy reports to the Web Access Server. See [Chapter 6, Installing Report Packs](#) for details.
- **Object Manager.** This application lets you view and change certain properties of the managed objects on your system. It also enables you to import node objects into the database and create export files containing nodes. See [Chapter 8, Managing Managed Objects](#) for details.

- **Group Manager.** This application lets you create and maintain groups of managed objects referenced by the polling policies for your enterprise. See [Chapter 9, Grouping Managed Objects](#) for details.
- **Polling Policy Manager.** This application lets you set up and modify the polling policies that collect data, import and manage nodes, and manage polling groups. It also lets you update the values in a property table. See [Chapter 7, Managing Polling Policies](#) for details.
- **Table Manager.** This application lets you view all the tables in the database schema by their functional type, and examine property, collectible (raw and rate), and summary tables. See [Chapter 11, Managing Tables](#) for details.
- **System Manager.** This application lets you manage your OVPI databases, generic databases, and Web Access Servers. See [Chapter 10, Managing Your System](#) for details.

Starting the Management Console

You can use the Management Console on either a UNIX or Windows operating system. For a list of supported operating systems, refer to the *HP OpenView Performance Insight Installation Guide*.

UNIX System

To start the Management Console on a UNIX system:

- 1 Open a new shell window on the system where the application is installed.
- 2 Type the following command, and then press **ENTER**:

```
piadmin
```

Refer to the *Performance Insight Reference Guide* for more information about the **piadmin** command.

Windows System

To start the Management Console on a Windows system, do one of the following:

- On the Windows taskbar, click the Start button, point to Programs, HP OpenView, Performance Insight, and choose **Management Console** from the menu.
- From a Windows Command Prompt window, type the following command at the MS-DOS command prompt, and then press **ENTER**:

```
piadmin
```

Refer to the *Performance Insight Reference Guide* for more information about the **piadmin** command.

- Using Windows Explorer, do the following:

- a Navigate to the following directory:

```
installation_directory\bin
```

In this instance, *installation_directory* is the directory into which you installed OVPI.

- b Double-click on the following:


```
piadmin.exe
```

Logging On

To log on to the Management Console:

- 1 Start the Management Console.

The Login dialog box opens.



- 2 Enter your username, your password, and then select the system on which you want to use the application.

For more information about user accounts, see [Maintaining User Accounts](#) on page 295.

- 3 Click **Login**.

The Management Console Interface

Figure 1 shows the Management Console window

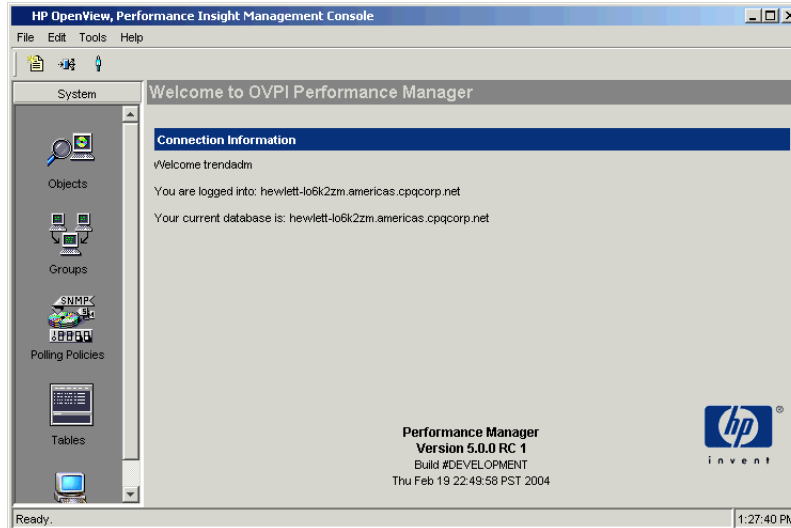


Figure 1 Management Console Window

The interface comprises a menu bar, a toolbar, a System pane, and a viewing pane.



The System pane contains the following icons, which correspond to Management Console applications:

- **Objects** (Object Manager)
- **Groups** (Group Manager)
- **Polling Policies** (Polling Policy Manager)
- **Tables** (Table Manager)
- **Systems** (System Manager)

After you open any of the applications available from the System pane or open the MIB Browser, which is available from the Tools menu, the application's window will display the following panes:

- **Navigation.** The navigation pane lets you select an item.
You expand the hierarchical levels of the navigation pane just as you do when you view the contents of a Windows system using Windows Explorer. You can expand the folders to see the additional pages or folders it may contain, or collapse those levels.
- **Viewing.** The viewing pane displays the results of the selection that you made in the navigation pane.

You can adjust the size of the panes by moving the vertical line that separates the panes in either direction. Do one of the following depending on your operating system:

- **Windows.** Place the cursor over the vertical line that separates the panes until you see a double arrow () and then drag it in either direction.
- **UNIX.** Point to the box on the vertical line that separates the panes () , left-click and hold, and then drag it either direction.

Menu Bar



The menu bar will change depending on which application is opened. See the chapter for the application you are using for a description of its menu bar.

The Management Console has the following menus: File, Edit, Tools, Window, and Help.

File Menu

Table 1 lists the options on the File menu.

Table 1 Management Console File Menu Options

Menu Option	Description
Change Server	Displays the Select Admin Server dialog box, which lets you select another system to manage. See Changing Database Servers on page 83 for more information.
Logout	Logs you off of the system, but leaves the Management Console running. You can still access the Help menu and some other basic functions. After you log off, this option changes to Logon ; select this option to log on to the system again.
Exit	Exits from the Management Console.

Edit Menu

Table 2 lists the option on the Edit menu.

Table 2 Management Console Edit Menu Options

Menu Option	Description
Options	Displays the Options dialog box, which lets you change Management Console options and options that are specific to each application.

Tools Menu

Table 3 lists the options on the Tools menu.

Table 3 Management Console Tools Menu Options

Menu Option	Description
Database Tuner	Lets you view and modify Sybase configuration parameters to improve system performance.
MIB Browser	Displays the MIB Browser application, which lets you analyze, query, and troubleshoot Structure of Management Information (SMI)-compliant SNMP MIBs.
SNMP Discovery	Displays the SNMP Discovery wizard, which lets you verify the existence of a device at a particular IP address and identify whether or not that device is SNMP manageable.
Type Discovery	Displays the Type Discovery wizard, which lets you determine the type of a discovered node.
Package Manager	Displays the Package Manager application, which lets you install and uninstall packages (report packs) and deploy reports to the Web Access Server. See Chapter 6, Installing Report Packs for details.

Window Menu

Table 4 lists the options on the Window menu. This menu appears only when the Object Manager and Group Manager applications are open. The options on this menu are active when you open multiple forms or reports.

Table 4 Management Console Window Menu Options

Menu Option	Description
Tile	Creates a tiled presentation in which the open windows are sized to fill the display the space, and the window borders touch without overlapping.
Tile Vertically	Creates a tiled presentation in which the open windows are arranged side-by-side in the display space.
Tile Horizontally	Creates a tiled presentation in which the open windows are arranged one above the other.
Cascade	Creates an overlapping presentation of windows in which the first window is presented in the upper left corner of the display and each subsequent window is presented lower and further right on the display than the previous window.

Help Menu

Table 5 lists the options on the **Help** menu.








Table 5 Management Console Help Menu Options

Menu Option	Description
Help Topics	Provides access to the OVPI online documentation. See Displaying Online Documentation on page 85 .
About	Displays version and copyright information for OVPI.

Toolbar

The buttons that appear on the toolbar vary depending on the application you are using. [Table 6](#) lists the toolbar buttons.

Table 6 Management Console Toolbar Buttons

Button	Menu Option	Description
	New	Lets you create a new item (for example, an object or a group).
	Open	Lets you open an existing group; this applies to the Group Manager application only.
	Delete	Lets you delete an item (for example, a group or a polling policy).
	Refresh	Refreshes the display.
	Change Server	Displays the Select Admin Server dialog box, which lets you select another system to manage. See Changing Database Servers on page 83 for more information.
	Views	Lets you change the way items are displayed. You can choose from small icons, large icons, a simple list, and a detailed list.
	Logout or Login	Lets you log off or to log on the Management Console.

Changing Database Servers


OVPI can read and store data on any OVPI server in your network. You can poll all servers from a single location. By changing database servers, you can log on to different servers from one location.

For information about creating, modifying, and deleting connections to other systems, see [Changing the Web Access Server Settings on page 329](#).

To change database servers:

- 1 From the File menu, choose **Change Server**.

The Select Admin Server dialog box opens, which displays the available OVPI servers.



- 2 To change servers, select a server name, and then click **Connect**.

The Management Console will now display information from the database of the server that you selected.

To verify that you are accessing the new database server, click the **Polling Policies** icon or **Tables** icon. The folder label at the top of the navigation tree identifies the current server.

Saving User Preferences

OVPI lets you maintain a unique set of user preferences for the following applications: Report Builder, Report Viewer, and the Management Console.

OVPI saves the preferences in an `.ini` file. It will maintain an `.ini` file for each user. When you modify preferences for the application, OVPI automatically saves those preferences to the appropriate `.ini` file.

For example, you can specify a directory in which OVPI should save Management Console log files, and that setting will be saved in your `.ini` file for the Management Console application.

Your `.ini` files are stored your home directory, using the following naming convention:

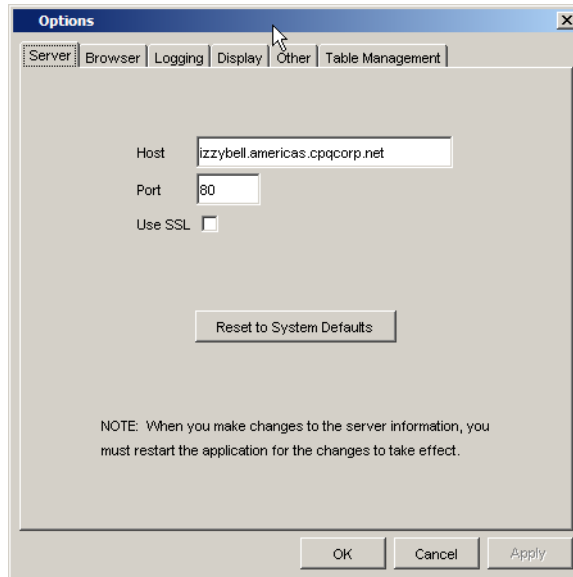
```
.ovpi/application_name/defaults.ini
```

```
.ovpi/application_name/preferences.ini
```

To save user preferences:

- 1 From the Edit menu, choose **Options**.

The Options dialog box opens.



- 2 Click a tab and change the values you want.

You will see an application-specific tab when you use any of the following Management Console applications: Object Manager, Group Manager, or Table Manager. For example, in the figure above, the Table Management tab appears.

- 3 Click **OK**.

Displaying Online Documentation

You must have Acrobat Reader installed on your system to display the online documentation that is available with OVPI. If you do not already have Acrobat Reader, you can install it from the `docs` folder, which is located in the directory into which you installed OVPI.


To display online documentation:

- 1 From the Help menu, choose **Help Topics**.

One of the following will occur:

- If you have not already chosen a browser for the client application (for example, Management Console or Report Builder) you are using, the Options dialog box opens, displaying the Browser tab. Go to [Step 2](#).
- If you have already chosen a browser for the application you are using, a browser window displays the list of online documents. Go to [Step 4](#).

2 In the Browser tab, do one of the following:

- Enter the Web browser's full path and application file name in the Browser Application box.
- Click . The Open dialog box opens. Locate the Web browser's application file name and double-click it.

For example, the full path and application file name for Microsoft Internet Explorer might be as follows:

```
C:\Program Files\Plus!\Microsoft Internet\IEXPLORE.EXE
```

3 Click **OK**.

A browser window displays a list of the available documents.



If Acrobat Reader is not installed on your system, click the **Acrobat Reader** icon to install it.

4 Click the title of the document you want to open.

The portable document format (PDF) document opens in Acrobat Reader.

Discovering the Network

This chapter describes how to retrieve information about the nodes on your network. Specifically, it describes how to use the HP OpenView Performance Insight discover process, which you perform using the following wizard-based application or commands:

- **SNMP Discovery.** Lets you find the nodes in your network and determine whether or not they are SNMP manageable. See [Using the Discovery Process on page 88](#) for details.
- **Type Discovery.** Lets you determine the device type of a node (for example, router, hub, or switch). See [Using the Type Discovery Process on page 96](#) for details.



If you installed report packs and decided not to run the type discovery process when you installed them, you must run the Type Discovery wizard after you install them.

- **The `trend_discover` command.** Lets you perform the same tasks as the Discovery (`trend_discover -l`) and Type (`trend_discover -t`) wizards; however, it also lets you control additional features of the discovery process by specifying different options.

Refer to the *Performance Insight Reference Guide* for more information about this command.

You can also use the `node_manager` command to add nodes to your system. This command lets you populate and update your various node tables. Refer to the *Performance Insight Reference Guide* for more information about this command.



The Discovery and Type wizards, the `trend_discover` command, and the `node_manager` command are only available on an OVPI server system.

Using the Discovery Process

The discovery process pings a range of IP addresses, one at a time, and then analyzes and identifies the responses it gets. It uses that information to populate or update tables in the OVPI database. The discovery process answers the following questions:

- Is there a device at a specific IP Address?
- Is the device SNMP manageable?
- If a device is SNMP manageable, what kind of device is it?

You can perform a discovery using either of the following methods:

- **Management Console.** Use the SNMP Discovery wizard.
- **Command line.** Using the `trend_discover -l` command. Refer to the *Performance Insight Reference Guide* for more information about the `trend_discover` command.

Using the SNMP Discovery Wizard

You can use the SNMP Discovery wizard to define a range of IP addresses, a Subnet mask, and one or more read community strings that you want to search. The discovery process then pings each IP address (except beginning subnet addresses [0] and broadcast addresses [255]) in that IP address range and *one* of the following occurs:

- If there is no device at the IP address, there is no response.
- If there is a device at that IP address that recognizes the ping, it responds to the discovery process. From this response, the discovery process knows that there is a device at that IP address.


The discovery process then does the following:

- Sends an SNMP GET message to the responding device.

If the device recognizes and responds to the GET, the discovery process identifies the device as SNMP manageable.

- Attempts to discover the host name of the device (if you have specified host name translation).

If your system has a protocol to translate IP addresses to host names, and the host name is found, the discovery process associates it with the discovered device. If the host name is not found, the discovery process lists the node by its IP address.

-  Ensure that Package Manager is not running before you start the SNMP Discovery wizard.

To start SNMP Discovery:

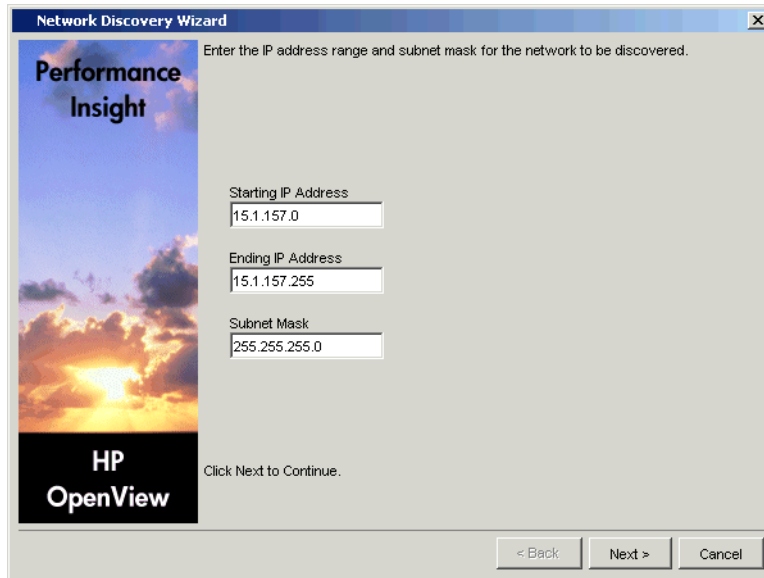
- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.



You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).

- 2 From the Tools menu, choose **SNMP Discovery**.

The following window opens, which asks you to specify a range of IP addresses to discover:



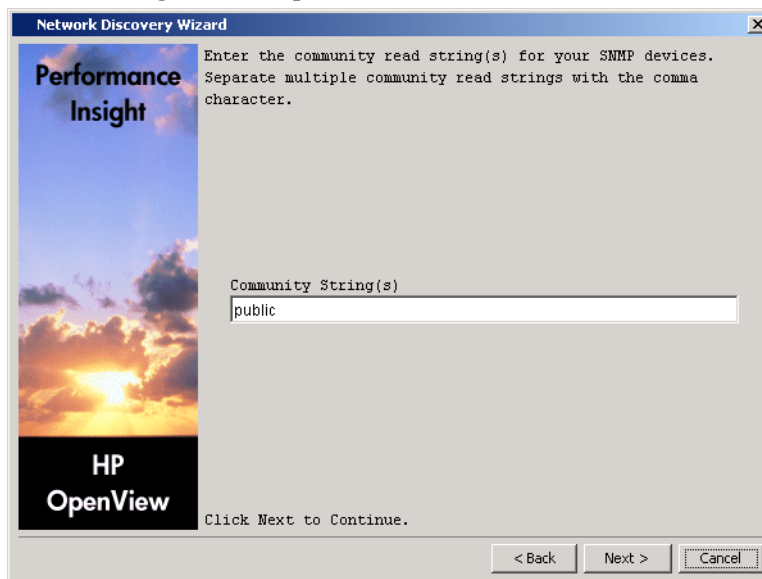
- 3** Do the following:
 - a** Specify the minimum value for the IP Address in the Starting IP Address box.
 - b** Specify the maximum value for the IP Address in the Ending IP Address box.

If you want to discover nodes that are in your local Subnet, you only have to specify addresses for the Starting IP Address box and the Ending IP Address box.

▶ The range you of addresses that you can discover depends on your subnet mask. You must ensure that the value in the Subnet Mask box is set appropriately for your network subnetting configuration.

- c** Enter a subnet mask in the Subnet Mask box.
- d** Click **Next**.

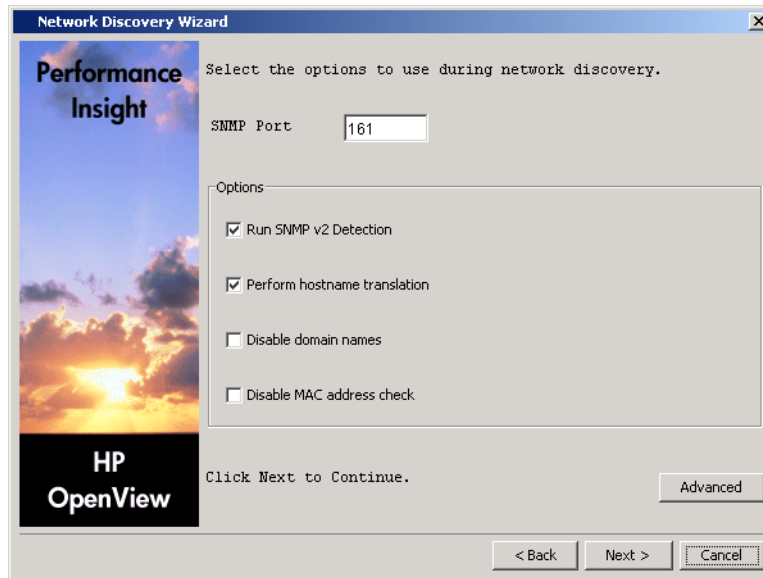
The following window opens:



When the SNMP Discovery wizard finds a device at an IP address, it sends an SNMP GET to determine whether or not that device is SNMP manageable. This SNMP GET includes a read community string. The targeted device will respond to the SNMP Discovery wizard only if its community string matches the string in the GET.

- 4 Do the following:
 - a Type the community read string or strings for your SNMP devices in the **Community String(s)** box.
Separate multiple community strings with a comma.
 - b Click **Next**.

The following window opens, which asks you to select the options that you want to use during the network discovery:



5 Do the following:

- a Enter the SNMP port in the SNMP Port box.
- b Select or click to clear the Run SNMP v2 Detection box.

When you select this option, the SNMP V2 Discovery utility will run. For more information about using the SNMP V2 Discovery utility, refer to the *Performance Insight Reference Guide*.

- c Select or click to clear the **Perform hostname translation** box.

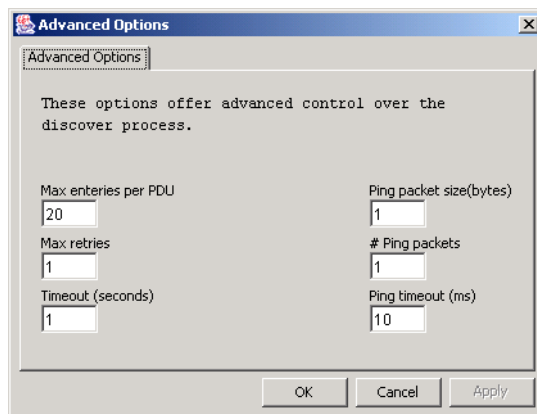
Host names are typically textual names that are used to identify devices. When you select this option, the wizard attempts to locate the host name for the IP address of each node that it finds.

Different systems use different methods to perform this translation. The wizard will try to use whichever method is active on your system. If the host name is discovered, the wizard uses it to identify the device. If it cannot find any host name for the target device, it will use the IP address.

- d Select or click to clear the Disable domain names box.

When you select this option, the discovery process will only show the hostnames of the nodes it finds; it will not show their domain names.

- e Select or click to clear the **Disable MAC address check** box.
- f Click **Advanced**. The Advanced Options dialog box opens, which lets you specify additional options for the discovery process.



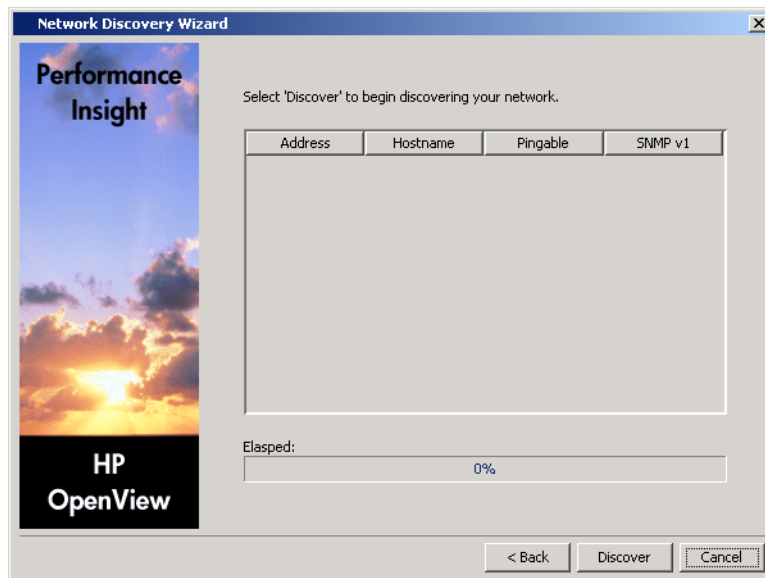
Do the following:

- Type the maximum number of entries that can be in a Protocol Data Unit (pdu) in the Max entries per PDU box.
- Type the number of times that the discovery process should retry its SNMP GET request if there is no response in the Max retries box.
- Type the number of seconds that the discovery process should wait before timing out an SNMP GET request in the Timeout (seconds) box.

- Type the byte size of the ping packet in the Ping packet (bytes) box.
- Type the number of packets that can be in a ping in the # Ping packets box.
- *Windows only.* Type the number of milliseconds that should pass before the discovery process times out a ping when the process is running in the Ping timeout (ms) box.
- Click **OK** when you are finished.

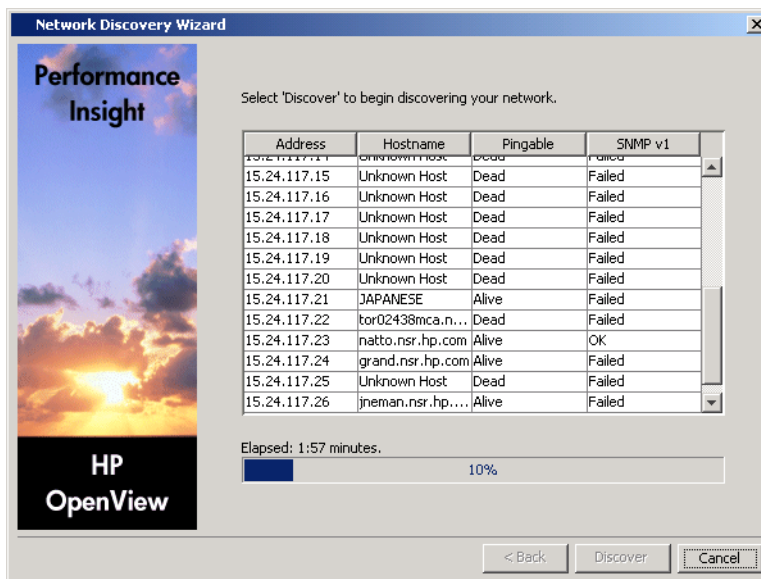
g Click **Next**.

The following window opens.




6 Click **Discover** to begin the discovery process.

As the discovery process runs, a list will appear showing each node's address, hostname, whether it is pingable, and whether it is SNMP V1; the progress bar at the bottom of the page will show the elapsed time for the process, for example:



When the discovery process is complete, the following confirmation message displays.



- 7 Click **OK**.
- 8 Click **Cancel** in the wizard window.

Using the Type Discovery Process

The Type Discovery wizard determines the device type of a discovered node and then records this information in the database. This information lets OVPI target a specific type of device for data collection.



If you installed report packs and decided not to run the type discovery process when you installed them, you must run the Type Discovery wizard after you install report packs.

You can perform a type discovery using any of the following methods:

- **Management Console.** Use the Type Discovery wizard.
- **Package Manager.** Start the type discovery process during report pack installation. See [Step 7 on page 108](#) for details.
- **Command line.** Using the `trend_discover -t` command. Refer to the *Performance Insight Reference Guide* for more information about this command.

By default, there is a `trend_discover -t` command in the `trendtimer.sched` file; it is scheduled to run once a day at 2 A.M. local time.

Starting the Type Discovery Wizard

To start the Type Discovery Wizard:

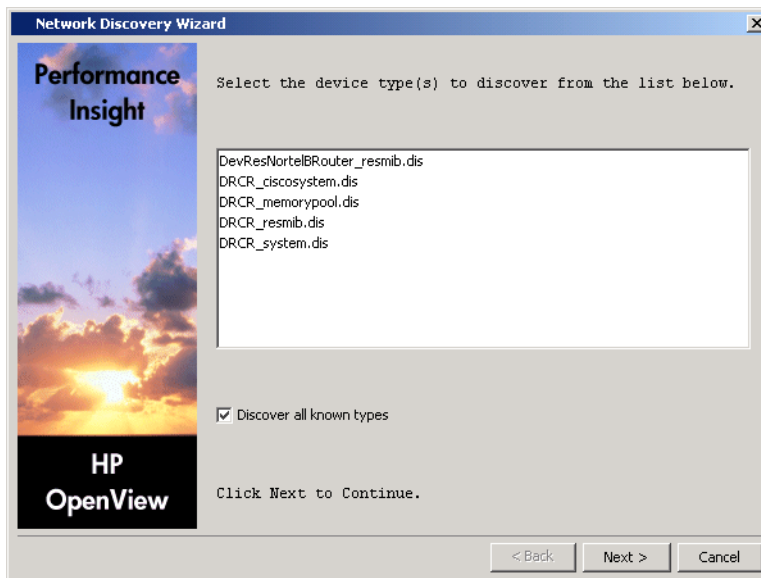
- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.



You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).

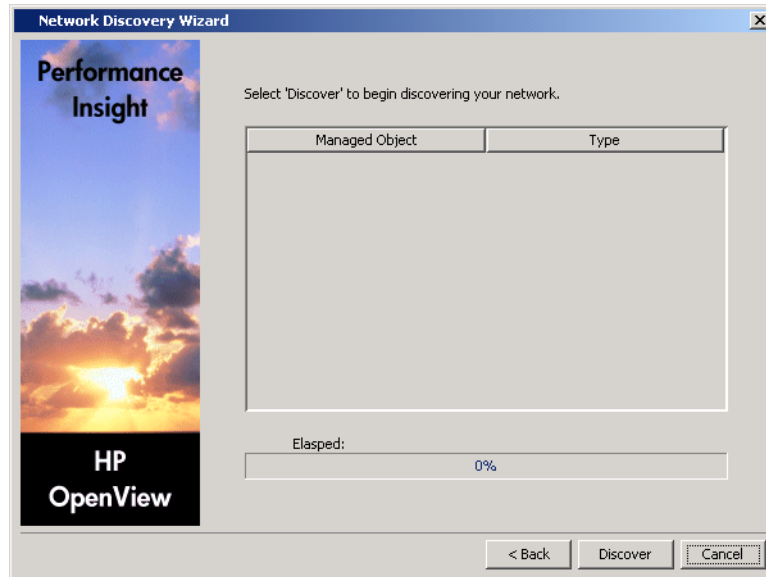
- 2 From the Tools menu, choose **Type Discovery**.

The following window opens, which asks you to specify the types of devices that you want to discover:



- 3 Do the following:
 - a Select the device type you want to discover from the list.
 - b Select the **Discover all known types** box to discover all known types of the device.
 - c Click **Next**.


The following window opens:



- 4 Click **Discover** to begin the type discovery process.

As the discovery process runs, a list will appear showing each managed object and its type; the progress bar at the bottom of the page will show the elapsed time for the process.

When the discovery process is complete, the following confirmation message displays.



- 5 Click **OK**.
- 6 Click **Cancel** in the wizard window.

Installing Report Packs

This chapter explains how to use Package Manager to do the following:

- Install and uninstall packages. (The most common packages are report packs and datapipes.)
- Deploy reports to the Web Access Server during installation and undeploy reports from the Web Access Server when packages are uninstalled.

The *Web Access Server* contains the HP OpenView Performance Insight (OVPI) Web site (the Web Access Server application) and domain.

Once you deploy reports, you can view them using the Web Access Server, Report Viewer, or Report Builder. For more information about these applications, see the *Performance Insight Guide to Building and Viewing Reports*.

Package Overview

This section discusses report packs, datapipes, and their dependencies.

Report Packs

A *report pack* is a collection of interactive reports that analyze performance. You use report packs to monitor the current performance of a particular type of device and compare performance to historic trends. That comparison lets you identify abnormal conditions and troubleshoot problem areas. Additionally, you can also compare current performance to future performance and thereby know, well in advance, when a device is likely to cross an important performance threshold.

Nearly all report packs are distributed with a demo package. If you want to familiarize yourself with the contents of the package, and you do not want to wait for the actual reports to collect data for several days or weeks, install the demo package. The demo package is complete, containing one instance of each report in the package, and the linking between tables and graphs in a demo report works just like it does in an actual report. You can also experiment with the various view options for tables and graphs by right-clicking the object and selecting a different view. The data in a demo report is static and will not change from day to day.

Datapipes

A *datapipe* creates collection tables and provides instructions for collecting data to populate the tables. It also transforms the raw data into a form that can be used by report packs.

Some, but not all, datapipes also provide custom, application-specific scripts used to massage data into a format compatible with OVPI. A datapipe allows OVPI to collect data from MIB tables and flat files and provides one or more tables for data storage.

Report Pack and Datapipe Dependencies

Some packages depend on other packages. For example, the Device Resource report pack requires the Common Property Tables package. Package Manager is aware of dependencies and will always install packages in the proper sequence whenever possible. That is, if you are installing two packages at the same time, and one package requires the other, Package Manager will know

which package to install first. Furthermore, if you are installing a package, and the prerequisite package is available to be installed, Package Manager will automatically install the required package for you.

Starting Package Manager

This section explains how to start Package Manager on Windows and UNIX systems.

Windows Desktop

To start Package Manager:

On the Windows taskbar, click **Start**, point to Programs, HP OpenView, Performance Insight, and select **Package Manager**.

The Welcome window (Figure 2 on page 103) opens.

OVPI Management Console (UNIX and Windows)

To start Package Manager, do one of the following depending on your operating system:

- Windows
 - a On the Windows taskbar, click **Start**, point to Programs, HP OpenView, Performance Insight, and select **Management Console**.
The Management Console window opens.
 - b From the Tools menu, choose **Package Manager**.
The Welcome window (Figure 2 on page 103) opens.
- UNIX
 - a Log on as `trendadm`.
 - b Type the following command:

```
piadmin
```


The Management Console window opens.

- c From the Tools menu, choose **Package Manager**.
The Welcome window ([Figure 2 on page 103](#)) opens.

Command Line (UNIX and Windows)

To start Package Manager, do one of the following depending on your operating system:

- Windows

Type the following at the command line:

```
packagemgr
```

The Welcome window ([Figure 2 on page 103](#)) opens.

- UNIX

a Log on as trendadm.

b Type the following at the command line:

```
packagemgr
```

The Welcome window ([Figure 2 on page 103](#)) opens.

Installing Packages

To install packages:

- 1 Start Package Manager. (See [Starting Package Manager on page 101](#)).

The Welcome window (Figure 2) opens.



Figure 2 Welcome Window

2 Click **Next**.

The OVPI Packages Location window (Figure 3) opens.

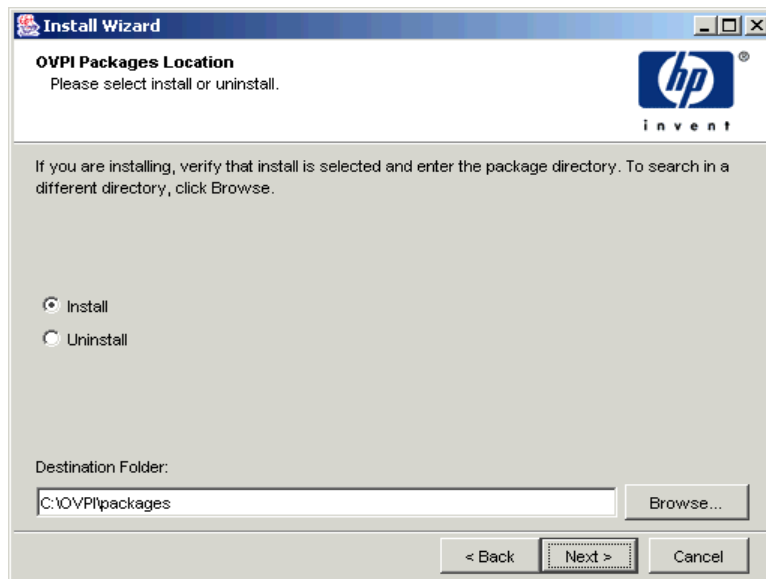


Figure 3 OVPI Packages Location Window

- 3** Do the following:
 - a** Select **Install**, and then enter the directory in which you installed the package or packages.

This should be the packages folder of the directory into which you installed OVPI. For example, on a Windows system, this directory might be C:\OVPI\packages.
 - b** Click **Next**.

The OVPI Report Deployment window (Figure 4) opens.

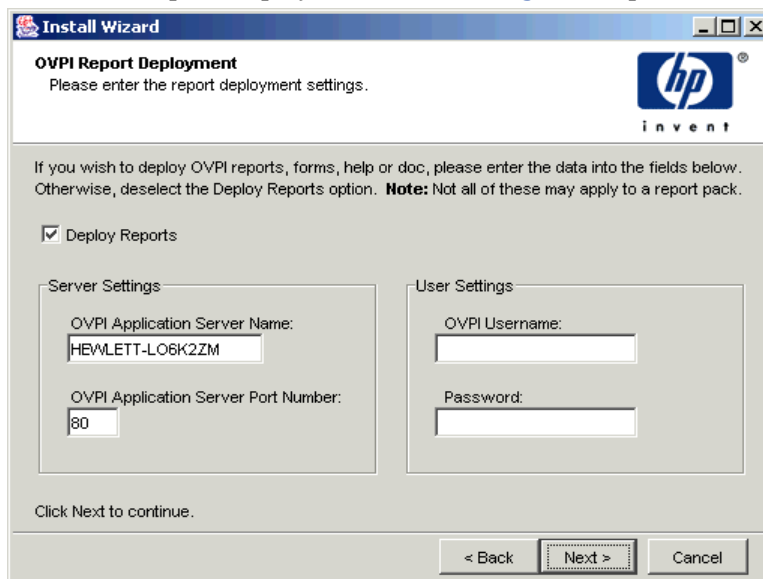


Figure 4 OVPI Report Deployment Window

This window gives you the option of deploying a package's reports to the Web Access Server. If you deploy reports, users will be able to view them using the client applications (Web Access Server, Report Viewer, and Report Builder).

You may want to collect data for the report packs but deploy them at some other time. If you decide not to deploy reports now, you can do so later using the Deployment wizard or Deployment tool. Refer to the *Performance Insight Guide to Building and Viewing Reports* for details.

- 4 Do one of the following:
 - a Click to clear the **Deploy Reports** check box if you want to install packages but do not want to deploy their reports to the Web Access Server.
 - b Select the **Deploy Reports** check box to deploy the reports to the Web Access Server, and then do the following:
 - In the OVPI Web Access Server Name box, type the name of the system (the Web Access Server system) to which you want to deploy the reports.

- In the OVPI Web Access Server Port Number box, type the port number of the system (the Web Access Server system) specified in the OVPI Web Access Server Name box.
- Enter your Administrator user name and password. You must have an account with administrator privileges to install packages.

OVPI creates a default Administrator account called `trendadm` for the Web Access Server when it is installed. Use this account or use another account that you may have created that has administrator privileges. For more information about creating accounts, see [Maintaining User Accounts on page 295](#).

5 Click **Next**.

➤ If the packages you are installing cannot be installed on the current database, a message will display, indicating that you cannot install the selected report packs on this system. For example, you will receive a message if you try to install a report pack that is only available for Sybase on a system that uses Oracle for its database.

The OVPI Package Selection window opens (Figure 5).

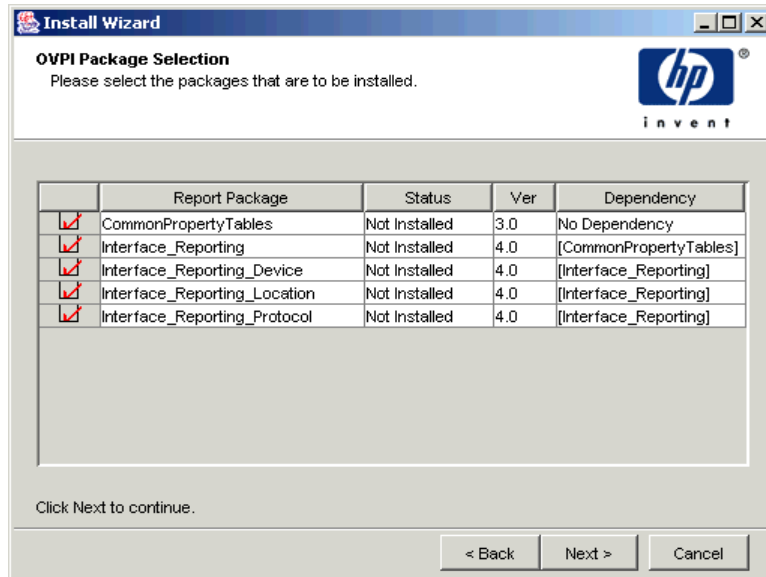


Figure 5 OVPI Package Selection Window

This window lists the following:

- Report Package. Lists the report pack packages that are currently located in the packages folder.
- Status. Lists whether or not the report pack package is installed.
- Version. Lists the version number of the report pack package.
- Dependency. List the report packs that require the selected report pack. If the dependency is `None`, this mean the report pack package does not require any other report pack packages.

6 Do the following:

- a** Select the check box next to each package that you want to install, ensuring that a red check mark (✓) appears in the box next to each package.



If you select a package that has dependencies on other packages, Package Manager automatically selects those packages for installation. A black check mark will appear in the box next to each dependent package.

- b** Click **Next**.

The OVPI Type Discover window (Figure 6) opens.

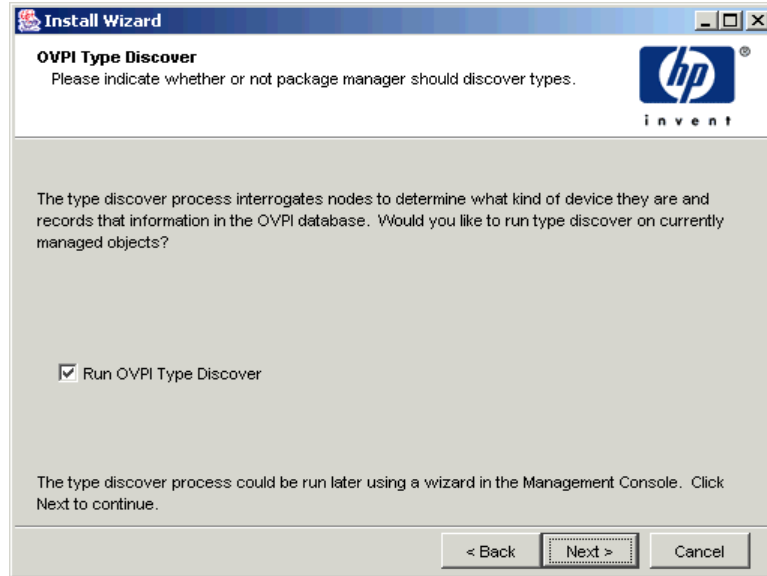


Figure 6 OVPI Type Discover Window

Type Discovery locates the type of managed objects that will be used by the packages that you install. Use this option only if you want to discover types for all the packages that you are installing. Some packages add in types that may not be of interest to you; thus, if you know which types you have in your environment, you may want to run the type discovery process after you install the packages.

7 Do *one* of the following:

- a To run Type Discovery after the package installation, click **Next**.
- b To prevent Type Discovery from running after the package installation, click to clear the **Run OVPI Discover** check box, and then click **Next**.

➤ If you want to run the type discovery process after you install the report packs, discover your network ([page 88](#)), and then run type discovery ([page 96](#)).

A Selection Summary window opens, which is similar to the one in Figure 7.

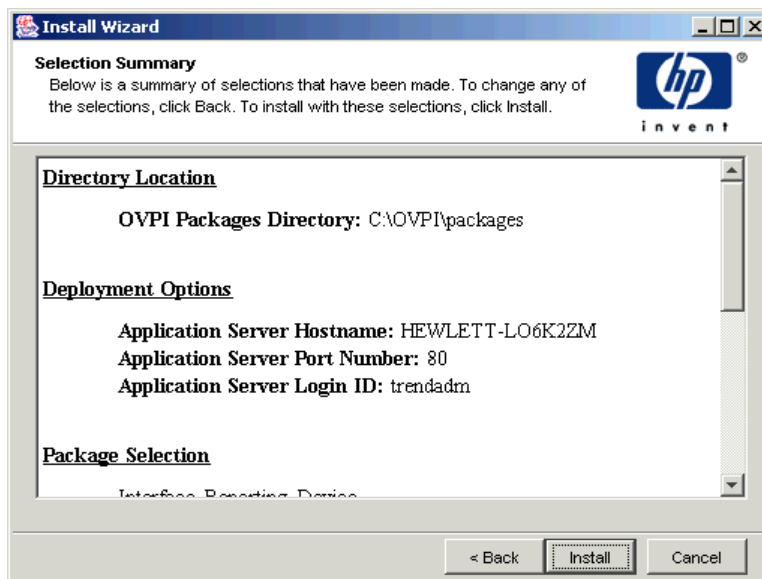


Figure 7 Selection Summary Window

- 8 Verify the information listed in the Selection Summary window, and then do one of the following:
 - If you are satisfied with the information, click **Install**.
 - If you want to change any of the information listed, click the **Back** button to return to any of the previous screens and make corrections. When finished, return to the Selection Summary window, and then click **Install**.

The Installation Progress window (Figure 8) opens, which provides information about the progress and completion of the installation.

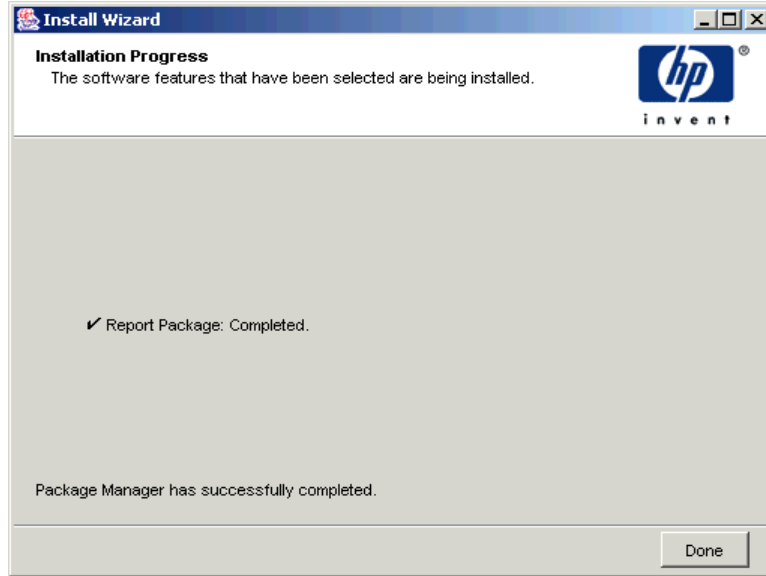


Figure 8 Installation Progress Window

If the installation of a package fails, the Report Pack Installation Error dialog box will display.

It lists the package that Package Manager could not install and asks if you want to continue with the installation. If this dialog box opens, do one the following:

- Click **Yes** to continue with the installation of the next package regardless of the error.

Select the **Don't ask me again** check box if you do not want this dialog box displayed if there is another error. (To check for any further errors after Package Manager has finished, click the **Details** button on the Installation Progress window.)

— Click **No** to end the installation.

► Each package may have several directives associated with it. A directive is a set of instructions that tell Package Manager how to install the package. If the installation of one directive fails, Package Manager flags the error by displaying the Report Pack Installation Error dialog box. If you select **Yes** to continue, Package Manager will process the rest of the directives of that report pack.

- 9 When the installation completes successfully, click **Done** on the Installation Progress window to complete the installation.

Uninstalling Packages

To uninstall packages:

- 1 Log on to the system as the OVPI Administrator.
- 2 Start Package Manager (see [Starting Package Manager on page 101](#)). The Welcome window ([Figure 2 on page 103](#)) opens.
- 3 Click **Next**.

The OVPI Packages Location window ([Figure 3 on page 104](#)) opens.

- 4 Select **Uninstall** and click **Next**.

The OVPI Report Undeployment window opens. See [Figure 3 on page 104](#).

This window gives you the option of undeploying (removing) a package's reports from the Web Access Server. If you choose to undeploy these reports, users will no longer be able to view these reports using the client applications (Web Access Server, Report Viewer, and Report Builder).

- 5 Do one of the following:
 - a Click to clear the **Undeploy Reports** check box if you do not want to remove the reports from the Web Access Server.
 - b Select the **Undeploy Reports** check box to remove reports from the Web Access Server, and then do the following:
 - In the OVPI Web Access Server Name box, type the name of the system (the Web Access Server system) from which you want to undeploy reports.

- In the OVPI Web Access Server Port Number box, type the port number of the system (the Web Access Server system) specified in the OVPI Web Access Server Name box.
- Enter your Administrator user name and password. You must have an account with administrator privileges to undeploy reports.

OVPI creates a default Administrator account called `trendadm` when it is installed. Use this account or use another account that you may have created that has administrator privileges.

6 Click **Next**.

The OVPI Package Selection window opens. See [Figure 5 on page 106](#).

7 Do the following:

- a** Select the check box next to each package that you want to uninstall, ensuring that a red check mark (✓) appears in the box next to each package.



If you select a package that has dependencies on other packages, Package Manager automatically selects those packages for removal. A black check mark will appear in the box next to each dependent package.

- b** Click **Next**.

The Selection Summary window opens, which is similar to the one in [Figure 7 on page 109](#).

8 Verify the information listed in the Selection Summary window, and then do one of the following:

- a** If you are satisfied with the information, click **Uninstall**.
- b** If you want to change any of the information listed, click the **Back** button to return to any of the previous screens and make corrections. When finished, return to the Selection Summary window, and then click **Uninstall**.

The Progress window opens ([Figure 8 on page 110](#)), showing the individual and overall progress of the uninstall.

9 Click **Done** on the Progress window when the uninstall is complete.

Managing Polling Policies

Polling Policy Manager lets you set up and modify the polling policies that collect data, import and manage nodes, and manage polling groups. It also lets you update the values in a property table. This chapter describes how to do the following:

- View current polling policies ([page 119](#)).
- Create new polling policies ([page 120](#)).
- Edit existing polling policies ([page 123](#)).
- Delete existing polling policies ([page 125](#)).
- Import nodes ([page 126](#)).
- Manage polling groups ([page 129](#)).
- Manage nodes ([page 140](#)).
- Manage community string profiles ([page 144](#)).
- Manage SNMP profiles ([page 147](#)).
- Manage Datapipe installations ([page 149](#)).
- Modify property table values ([page 152](#)).

A polling policy lets you configure data collection for maximum efficiency when you use an appropriate polling strategy. Each polling policy requires a polling group in its definition. Each polling group is a list of devices, such as nodes or instances to poll. You can define these polling groups using various

strategies to maximize efficiency and simplify administration. Polling strategies also require that you set the appropriate polling frequency for the data that you want to collect.

Polling Strategies

Your polling strategy should take into account the following:

- The most effective way to group the nodes in your network
- How frequently you need to poll each group

For example, you may wish to poll all of your routers every hour. You may want to poll all the devices at a particular data center once each day. You can create two groups that include the appropriate devices (and some devices will probably be in both groups) and define a separate polling frequency for each group.



When you install datapipes by way of report packs, the polling strategies are already set up.

Polling Groups

OVPI offers great flexibility in defining groups of nodes for polling.

Following are examples of the kinds of polling groups you can create using Polling Policy Manager:

- Group By Function

You can group all devices that perform the same function; for example, file servers.

- Group By Information Available

You may choose to create node groups according to a particular type of information that is available from all the nodes in the group.

- Group By Location

In a large network, it may make sense to subdivide groups based on their physical location. This way, target nodes are always polled by the closest polling system.

- Group By Vendor

You may have many devices that are of the same type, but they come from different vendors. You can create a separate type group for each vendor. For example, the type used for Vendor A's router may be different than the type used for Vendor B's router.

You can also create a vendor functional group that puts all devices from a particular vendor in the same group regardless of the type of device.

- Group By Priority

This can be viewed as a way to subdivide larger groups. For example, in a large network, there may be dozens or hundreds of routers. You may want to focus on the core routers for polling purposes. You could create a Core Routers group, which is a subset of the Routers Group, which is polled separately.

Remember that the polling processes, `ee_collect`, `mw_collect`, and `rmon_collect`, read the membership list for the group being polled each time they start a polling cycle. Changes made to the group are automatically reflected in the list of nodes that are actually polled for the data.



When you install datapipes using report packs, the definitions for the groups are included.

Polling Frequency

The type of data collected in a polling request, and the reason it is collected, directly affects your decision regarding how often to poll for the data. For example:

- System configuration data that rarely changes, or disk utilization data that changes gradually, may need to be collected only once a day.
- Dynamic performance statistics, with characteristics that vary with the time of day, should be collected many times each day to give an accurate picture.
- The number of minutes between polls also depends on how much the data is likely to change in the interval.

Other factors affecting polling frequency decisions include:

- Whether you are likely to miss anything if the polls are farther apart
- The load placed on both the polled and polling systems (and on the network by the data collection, especially if target nodes are at remote sites)
- The amount of database space you have for storing data

Remember that changing polling frequency from 10 to 5 minutes doubles the volume of data collected. Typically, performance statistics are collected every 10 or 15 minutes, while larger amounts of data, such as RMON host and matrix, are collected hourly.

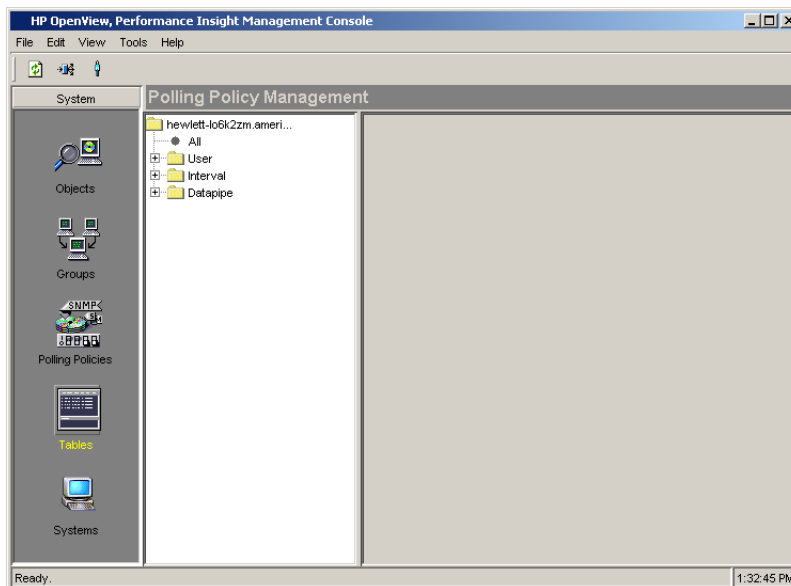
Also, remember that OVPI is not intended for reviewing real-time, current statistics to solve an immediate problem. That is the role of dashboard or real-time applications that collect data at a high frequency rate (generally measured in seconds), display the most recent data, and discard it when new data arrives or soon afterward. Such real-time applications are focused on troubleshooting, with a few key statistics collected from a few devices. OVPI polling is aimed at providing a baseline to support the troubleshooting activity and at identifying patterns generally occurring over hours, days, and weeks.

Starting Polling Policy Manager

To start Polling Policy Manager:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.
 - ▶ You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).
- 2 Click the **Polling Policies** icon in the navigation pane.

The Polling Policy Management window opens.



The Polling Policy Management window has the following panes:

- **Viewing pane.** The viewing pane is on the right side of the window.
- **Navigation pane.** The navigation pane is on the left side of the window. It shows a tree for the database that you are currently viewing. The tree contains folders; the folders contain pages, which may include other pages or additional folders.

You expand the hierarchical levels of the database just as you do when you view the contents of a Windows system using Windows Explorer. You can expand the folders to see the additional pages or folders it may contain, or collapse those levels.

You open the folders in the navigation frame's tree as follows:

- To preview the contents of a folder, click the plus sign (+).
The folder's contents is displayed in the viewing frame.
- To close an open folder, click the minus sign (-).

Viewing Functions

The Polling Policy Manager program has functions that allow you to modify the way the information appears in the window. These functions are as follows: sorting in ascending or descending order, changing the size of the columns, and refreshing a policy list.

Sorting Columns

To sort the information in the column in *ascending* order, place the cursor in the heading of the column to sort and click. Similarly, to sort the information in the column in *descending* order, place the cursor in the heading of the column to sort, hold the shift key down, and click.



When you sort in ascending order in a text field, the UPPERCASE letters appear in alphabetical order before the lowercase letters in alphabetical order. Also, the numbers in a text field appear in an alphanumeric order, which is the standard ASCII sort sequence, for example: 1, 11, 112, 3, 34, 5.

Changing the Size of a Column

To change the size of a column:

- 1 Place the cursor on the border on the right side of the column.
A double-ended arrow appears.
- 2 Click and hold the left mouse button.
- 3 Move the arrow to the right or to the left until the column has reached the desired width.
- 4 Release the left mouse button.

Refreshing a Policy List

Polling Policy Manager lets you refresh the list of polling policies in the viewing pane. When Polling Policy Manager refreshes the list, it sorts the columns in the original order, sets the sizes of the columns back to the original sizes, and removes the highlighting from any highlighted rows.

To refresh the viewing pane, do one of the following:

- Click .
- Choose **Refresh** from the View menu.

Selecting Items

You can select items in a list by doing any of the following:

- You can select a group of contiguous rows by selecting the first row and then holding the **SHIFT** key at the last row you select or holding the left mouse button down as you scroll.
- You can select a group of non-contiguous rows by pressing the **CTRL** key for each row that you select.
- In some cases, you can double-click a single row to edit that particular row.

You can combine some of these methods for a selection.

Viewing Polling Policies

The viewing pane contains the following page and folders (which contain polling policies in various categories):

- **All** page. This page provides information about all the current polling policies.
- **User** folder. This folder contains polling policies sorted by username. A page for the `trendadm` user appears in this folder; this page shows all the polling policies created by the `trendadm` user.
- **Interval** folder. This folder contains polling policies sorted by polling interval.
- **Datapipe** folder. This folder contains polling policies sorted by the datapipe name.

Viewing the Current Polling Policies

To view the current polling policies, click the **All** page.

The following polling policy information appears in the viewing pane:

Name	Datapipe	Collector	Table Alias Name	Group Name	Poll From	Server	Interval	User	Description
2003-05-15_11:13:57:310_875	dpipesnmp	MW	DevResCiscoRouter_MemoryPool	DevResCiscoRouter_Memory	POWDER2	POWDER2	15 Minutes	trendadm	
2003-05-15_11:14:00:670_317	dpipesnmp	MW	DevResCiscoRouter_CPUtotal	DevResCiscoRouter_Process	POWDER2	POWDER2	15 Minutes	trendadm	
2003-05-15_11:14:03:920_69	dpipesnmp	MW	DevResCiscoRouter_Buffer	DevResCiscoRouter_System	POWDER2	POWDER2	15 Minutes	trendadm	
2003-05-15_11:14:07:045_301	dpipesnmp	MW	DevResCiscoRouter_DeviceSystem	DevResCiscoRouter	POWDER2	POWDER2	1 Hour	trendadm	
2003-05-15_11:14:34:373_675	dpipesnmp	MW	DevResNortelBRouter_wvfResourc...	DevResNortelBRouter	POWDER2	POWDER2	15 Minutes	trendadm	
2003-05-15_11:14:37:498_108	dpipesnmp	MW	DevResNortelBRouter_DeviceSystem	DevResNortelBRouter	POWDER2	POWDER2	1 Hour	trendadm	
SysResOWPA_Config	dsi_sr	SR	SysResOWPA_Config	OWPA	powder2	POWDER2	1 Day	trendadm	The configu...
SysResOWPA_File_System	dsi_sr	SR	SysResOWPA_FileSystem	OWPA	powder2	POWDER2	1 Hour	trendadm	The file syst...
SysResOWPA_Global	dsi_sr	SR	SysResOWPA_Global	OWPA	powder2	POWDER2	1 Hour	trendadm	The global c...

The items in each column are as follows:

Name	The name of the polling policy
Datapipe	The name of the datapipe that contains the data
Collector	The name of the program that collects the data from sources outside of OVPI
Table Alias Name	The user visible name for the table that contains the data in the datapipe
Group Name	The name of the group that contains the source nodes or instances for polling the data
Poll From	The name of the machine that is doing the polling
Server	The target server name, which is the server that contains the collected data
Interval	The length of time between polling requests
User	The name of the user that created the polling policy
Description	The text that describes the polling policy

Creating Polling Policies

When you install a report pack, it provides the polling policies that it needs; however, there may be situations where you may need to create a new polling policy.

Each polling policy requires a *polling group*, which identifies where the system will locate the data to poll. Each polling group requires a list of nodes or a list of instances. Each node requires an SNMP profile and a community string profile.

When you create a polling policy, you can create each item ahead of time or you can create each item at the same time by clicking the corresponding **Create** button.

To create a polling policy, do the following steps:

- 1 From the File menu, select **Create Polling Policy**.

The Create Collection dialog box opens.

- 2 Enter a distinct and unique name for the polling policy in the **Policy Name** box.
- 3 From the **Data to Poll For** list, select the table that will contain the collected data.

The tables that appear in this list are located in the database; the name of the current database appears at the top of the navigation pane.

If you want to collect data in a table that is located on another server, select **Change Server** from the File menu to change the database. See [Changing Database Servers on page 83](#) for details. If you want to view the items that you can collect in a table, use Table Manager to view which columns are in the table.

- 4 From the **Collect Data From** list, select the type of group to collect the data from.

Following are the options on the **Collect Data From** list:

- **A Single Node.** Displays the groups that contain just one node. The system will poll only the one node in the specified group.
 - **All Nodes of the Same Type.** Displays the groups that contain nodes designated as the same type. The system will poll only the nodes in the specified group. This group designation usually identifies items that are similar devices, such as routers.
 - **All Nodes in Same View.** Displays the groups that contain nodes designated as the same view. The system will poll only the nodes in the specified group. This group designation usually identifies items that are related for a specific purpose, such as by location or by customer.
 - **A Combination of Type and View.** Displays the groups that contain nodes designated as combined types and views. The system will poll only the nodes in the specified group.
 - **Specific Instances.** Displays the groups that contain instances in a property table selected by a rule or enumerated list. The system will poll only the nodes in the specified group.
 - **Custom Groups.** Displays special groups, such as rule-based groups, defined by some other method outside the Polling Policy Manager application.
- 5** In the **Select Group to Poll From** list, select the group that contains the sources of data to poll from.

To create a new group, click **Create**. See [Managing Polling Groups on page 129](#) for details.

- 6** From the **Poll Interval** list, select the polling interval, which is the amount of time between polling requests.

Following are the options on the **Poll Interval** list:

- **Off.** Stops all polling for the data.
- **5 Minutes.** Polls for the data at 5-minute intervals starting from the zero hour.
- **10 Minutes.** Polls for the data at 10-minute intervals starting from the zero hour.
- **15 Minutes.** Polls for the data at 15-minute intervals starting from the zero hour.

- **20 Minutes.** Polls for the data at 20-minute intervals starting from the zero hour.
 - **1 Hour.** Polls for the data at the zero hour.
 - **1 Day.** Polls for the data each day at the hour specified in the `trendtimer.sched` file. The default polls for the data at 1:00 AM each day.
- 7 From the **Use Datapipe** list, select the datapipe, which specifies the method for collecting the data.
- If you want more information about datapipes, contact your HP OpenView sales representative to obtain the *Performance Insight TEEL Reference Guide*.
- 8 From the **Polling Assigned to** list, select the server that will collect the data.
- To add another collection station to the list, click **Create**. See [Managing Datapipe Installations on page 149](#) for details.
- 9 In the **Description** box, enter a text description of the polling policy.
- 10 Click **OK**.
- The polling policy is created.

Editing Polling Policies

In certain situations, you may need to modify a polling policy, for example, when you only want to poll a subset of nodes. To do this, create a view containing the subset of nodes, and then modify the polling policy to use that list of nodes instead.

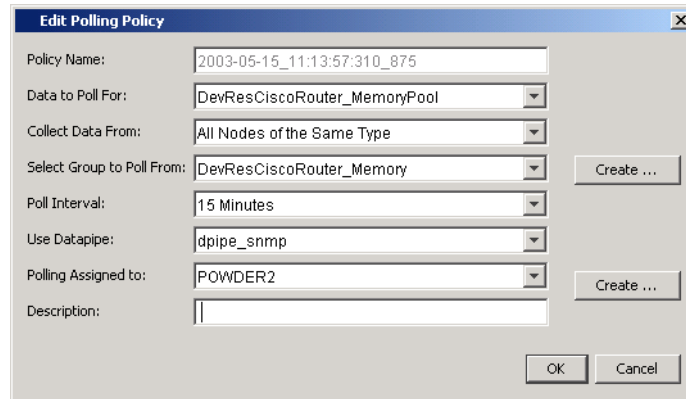


A change to an existing polling policy may prevent the installed solutions (report packs and datapipes) from working properly.

To edit a polling policy:

- 1 Do one of the following:
 - Double-click a polling policy in the viewing pane.
 - Select a polling policy in the viewing pane, and then choose **Polling Policy** from the Edit menu.

The Edit Polling Policy dialog box opens.



- 2 Do any of the following:
 - a To change the table that will contain the collected data, select another table from the **Data to Poll For** list.
 - b To change the type of group to collect the data from, select another type from the **Collect Data From** list.
 - c To change the group that contains the sources of data to poll from, select another group from the **Select Group to Poll From** list.
 To create a new group, click **Create**. See [Managing Polling Groups on page 129](#) for details.
 - d To change the polling interval, select another interval from the **Poll Interval** list.
 - e To change the datapipe that specifies the method for collecting the data, select another datapipe from the **Use Datapipe** list.

If you want more information about datapipes, contact your HP OpenView sales representative to obtain the *Performance Insight TEEL Reference Guide*.

- f To change the server that will collect the data, select another server from the **Polling Assigned to** list.

To add another collection station to the list, click **Create**. See [Managing Datapipe Installations on page 149](#) for details.

- 3 Click **OK**.

Deleting a Polling Policy

To delete one or more polling policies:

- 1 Select one or more polling policies from the polling policy list that appears in the viewing pane.

See [Selecting Items on page 119](#) for details.

- 2 From the Edit menu, choose **Delete Polling Policy**.

A deletion confirmation window displays; its appearance will differ depending on whether you selected one or several polling policies.

- 3 Do one of the following:

- If you selected only one polling policy, click **Yes** to delete the polling policy.
- If you selected multiple polling policies, do one of the following.
 - Click **Yes to All** to delete all of the selected polling policies.
 - Click **Yes** to delete the polling policy that currently is displayed in the confirmation window.
 - Click **No** to skip the polling policy that currently is displayed in the confirmation window.
 - Click **Cancel** to close the window without deleting any of the selected polling policies.

Importing Nodes

There are two steps involved when importing nodes. You must first create an ASCII file that contains the nodes you want to import, and then you must import that file.

Creating the Import File

Create an ASCII file that contains the list of nodes that you want to import. Each node that you want to import must have an entry in the file formatted as follows:

node_name, read_community, write_community, type_name, view_name, snmp_v1_flag, snmp_v2_flag, node_status, desc, snmp_profile, port_num, num_retries, timeout_sec, num_oids, get_bulk_size

Use a comma as a placeholder for any parameter that you do not include; otherwise, the node record will be skipped. [Table 7](#) describes the parameters for the import file.

Table 7 Import File Parameters

Parameter	Description
<i>node_name</i>	This is the name of the node or the IP address for the node.
<i>read_community</i>	This is the read community string for the node. If the profile does not exist, the system will create it.
<i>write_community</i>	This is the write community string for the node. If the profile does not exist, the system will create it. The default value is private.

Table 7 Import File Parameters (cont'd)

Parameter	Description
<i>type_name</i>	<p>This is the name of the type list that contains the list of nodes to poll. If you specify this field in the file, you must specify the group name. If this name does not exist, Polling Policy Manager will create a type list with the specified name.</p> <p>The group name is equivalent to the one that appears in the Group Name box of the Edit Type Group dialog box. To display this dialog box, select the Polling Groups option from the Edit menu, then select the All Nodes of Same Type option, and click Edit.</p> <p>If this parameter is blank, Polling Policy Manager will not assign the node to a type list group or remove the node from a type list group.</p>
<i>view_name</i>	<p>This is the name of the view list that contains the list of nodes to poll. If you specify this field in the file, you must specify the group name. If this name does not exist, Polling Policy manager will create a view list with the specified name.</p> <p>The view name is equivalent to the one that appears in the Group Name box of the Edit View Group dialog box. To display this dialog box, select the Polling Groups option from the Edit menu, then select the All Nodes in Same View option, and click Edit.</p> <p>If this parameter is blank, Polling Policy Manager will not assign this node to a view list group or remove this node from a view list group.</p>
<i>snmp_v1_flag</i>	<p>This specifies whether or not the node supports the SNMP protocol. Set the value to 1 (supports) or to 0 (no support). The default value is 1.</p>
<i>snmp_v2_flag</i>	<p>This specifies whether or not the node supports the SNMPv2 protocol. Set the value to 1 (supports) or to 0 (no support). The default value is 1.</p>

Table 7 Import File Parameters (cont'd)

Parameter	Description
<i>node_status</i>	This specifies whether the status of the node will be inactive or active when imported. Set the value to 1 (active) or to 0 (inactive). The default value is 1.
<i>desc</i>	This is a description of the node. If this parameter is blank, the value is blank, which appears as a space in the export file. The length can be up to 255 characters.
<i>snmp_profile</i>	<p>This is the name of the SNMP profile assigned to the node. If this profile exists, the profile is assigned to the node and parameters that follow the <i>snmp_profile</i> parameter (<i>port_num</i>, <i>num_retries</i>, <i>timeout_sec</i>, <i>num_oids</i>, <i>get_bulk_size</i>) are ignored.</p> <p>If this profile does not exist, Polling Policy Manager creates the profile using the parameters that follow the <i>snmp_profile</i> parameter, and then assigns the profile to the node.</p> <p>When the <i>snmp_profile</i> parameter is blank, the value is default, and Polling Policy Manager ignores the parameters that follow it.</p>
<i>port_num</i>	This is the SNMP port number for the new SNMP profile. The default value is 161.
<i>num_retries</i>	This is the number of retries for the new SNMP profile. The default value is 5.
<i>timeout_sec</i>	This is the number of milliseconds allowed for time out after an SNMP request. The default value is 1000.
<i>num_oids</i>	This is the number of OIDs to include in a PDU for the new SNMP profile. The default value is 20.
<i>get_bulk_size</i>	This is the size to use for a get bulk SNMPv2 request. The default value is 50.

Importing the File

To import a file containing nodes to be managed:

- 1 From the File menu, choose **Import Nodes**.
The Import from dialog box opens.
- 2 Navigate to the file that contains the nodes you want to import.
- 3 Double-click the file name, or select the name and click **Open**.

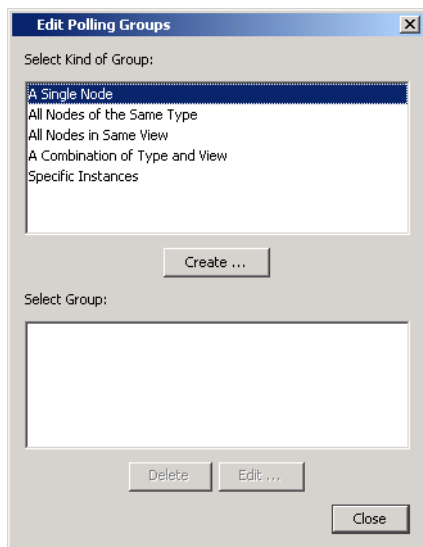
Alternatively, you can use the `node_manager` command to import nodes using the command line. For more information about the `node_manager` command, refer to the *Performance Insight Reference Guide*.

Managing Polling Groups

To create, edit, or delete a polling group:

- 1 From the Edit menu, choose **Polling Groups**.

The **Edit Polling Groups** dialog box opens.



- 2 Select the kind of group that you want to create, edit, or delete from the **Select Kind of Group** box.

If there are any corresponding groups for the selected type, they will appear in the **Select Groups** box.

The **Select Kind of Group** box contains the following options:

- **A Single Node.** Contains just one node. The system will poll only the one node in this group.
- **All Nodes of the Same Type.** Contains nodes designated as the same type. The system will poll only the nodes specified in this group. This group designation usually identifies items that are similar devices, such as routers.
- **All Nodes in Same View.** Contains nodes designated as the same view. The system will poll only the nodes specified in this group. This group designation usually identifies items that are related for a specific purpose, such as by location or by customer.
- **A Combination of Type and View.** Contains nodes designated as combined types and views. The system will poll only the nodes specified in this group.
- **Specific Instances.** Contains instances in a property table selected by a rule. The system will poll only the instances specified in this group.

- 3 To create a new group for the selected type, click **Create**.

See [Creating a Polling Group on page 131](#) for the applicable procedure for the selected group.

- 4 To edit a group for the selected type, do one of the following:
 - Double-click the group you want to edit in the **Select Group** box.
 - Select the group you want to edit in the **Select Group** box, and then click **Edit**.

For information about how to edit a polling group, see [Editing a Polling Group on page 139](#).

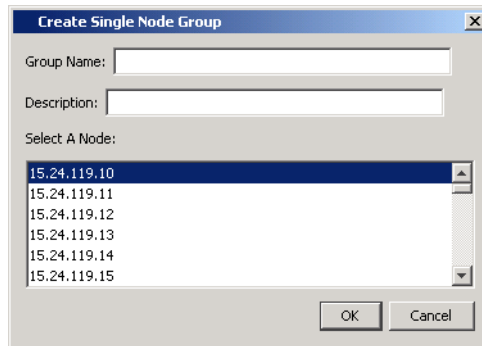
- 5 To delete a group for the selected type:
 - a Select the group you want to delete from the **Select Group** box.
 - b Click **Delete**. A deletion confirmation window displays.
 - c Click **Yes** to delete the polling group
- 6 When you are finished, click **Close**.

Creating a Polling Group

When you create a polling group, you identify the nodes or instances from which data will be collected during the polling process. You can create the following types of polling groups: a single node, nodes of the same type, nodes in the same view, a combination of types and views, and specific instances.

Creating a Single Node Group

When you select the **A Single Node** option in a dialog box and click **Create**, the Create Single Node Group dialog box opens.

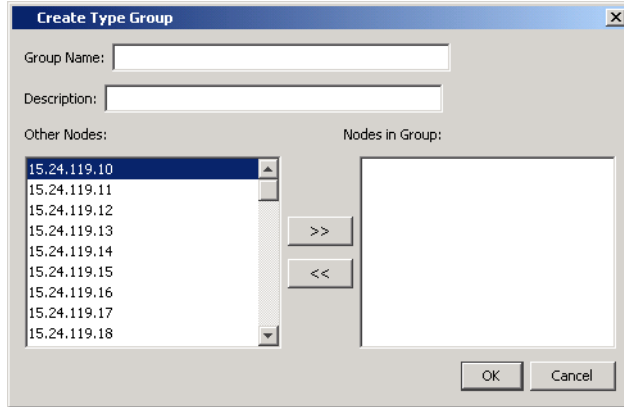


To create a single node group:



- 1 Type the name of the group in the **Group Name** box. This name can be a maximum of 30 characters.
- 2 Type a description for the group in the **Description** box. The length of the description can be 255 characters.
- 3 Select one node from the **Select A Node** box.
- 4 Click **OK**.

Creating All Nodes of the Same Type Group

When you select the **All Nodes of the Same Type** option in a dialog box and click **Create**, the Create Type Group dialog box opens.

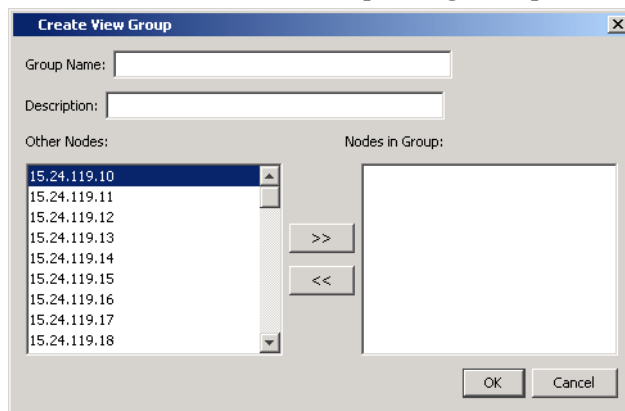


To create a group that contains nodes of the same type:



- 1 Type the name of the group in the **Group Name** box. This name can be a maximum of 30 characters.
- 2 Type a description for the group in the **Description** box. The length of the description can be 255 characters.
- 3 To select nodes for the group, do *one* of the following:
 - Select a node in the **Other Nodes** box, and then click  to add the node to the **Nodes in Group** box.
 - Double-click a node in the **Other Nodes** box; the node moves to the **Nodes in Group** box.
- 4 To remove a node from group, do one of the following:
 - Select a node in the **Nodes in Group** box, and then click  to return the node to the **Other Nodes** box.
 - Double-click a node in the **Nodes in Group** box; the node moves to the **Other Nodes** box.
- 5 When you are finished, click **OK**.

Creating All Nodes in Same View Group

When you select the **All Nodes in Same View** option in a dialog box and click **Create**, the Create View Group dialog box opens.

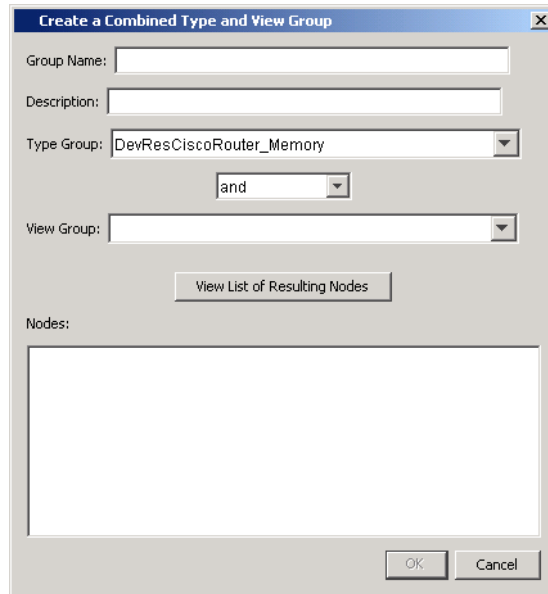


To create a group that has all nodes in the same view:

- 1 Type the name of the group in the **Group Name** box. This name can be a maximum of 30 characters.
- 2 Type a description for the group in the **Description** box. The length of the description can be 255 characters.
- 3 To select nodes for the group, do *one* of the following:
 - Select a node in the **Other Nodes** box, and then click  to add the node to the **Nodes in Group** box.
 - Double-click a node in the **Other Nodes** box; the node moves to the **Nodes in Group** box.
- 4 To remove a node from group, do one of the following:
 - Select a node in the **Nodes in Group** box, and then click  to return the node to the **Other Nodes** box.
 - Double-click a node in the **Nodes in Group** box; the node moves to the **Other Nodes** box.
- 5 When you are finished, click **OK**.

Create Combination of Type and View Group

When you select the **A Combination of Type and View** option in a dialog box and click **Create**, the Create a Combined Type and View Group dialog box opens.



The screenshot shows a dialog box titled "Create a Combined Type and View Group". It features the following elements:

- Group Name:** A text input field.
- Description:** A text input field.
- Type Group:** A dropdown menu with "DevResCiscoRouter_Memory" selected.
- Connector:** A dropdown menu with "and" selected.
- View Group:** A dropdown menu.
- View List of Resulting Nodes:** A button.
- Nodes:** A large empty rectangular area for displaying results.
- OK** and **Cancel** buttons at the bottom right.

To create a group of nodes that comprises a combination of type and view node groups:

- 1 Type the name of the group in the **Group Name** box. This name can be a maximum of 30 characters.
- 2 Type a description for the group in the **Description** box. The length of the description can be 255 characters.
- 3 Select a group of nodes from the **Type Group** list.
- 4 Select a group of nodes from the **View Group** list.

- 5 Select an operator from the **Operator** list, which is located between the **Type Group** and **View Group** lists.

The **Operator** list has the following options:

- **and**. The resulting list of nodes contains the same nodes that are in both the Type and View groups. For example, if the Type group had nodes 1, 2, 3, and the View group had nodes 3, 4, 5, then the resulting group would have node 3.
 - **or**. The resulting list of nodes contains all the nodes that are in both the Type and View groups. For example, if the Type group had nodes 1, 2, 3, and the View group had nodes 3, 4, 5, then the resulting group would have nodes 1, 2, 3, 4, 5.
 - **type not in view**. The resulting list of nodes contains the nodes that are in the Type group and not in the View group. For example, if the Type group had nodes 1, 2, 3, and the View group had nodes 3, 4, 5, then the resulting group would have nodes 1, 2.
 - **view not in type**. The resulting list of nodes contains the nodes that are in the View group and not in the Type group. For example, if the Type group had nodes 1, 2, 3, and the View group had nodes 3, 4, 5, then the resulting group would have nodes 4, 5.
- 6 Click **View List of Resulting Nodes** if you want to see the list of nodes.
 - 7 When you are finished, click **OK**.

Creating a Specific Instance Group

When you select the **Specific Instances** option in a dialog box and click **Create**, the Create a Specific Instance Group dialog box opens.

The screenshot shows a dialog box titled "Create a Specific Instance Group". It features the following elements:


- Group Name:** A text input field.
- Description:** A text input field.
- Select a Property Group:** A dropdown menu currently showing "K_Customer".
- Rule:** A text input field with an "Edit ..." button to its right.
- Buttons:** "View Resulting Instances" and "Edit Instance Values ..." are positioned below the Rule field.
- Instances:** A large, empty rectangular area labeled "Instances:".
- Footer:** A checkbox labeled "Save Selected Instances Only", and "OK" and "Cancel" buttons.

A specific instance group and a rule-based group refer to the same type of group.

To create a specific instance group:

- 1 Type the name of the group in the Group Name box. This name can be a maximum of 30 characters.
- 2 Type a description for this group in the Description box. The length of the description can be 255 characters.
- 3 Select the property table that contains the instances from the Select a Property Group list.

- 4 Click **Edit** to enter the selection criteria for the instances. The Create Rule dialog box opens.



You can define the rule by entering the selection criteria for up to two columns.

- 5 Do the following:
 - a Select the column name from the Select Columns list. These are the columns in the selected property table.
 - b Select the operator for the value in the selected column. The following operators are available:

=	equal
<	less than
>	greater than
>=	greater than or equal
<=	less than or equal
<>	not equal

like

not like
 - c Type the value for the selected column in the box to the right. This value can be an alphanumeric string.

When using the **LIKE** operator, its value must be a string and can include the following wildcard characters: underscore (**_**) and percent (**%**). The underscore character represents any single character. The percent character represents zero or more characters. For example:

12%3 would match the following: 1233, 123, and 1244443.

_ime would match the following: time, lime and dime.

- d** If you select two columns, or the same column name twice, select the joining operator for the two columns from the *Operator* list that appears between the columns.


If you select the same column name twice, you can create a compound rule that will select only a portion of that column.

- e** When you are finished, click **OK**.

After you create the rule, it appears in the **Rule** box in the Create Rule dialog box.

- 6** If you want to see all the columns for the selected rows in the property table, click **View Resulting Instances**.
- 7** If you want to modify the values in the selected rows, do the following:
 - a** Select one or more rows by highlighting them. See [Selecting Items on page 119](#) for details.
 - b** Click **Edit Instance Values** to modify the values in one or more columns of the selected rows.

The **Edit Instance** dialog box opens.



- c** Select the column to modify from the **Column** list.
- d** Type the value for the selected column in the **Column** box, which is to the right of the **Column** list.

- e Verify that you want to make the same change for the selected rows.
 - ▶ When you click **Apply**, the system updates the database immediately for the selected rows.
- f If the change is correct, click **Apply** to update all the selected rows in the property table immediately.
- 8 If you want the group to have only some of the resulting instances in it, do the following:
 - a Select one or more rows by highlighting them. See [Selecting Items on page 119](#) for details.
 - b Select the **Save Selected Instances Only** check box. Verify that the check mark appears in the box.
- 9 When you are finished, click **OK**.

Editing a Polling Group

When you edit a polling group, you can modify any value in any field in the window except the value in the **Group Name** field.

When the Edit a Combined Type and View Group dialog box opens, the list of nodes do not appear until you click the **View List of Resulting Nodes** button. If there are a lot of nodes, you may have to wait longer than you are willing just to change a description.

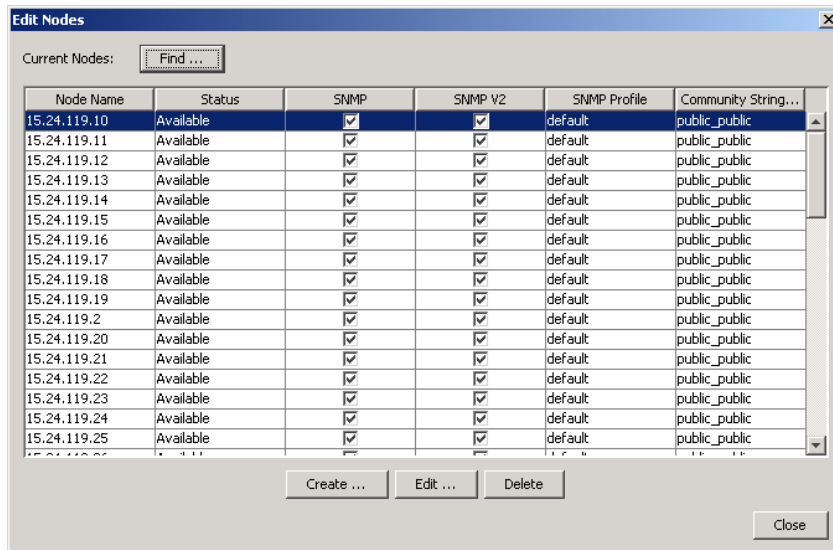
When the **Edit a Specific Instance Group** dialog box opens, the list of instances do not appear until you click the **View Resulting Instances** button unless you saved the list of instances by clicking the **Save Selected Instances Only** box and not the rule. If there are a lot of instances, you may have to wait longer than you are willing just to change a description. You must click the **Edit** button to modify the **Rule** field.

Managing Nodes

To create, edit, or delete nodes:

- 1 From the Edit menu, choose **Nodes**.

The Edit Nodes dialog box opens.



This window shows a list of the current nodes with the following fields:

- **Node Name.** This is the name of the node. It can be an IP address or an alphanumeric name.
 - **Status.** The status specifies whether the node is available for polling.
 - **SNMP.** If the box is selected, the node supports SNMP v1 data.
 - **SNMP V2.** If the box is selected, the node supports SNMP v2 data.
 - **SNMP Profile.** This is the name of the SNMP profile for the node.
 - **Community String Profiles.** This is the name of the Community String profile for the node.
- 2 To create a new node, see [Creating Nodes on page 141](#).
 - 3 To edit an existing node, see [Editing Nodes on page 143](#).
 - 4 To delete an existing node, see [Deleting Nodes on page 143](#).

- 5 When you are finished, click **Close**.


Finding a Node

To locate an existing node from the list:

- 1 From the Edit menu, choose **Nodes**.

The Edit Nodes dialog box opens.

- 2 Click **Find**. The Find dialog box opens.



- 3 Enter the name of the node that you are looking for in the **Node Name** box.
- 4 If you want to only locate complete words in the name instead of partial words, select the **Match Whole Word Only** check box.
- 5 If you want to locate names that match the case you specified in the **Node Name** box instead of any case, select the **Match Case** check box.
- 6 Click **Find** to locate and highlight the node.

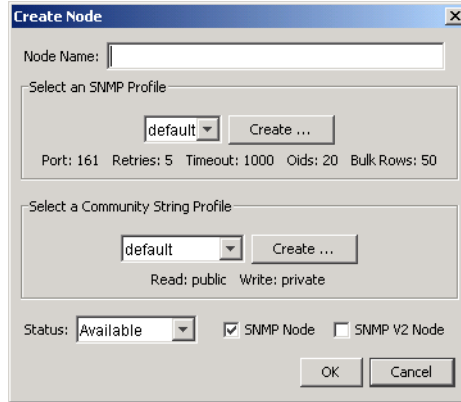
Creating Nodes

To create a node:

- 1 From the Edit menu, select **Nodes**.

The Edit Nodes dialog box opens.

- 2 Click **Create**. The Create Node dialog box opens.



- 3 Type the name for the node in the **Node Name** box. This name can be a maximum of 64 characters.
- 4 Select an SNMP profile from the **Select an SNMP Profile** list.
If the SNMP profile does not exist, click **Create** to create a new one. See [Managing SNMP Profiles on page 147](#) for details.
- 5 Select a Community String profile from the **Select a Community String Profile** list.
If the Community String profile does not exist, click **Create** to create a new one. See [Managing Community String Profiles on page 144](#) for details.
- 6 Select the status for the node from the **Status** list.
The **Status** list contains the following options:
 - **Not Available**. Use this value when you want to create the node but you do not want to poll it yet, or you want to make the node unavailable for polling at this time.
 - **Available**. Use this value to make the node available for polling.
- 7 Select the **SNMP Node** box if the node supports SNMP v1 data.
- 8 Select the **SNMP V2 Node** box if the node supports SNMP v2 data.
- 9 Click **OK**.

Editing Nodes

To edit an existing node:

- 1 From the Edit menu, choose **Nodes**.

The Edit Nodes dialog box opens.

- 2 Select a node that you want to edit by highlighting the corresponding row.

Alternatively, you can double-click a single node, and then go to [Step 4](#).

If you want to find a node in the list of existing nodes, see [Finding a Node on page 141](#) for details.

- 3 Click **Edit**.

The Edit Node dialog box opens.

- 4 Modify any value except the value in the **Node Name** box. See [Creating Nodes on page 141](#) for details.

- 5 Click **OK**.

Deleting Nodes

To delete one or more existing nodes:

- 1 From the Edit menu, choose **Nodes**.

The Edit Nodes dialog box opens.

- 2 Select one or more nodes that you want to delete by highlighting the corresponding rows. See [Selecting Items on page 119](#) for details.

If you want to find a node in the list of existing nodes, see [Finding a Node on page 141](#) for details.

- 3 Click **Delete**.

A deletion confirmation window displays. Its appearance will differ depending on whether you selected one or multiple nodes.

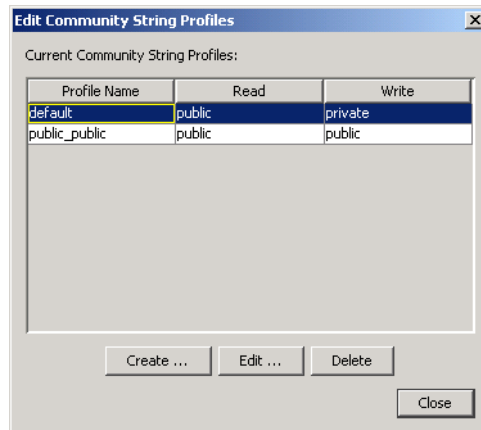
- 4 Do *one* of the following:

- If you selected only one node, click **Yes** to delete the polling policy.
- If you selected multiple nodes, do *one* of the following.
 - Click **Yes to All** to delete all of the nodes.

- Click **Yes** to delete the node that currently is displayed in the confirmation window.
- Click **No** to skip the node that currently is displayed in the confirmation window.
- Click **Cancel** to close the window without deleting any of the selected nodes.

Managing Community String Profiles

When you choose **Community String Profiles** from the Edit menu, the Edit Community String Profiles dialog box opens, which shows a list of the existing community string profiles with the name of the profile and the corresponding read and write permissions.



You can do the following:

- Create a community string profile. See [Creating a Community String Profile on page 145](#) for details.
- Edit a community string profile. See [Editing a Community String Profile on page 145](#) for details.
- Delete a community string profile. See [Deleting a Community String Profile on page 146](#) for details.

Creating a Community String Profile

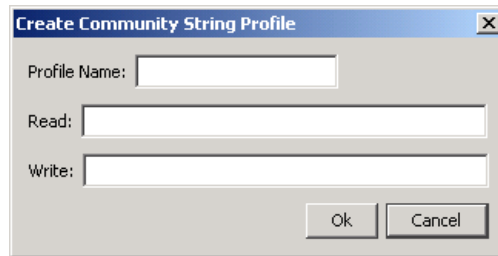
To create a community string profile:

- 1 From the Edit menu, choose **Community String Profiles**.

The Edit Community String Profiles dialog box opens

- 2 Click **Create**.

The Create Community String Profile dialog box opens.



- 3 Type the name for the profile in the **Profile Name** box. This name can be a maximum of 64 characters.
- 4 Type the value for the read permission in the **Read** box.
- 5 Type the value for the write permission in the **Write** box.
- 6 Click **OK**.

Editing a Community String Profile

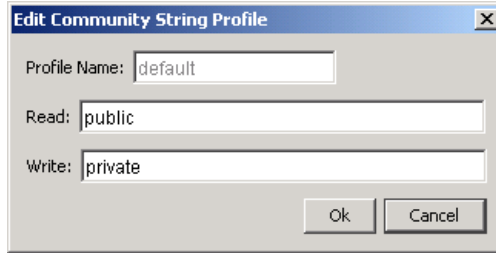
To edit a community string profile:

- 1 From the Edit menu, choose **Community String Profiles**.

The Edit Community String Profiles dialog box opens.

- 2 Select the community string profile you want to modify from the **Current Community String Profiles** box.

- 3 Click **Edit**. The Edit Community String Profile dialog box opens.



- 4 Modify the permission values in the **Read** and **Write** fields.

▶ You cannot change the name of the profile.

- 5 Click **OK**.

Deleting a Community String Profile

To delete a community string profile, do the following:

- 1 From the Edit menu, choose **Community String Profiles**.

The Edit Community String Profiles dialog box opens.

- 2 Select the community string profile you want to delete from the **Current Community String Profiles** box.

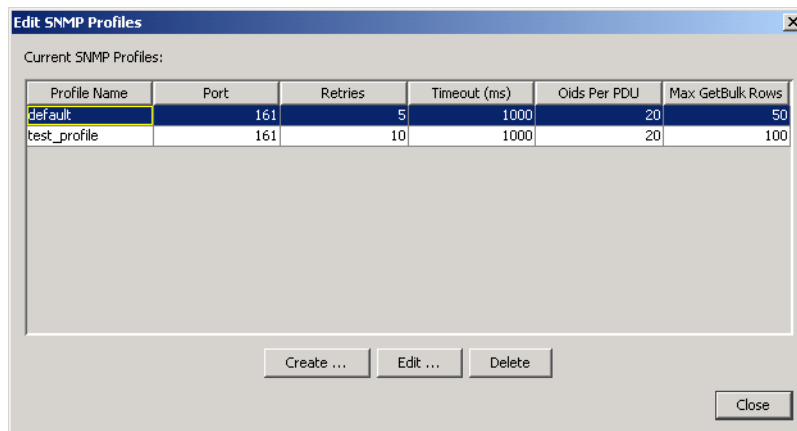
- 3 Click **Delete**.

A deletion confirmation window displays.

- 4 Click **Yes**.

Managing SNMP Profiles

You can create, edit, or delete an SNMP profile by choosing **SNMP Profiles** from the Edit menu. The Edit SNMP Profiles dialog box opens, which lists the existing SNMP profiles with names and their corresponding settings.

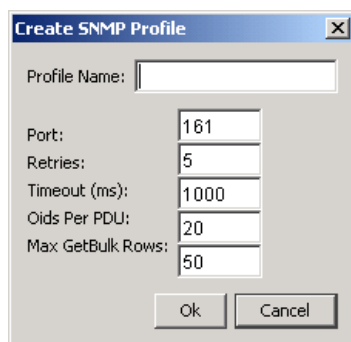


Creating an SNMP Profile

To create an SNMP profile:

- 1 From the Edit menu, choose **SNMP Profiles**. The Edit SNMP Profiles dialog box opens.
- 2 Click **Create**.

The Create SNMP Profile dialog box opens.



- 3 Type the name for the profile in the **Profile Name** box. This name can be a maximum of 64 characters.
- 4 Type the value for the port number in the **Port** box.
- 5 Type the number of retries allowed in the **Retries** box.
- 6 Type the number of milliseconds to wait for a timeout in the **Timeout (ms)** box.
- 7 Type the maximum number of OIDs that the agent will support for each PDU in the **Oids Per PDU** box.
- 8 Type the maximum number of table rows that will be retrieved in a single GetBulk request in the **Max GetBulk Rows** box. This field is supported only in SNMPv2 and SNMPv3 devices.
- 9 Click **OK**.

Editing an SNMP Profile

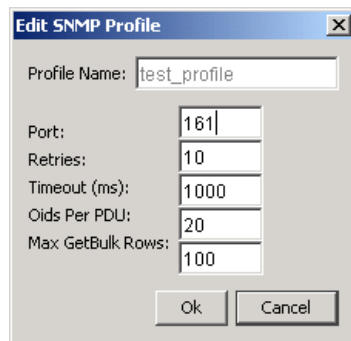
To edit an SNMP profile:

- 1 From the Edit menu, choose **SNMP Profiles**.

The Edit SNMP Profiles dialog box opens.

- 2 Select the SNMP profile you want to modify from the **Current SNMP Profiles** box.
- 3 Click **Edit**.

The Edit SNMP Profile dialog box opens.



The screenshot shows a dialog box titled "Edit SNMP Profile". It contains the following fields and values:

Profile Name:	test_profile
Port:	161
Retries:	10
Timeout (ms):	1000
Oids Per PDU:	20
Max GetBulk Rows:	100

At the bottom of the dialog box are two buttons: "Ok" and "Cancel".

- 4 Modify the values for the settings that you want to change.

You cannot change the name of the profile.

- 5 Click **OK**.

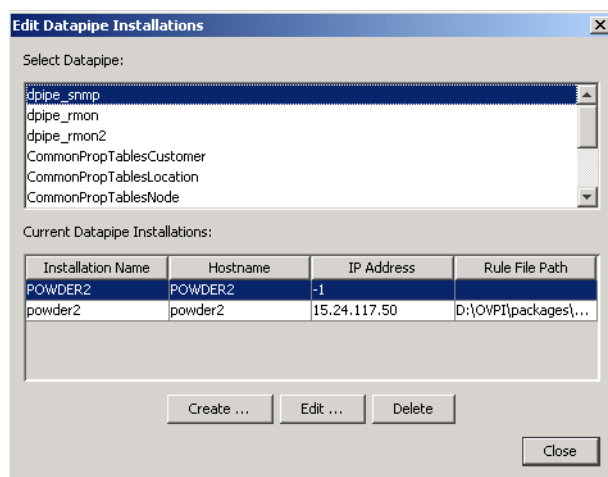
Deleting an SNMP Profile

To delete an SNMP profile:

- 1 From the Edit menu, choose **SNMP Profiles**.
The Edit SNMP Profiles dialog box opens.
- 2 Select the SNMP profile you want to delete from the **Current SNMP Profiles** box.
- 3 Click **Delete**.
A deletion confirmation window displays.
- 4 Click **Yes**.

Managing Datapipe Installations

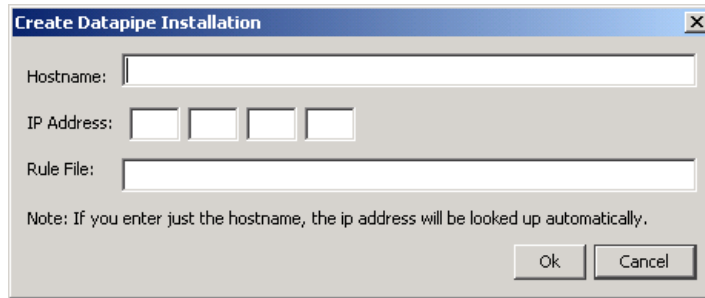
You can create, edit, or delete a datapipe installation by choosing **Datapipe Installations** from the Edit menu. The Edit Datapipe Installations dialog box will open. It lists the existing datapipes with a corresponding list of installations for each datapipe.



Creating a Datapipe Installation

To create a datapipe installation:

- 1 From the Edit menu, choose **Datapipe Installations**. The Edit Datapipe Installation dialog box opens.
- 2 Click **Create**. The Create Datapipe Installation dialog box opens.



- 3 Type the hostname for the collection station in the **Hostname** box.
If you enter only the hostname, the system will look up the IP address automatically. You may have to fully qualify the hostname for the system to look up the address properly. When you enter a hostname, it is the name of the collection station.
If the system cannot locate the IP address for the specified hostname, a window displays to confirm that it could not find the address.
- 4 Type the IP address for the collection station in the **IP Address** box.
You can leave this box blank if the system can locate the corresponding address for the hostname. If you enter only the IP address, then the address is the name of the collection station.
- 5 Type the name of the file that contains the instructions for the collection station in the **Rule File** box. This box can remain blank.
- 6 Click **OK**.

Editing a Datapipe Installation

To edit a datapipe installation:

- 1 From the Edit menu, choose **Datapipe Installations**. The Edit Datapipe Installation dialog box opens.

- 2 Select the datapipe you want to modify from **Select Datapipe** list.
- 3 Select the corresponding installation from the list in the **Current Datapipe Installations** box.
- 4 Click **Edit**. The Edit Datapipe Installation dialog box opens.

Edit Datapipe Installation

Hostname: powder2

IP Address: 15 24 117 50

Rule File: \\ResCiscoRouter\DevResCiscoRouter.ap\rDRCR_MemoryPool.teel

Note: If you enter just the hostname, the ip address will be looked up automatically.

Ok Cancel

- 5 Modify the values for the settings that you want to change.
- 6 Click **OK**.

Deleting a Datapipe Installation

To delete a datapipe installation:

- 1 From the Edit menu, choose **Datapipe Installations**. The Edit Datapipe Installation dialog box opens.
- 2 Select the datapipe to delete from the **Select Datapipe** box.
- 3 Select the corresponding installation from the **Current Datapipe Installations** box.
- 4 Click **Delete**.

A deletion confirmation window displays.

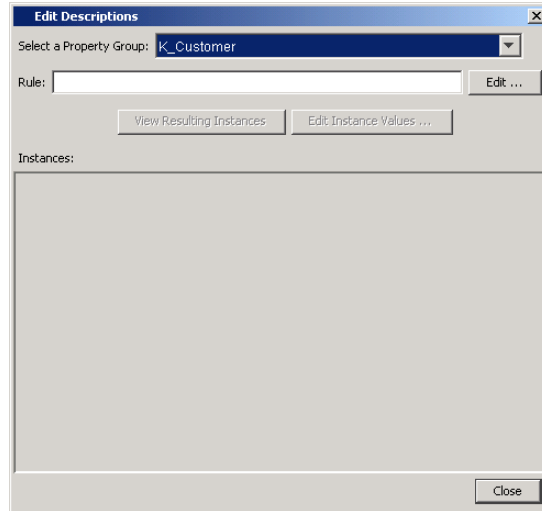
- 5 Click **Yes**.

Modifying Property Table Values


To modify any value in a property table:

- 1 From the Edit menu, choose **Descriptions**.

The Edit Descriptions dialog box opens.



- 2 Select the property table that contains the values you want to modify from the **Select a Property Group** list.
- 3 Click **Edit** to enter the selection criteria for the instances to modify. The Create Rule dialog box opens.



- 4 Define the rule by entering the selection criteria for up to two columns.
 - a Select the column name from the **Select Columns** list. These are the columns in the selected property table.

- b** Select the operator for the value in the selected column. The following operators are available.

=	equal
<	less than
>	greater than
>=	greater than or equal
<=	less than or equal
<>	not equal

like

not like

- c** Type the value for the selected column in the box on the right of the list.
- d** If you select two columns, or the same column name twice, select the joining operator for the two columns from the **Operator** list that is between the two columns.


If you select the same column name twice, you can create a compound rule that will select only a portion of that column.

- e** When you are finished, click **OK**.

After you enter the rule, it appears in the **Rule** box.

- 5** Click **View Resulting Instances** to see all the columns for the selected rows in the property table.
- 6** If you want to modify the values in the selected rows, do the following:
- a** Select one or more rows by highlighting them. You can select a group of contiguous rows or a group of non-contiguous rows by pressing the **CTRL** key for each row that you select.

- b Click **Edit Instance Values** to modify the values in one or more columns of the selected rows. The Edit Instance dialog box opens.



- c Select the column to modify from the **Column** list.
 - d Type the value for the selected column in the **Column** box.
 - e Verify that you want to make the same change for the selected rows.
- When you click **Apply**, the system updates the database immediately for the selected rows.
- f If the change is correct, click **Apply** to update all the selected rows in the property table *immediately*; otherwise, click **Close** to close the dialog box and leave the rows in the property table unchanged.
- 7 When you are finished, click **Close** in the Edit Descriptions dialog box.

Managing Managed Objects

Object Manager is an application that is available from the Management Console, which lets you view and change certain properties of the managed objects on your system. It also enables you to import node objects into the database and create export files containing nodes.



Managed objects are added to your system when you install report packs that support this feature; therefore, you must install these report packs before you can view managed objects.

Each report pack's property tables are associated with one or more properties; thus, you can view managed objects based on different properties. The available properties are as follows: Device (the default), Customer, Location, Vendor, and Cards.

In addition to viewing managed objects, you can do the following:

- **View reports.** You can view a list of the reports that contain data for the selected managed object. You can also generate and view a report for a selected object.
- **View and modify forms.** You can view a list of the forms associated with a selected managed object. You can also open forms and modify data, which updates the appropriate property tables.

A *form* is a type of report that gives you an interface to update properties in the database. For example, a form may allow you to update a customer in the database. You can also access forms that will create new objects in the database.

Viewing Managed Objects

To view managed objects:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.



You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).

- 2 Click the **Objects** icon in the navigation pane.

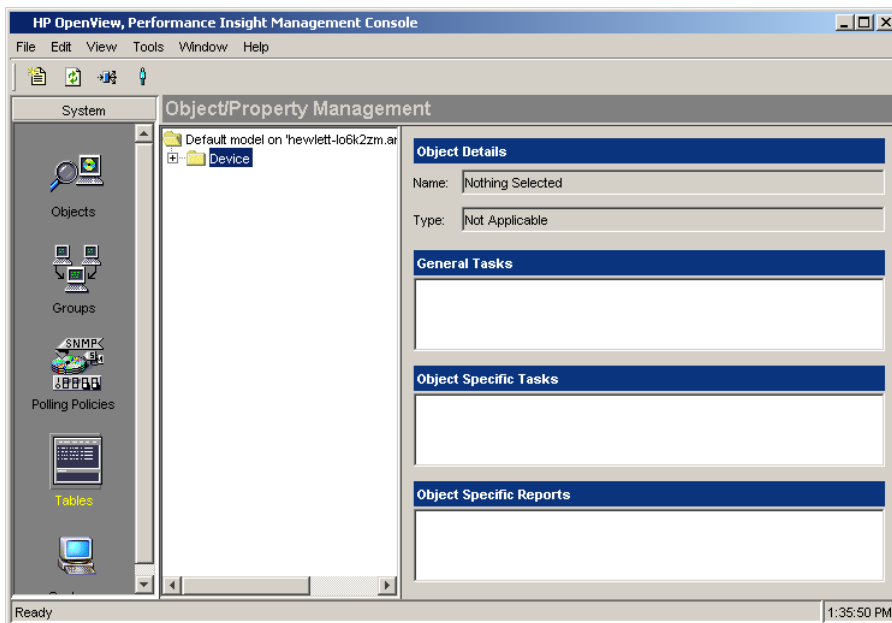
The Object/Property Management window opens in the Management Console. The left side of the window displays the Objects tree, which lists all the objects for the type of view that is selected.

In this tree, object names are arranged numerically, then alphabetically, with uppercase letters sorting before lowercase letters. Any objects that were discovered without associated names appear in the tree under the appropriate network address (for example, 15.24.115.205). These addresses are listed at the top of the tree.

By default, the first level below an object type (for example, Device) shows only one tree branch for each letter of the alphabet or for each numeral. Clicking the plus sign (+) next to the letter or numeral displays a list of objects that begin with that letter or numeral. You can achieve more differentiation in your tree by changing the Folder Uniqueness value. See [Modifying the Objects Tree Display on page 160](#).

If you know the exact name of the object you want, you can use the find feature. See [Finding a Specific Object on page 159](#).

The following figure shows the Object/Property Management window:



The following information appears on the right side of the window:

- The object's name and type appear in the **Object Details** section.
 - The forms that represent global tasks, not specific to an object, appear in the **General Tasks** section.
 - The forms that apply to the object you selected appear in the **Object Specific Tasks** section. This section displays only the forms that you use to modify data. To add new data, see [Creating New Managed Objects on page 165](#).
 - The reports that apply to the object you selected appear in the **Objects Specific Reports** section. To make any of your custom report appear in this section, you must create a parameter in the report and associate it with the table used in the report element (table or graph). Refer to Chapter 6 in the *Performance Insight Guide to Building and Viewing Reports*.
- 3 Select a managed object in the Objects tree or use the find feature. (See [Finding a Specific Object on page 159](#).)
 - 4 To view a report or form, double-click the name of the report or form.

The selected item appears in a separate window.

- If you select a report, OVPI generates and displays that report for the managed object you selected. You can view and print the report.
- If you select a task, OVPI displays the appropriate form with data filled in for the managed object you selected. You can change certain values (for example threshold values) in the form, and these values are updated in the associated property table. See [Modifying the Database on page 164](#).

Each time you select a different item in the Objects tree, the list of forms and reports is updated to reflect the object you selected.

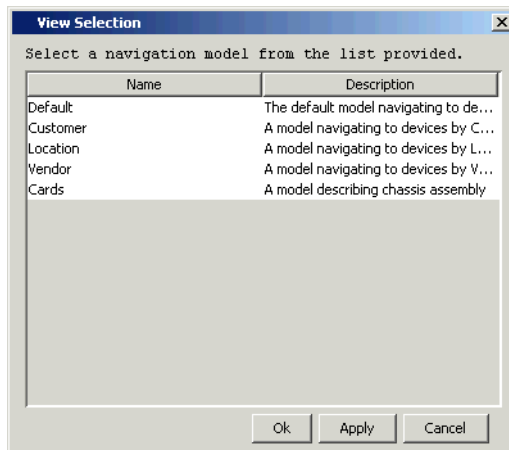
Changing the Objects View

You can change the property that is used to display managed objects. You can view objects by Device, Customer, Location, Vendor, or Cards.

To change the objects view:

- 1 From the View menu, choose **Change View**.


The View Selection dialog box opens.



- 2 Select a property, then click **OK**.

The Objects tree refreshes, displaying all the managed objects that have that property assigned to them.

For example, suppose you wanted to view all managed objects for customers; the Objects tree will appear similar to the following picture after you expand the Customer folder.




Finding a Specific Object

If you know the exact name of the object you want to display, you can use the find feature to access the information quickly.

- 1 From the Edit menu, choose **Find**.

The Find dialog box opens.



- 2 Type the object name in the text area and click **Find**.

If OVPI finds the name you typed, the corresponding object is highlighted in the Objects tree, and the associated forms and reports are listed in the respective areas.

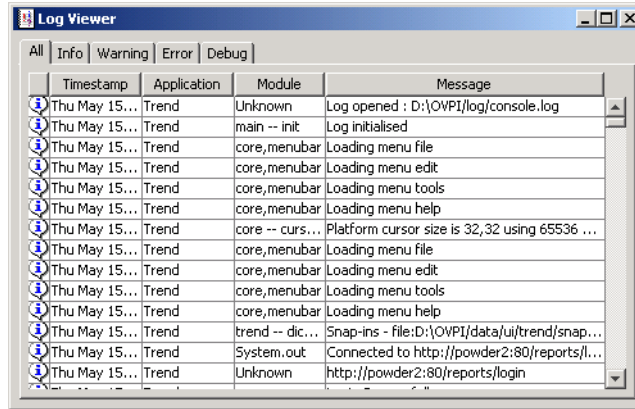
- 3 Do one of the following:
 - If OVPI did not find the object name you typed, you can try typing another name.
 - If the correct object name is highlighted in the tree, click **Close** in the Find dialog box.
- 4 When the desired object name is highlighted, go to [Step 4 on page 157](#).

Viewing Log Information

The View menu includes two options that enable you to view log file and debugging information.

To view log file information, choose **Log** from the View menu.

The Log Viewer window opens.



The Log Viewer window comprises five tabs, each displaying a different type of log information. There are four types of log messages: Info, Warning, Error, and Debug. The All tab displays all four types of messages.

To view output information, choose **Output** from the View menu.

The Output Viewer window displays the most current log messages as they are generated.

Modifying the Objects Tree Display

If you have a very large number of objects, you can set some values that help you control the length of the tree to minimize scrolling. On the other hand, if you are not concerned about scrolling, but you want to be able to quickly identify an object without expanding folders, other values help you accomplish this goal.

These options are under the Object Management tab of the Options page. The **Tree Nodes Per Level** value and the **Folder Uniqueness** value work together to determine how many folders or leaves you will see at any given level in the Objects tree.

Set **Tree Nodes Per Level** to the desired maximum number of folders/leaves to be displayed at each level of the tree. If the number of objects in the database exceeds this value, OVPI looks at the **Folder Uniqueness** field to determine how to reduce the number of displayed folders and leaves.

Because 1 is the smallest possible value for **Folder Uniqueness**, you may still see a very long list in the tree. For example, if you set **Tree Nodes Per Level** to 500 and Folder Uniqueness to 1, and you have 10,000 devices, OVPI will put all of the device names that begin with A in one folder, all the devices that begin with B in another folder, and so on. You might still have more than 500 folders at a given tree level, but OVPI has done its best to minimize the number of folders.

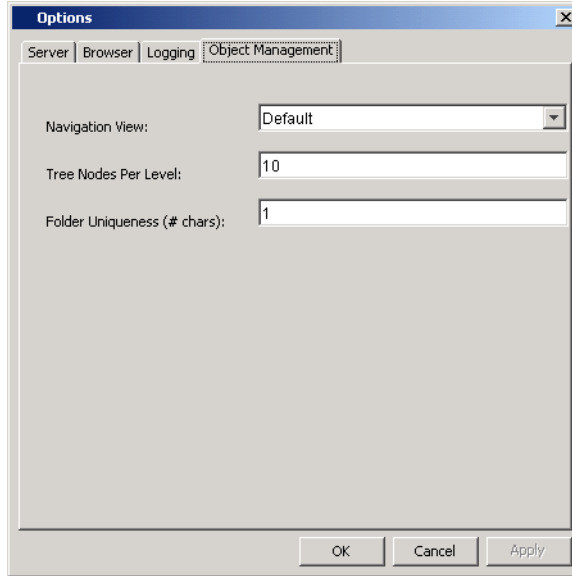
If you are not concerned with the number of folders and leaves that are displayed, and you want to be able to identify each node name quickly, you can set **Tree Nodes Per Level** to a number that is lower than the number of objects, and set **Folder Uniqueness** to a number larger than 1. OVPI will display more than the first character of each object name. (**Tree Nodes Per Level** must be a relatively low value; otherwise OVPI does not look at the **Folder Uniqueness** value.)

For example, if you have an object named Denver and another object named Destroy, you can immediately identify each one if you choose to display the first three characters of the name (**den** and **des**, respectively).

To modify the appearance of the Objects tree:

- 1 From the Edit menu, choose **Options**.
The Options window opens.
- 2 Select the **Object Management** tab.

The Object Management dialog box opens.



- 3 In the Navigation View box, click the down arrow to display the available views and select the desired view. You can select Device (the default), Customer, Location, Vendor, or Cards.

- 4 In the Tree Nodes Per Level box, enter the maximum number of levels that you want to appear at any given level in the Objects tree.

If the number of objects in the database exceeds this value, OVPI looks at the Folder Uniqueness value to determine how to reduce the number of displayed folders and leaves. See the discussion of these fields above.

- 5 In the Folder Uniqueness box, type the number of identifying characters that you want to be displayed for each folder or leaf in the tree at any given level.

OVPI implements this value only if the number of objects in the database exceeds the value in the **Tree Nodes Per Level** field. See the discussion of these fields above.

- 6 Click **Apply** (or **OK** if you have finished making all of your changes on the Options window).

- 7 From the View menu, choose **Refresh**.

OVPI updates the view in the Objects tree.

Importing Nodes

There are two steps involved when importing nodes. You must first create an ASCII file that contains the nodes you want to import, and then you must import that file.

Creating the Import File

Create an ASCII file that contains the list of nodes that you want to import. Each node that you want to import must have an entry in the file formatted as follows:

```
node_name, read_community, write_community, type_name, view_name,  
snmp_v1_flag, snmp_v2_flag, node_status, desc, snmp_profile, port_num,  
num_retries, timeout_sec, num_oids, get_bulk_size
```

Use a comma as a placeholder for any parameter that you do not include; otherwise, the node record will be skipped. [Table 7 on page 126](#) describes the parameters for the import file.

Importing the File

To import a file containing nodes to be managed:

- 1 From the File menu, choose **Import**.
The Import from dialog box opens.
- 2 Navigate to the file that contains the nodes you want to import.
- 3 Double-click the file name, or select the name and click **Open**.

Exporting Nodes

You can create a file that contains managed nodes, which you can use to export the nodes to another database.

To create an export file containing managed nodes:

- 1 From the File menu, choose **Export**.
The Export to dialog box opens.
- 2 Navigate to the location where you want to store the export file.
- 3 Type the file name in the **File name** box, and click **Open**.
A progress window displays while the file is being created.

Using Forms

A form lets you create new data or modify existing data in the database for a single database table. Creating and modifying managed objects from the Management Console is convenient and offers an alternative to importing property file information during report pack configuration.

Refer to the *Performance Insight Guide to Building and Viewing Reports* for more information about creating and deploying forms.

Modifying the Database

To use a form to modify data in the database for a selected managed object:

- 1 Click a managed object in the Objects tree. For example, select a device.
- 2 Double-click a form in the **General Tasks** or **Tasks specific to the selected objects** section.

The selected form appears in a separate window. For example, you might select a form called Change Protocol Defaults, which will display the default protocol settings that you can modify.

Interface Reporting
Change Protocol Defaults

Use this form to change the default settings for Protocols (interface types) for the Interface Reporting report pack. You can change the names for Unknown protocols, and modify the default duplex and threshold settings for each protocol. The system can be configured to not poll interface types that have the Monitor flag set to Off (0). Select the Protocol you want to change the settings for, and then enter the new settings below. Click the Apply button to save any changes. Click the Cancel button to cancel any changes. Click the OK button to save changes and close the form.

ifType	Protocol Name	Duplex	Monitor	Default Utilization Threshold	Default Discard Threshold	Default Error Threshold
1	other	2	1	90.00	1.00	2.00
2	regular1822	2	1	90.00	2.00	2.00
3	hdh1822	2	1	90.00	2.00	2.00
4	ddn-x25	2	1	90.00	2.00	2.00
5	rfc877-x25	2	1	90.00	2.00	2.00
6	ethernet-csma/cd	1	1	30.00	1.00	2.00
7	iso-88023-csma/cd	1	1	90.00	2.00	2.00
8	iso-88024-tokenBus	1	1	50.00	2.00	2.00
9	iso-88025-tokenRing	1	1	50.00	2.00	2.00
10	iso-88026-man	2	1	90.00	2.00	2.00
11	...	2	1	90.00	2.00	2.00

Protocol Name: Monitor:

Utilization Threshold: Error Threshold: Discard Threshold:

OK Apply Cancel

- 3 Make your changes to the form, and then click **OK**.

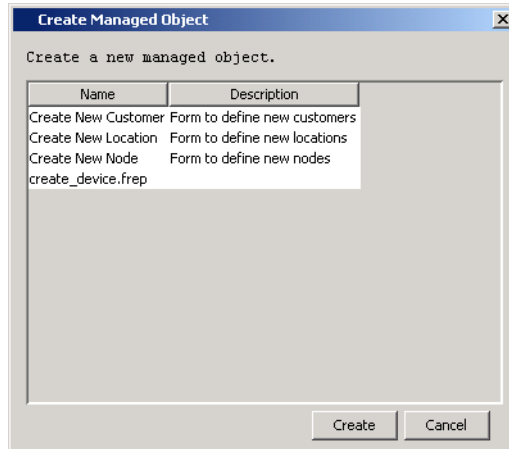
OVPI changes the property table values associated with the changes you made.

Creating New Managed Objects

To use a form to add a managed object to the database:

- 1 From the File menu, choose **New**.

The Create Managed Object dialog box opens.



- 2 Select the form you want to use, and then click **Create**.

The form opens. Note that all of the fields (input fields and combination input box) are blank.

- 3 Make your changes to the form, and then do one of the following:
 - If you are finished creating objects, go to [Step 4](#).
 - If you want to create another object, click **Apply**, then return to [Step 2](#).)
- 4 Click **OK**.

OVPI adds the objects to the database.

Grouping Managed Objects

Group Manager is an application that is available from the Management Console, which lets you create and maintain managed-object groups; these groups are referenced by the polling policies for your enterprise.

A *group* is a set of managed objects; that is, it corresponds to a set of rows in a property table. A *managed object* is an object that is managed by OVPI; primarily these are network devices and interfaces. Thus, a managed-object group is always a subset, not necessarily a proper subset, of the managed objects defined by a single property table.

Specifically, you can use Group Manager to do the following:

- Navigate a tree of the group elements.
- Create any of the following types of groups: enumerated list, property, derived, and folder.
- Delete and modify groups.
- Create, modify, and delete the polling policies associated with groups.

Grouping Types

This section describes the supported types of groups, where each group type defines a different means of specifying the group members:

- **All.** Comprises all the managed objects identified in a property table.

- **Enumerated list.** Comprises a specific set of managed objects that are identified by their object IDs from the property table. The composition of the group is static, unless you later add or delete specific objects.
- **Property.** Comprises a set of managed objects that share a common property or properties. This type of group is also called *rule-based*, because you define one or more rules that OVPI uses to select the objects for the group.

For example, a property group may comprise all the nodes located in Los Angeles. Unlike enumerated lists, rule-based groups are dynamic. If three new nodes are added in Los Angeles, the next time the Los Angeles node group is referenced it will have three new members.

➤ OVPI allows you to create a property group having any number of rules. However, property groups having more than two rules cannot be used by the remote poller. Property groups having any number of rules can be used to select nodes to be shown when viewing deferred node reports in Report Viewer or the Web Access Server (for example, the Near Real Time Snapshot report).

- **Derived.** Comprises a logical combination of objects from other groups. Derived groups are often used in the creation of polling policies.
- **Folder.** Supports the traditional file and folder view seen in an application such as Windows Explorer. You select objects for this type of group in the same way that you define an enumerated list group; the only difference is that a folder group can also have other folder groups inside it.

Grouping Methods

You can organize nodes by location, by administrative domain, by vendor, by supplier, by type, and so on. Alternatively, you may want to identify certain nodes for special handling; for example, you might want to poll backbone components for more data more frequently.

Grouping Nodes

You can subset the known node set the following ways:

- Base the grouping on a device property.

Using this method, nodes sharing a common value for an attribute are grouped together, but each node can be in only one group per attribute.

This method works well for some node organizations, such as those grouped by location, vendor, or supplier; however, this method does not work well for all manner of organizing nodes. For example, grouping nodes by type cannot be supported using a property or properties, because any given node may belong to many types. A node can be a router, a switch, a hub, and an RMON probe.

- Create an enumerated list of nodes for each group.

A node can appear in the list for each of routers, switches, and hubs. This list-based approach also supports the formation of groups on an ad hoc basis for specialized polling or more focused reporting. For example, a group of routers that are targeted for upgrade may be polled more frequently and extensively.

Alternatively, you might want to see a detailed report on a group of nodes that are directly affected by a new service being brought online.

- ▶ Creating an enumerated list is a more versatile form of grouping; however, maintaining enumerated lists is more complex and takes more work than maintaining rule-based groups.

Grouping Other Managed Objects

In addition to nodes, you can also group managed objects by property and by forming enumerated lists. For example, such groups may poll only WAN router interfaces, summarize by interface type, and report LAN and WAN interface information separately. The ability to create enumerated lists is extremely important in environments where there are large numbers of managed objects but few of real interest or for which data can be affordably processed and stored.

Group Categories

You create a group in the context of a specific *group category*. A group category is associated with only one property table, but a property table can be associated with multiple group categories.

For example, you can use the following categories for groups of managed nodes: a category corresponding to views, a category corresponding to types, and the required node group category.



You cannot use Group Manager to create new group categories or delete existing ones. Group categories are created when you create new property tables or install report packs.

Viewing Managed Object Groups

To view managed object groups:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.



You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).

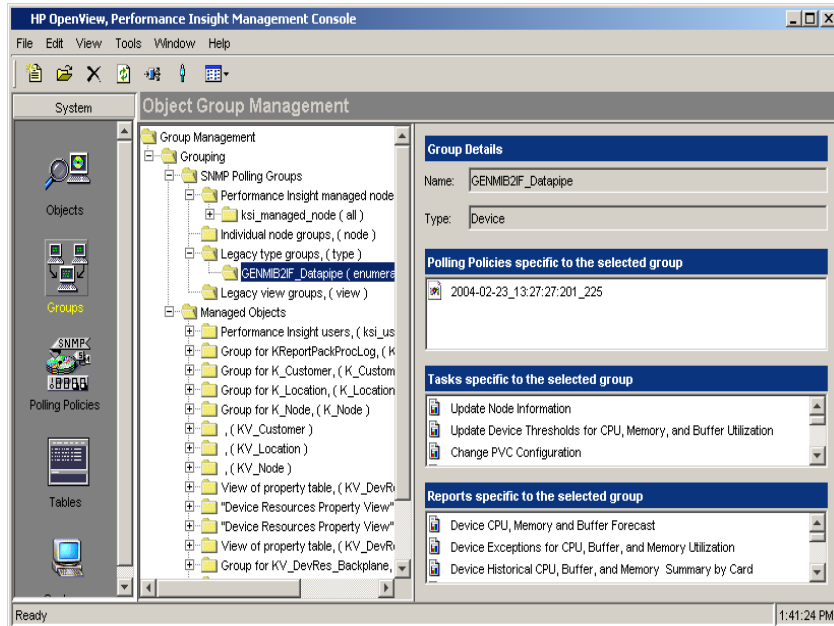
- 2 Click the **Groups** icon in the navigation pane.

The Object Group Management window appears in the Management Console.

The Object Group Management window has the following panes: The navigation and viewing panes.

The navigation window displays the group management tree, listing all the groups for the type of view that is selected. The viewing pane displays the result of your selection.

The following figure shows the Object Group Management window.



You expand the hierarchical levels of the tree just as you do when you view the contents of a Windows system using Windows Explorer. You open the folders in the navigation frame's tree as follows:

- To view the contents of a folder, click the plus sign (+).

The folder's contents are displayed.

- To close an open folder, click the minus sign (-).

- 3 Select a managed object group in navigation pane. (Alternatively, you can use the find feature; see [Finding a Specific Group on page 173](#)).

The following information appears in the viewing pane:

- The group's name and type appear in the **Group Details** section.
- The polling policies that are associated with that group appear in the **Polling Policies specific to the selected group** section.
- The forms that apply to the group you selected appear in the **Tasks specific to the selected group** section. This section displays only the forms that were created with the **Mode** property set to **Modify**.

- The reports that apply to the group you selected appear in the **Reports specific to the selected group** section.
- 4 To view a report or form, double-click the name of the report or form.
The selected item appears in a separate window.
 - If you select a report, OVPI generates and displays that report for the group you selected. You can view and print the report.
 - If you select a task, OVPI displays the appropriate form with data filled in for the group you selected. You can change certain values (for example, threshold values) in the form, and these values are updated in the associated property table. See [Modifying the Database on page 164](#).

Each time you select a different item in the group management tree, the information in the viewing pane is updated to reflect the group you selected.

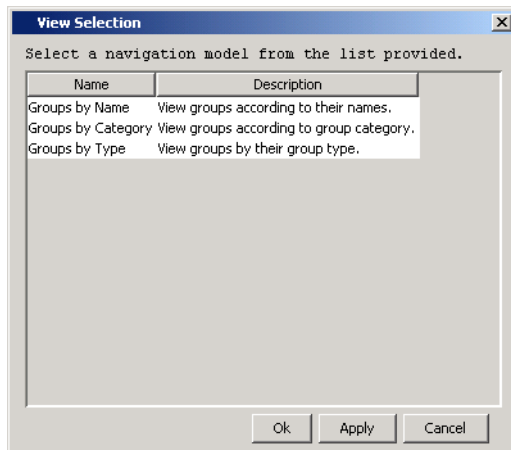
Changing the Groups View

You can change the property that is used to display managed object groups. (To set the default view, see [Setting the Display Options on page 187](#).)

To change the groups view:

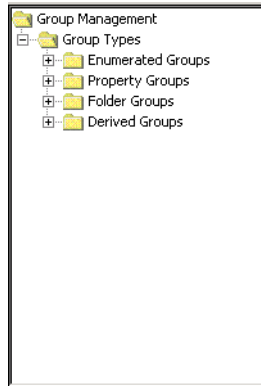
- 1 From the View menu, choose **Change View**.

The View Selection window opens.



- 2 Select a navigation view, then click **OK**.

The group management tree refreshes, displaying all the managed object groups that have that property assigned to them. For example, suppose you choose to view all groups according to their types; the group management tree will appear similar to the following picture after you expand the Group Types folder.




Finding a Specific Group

If you know the exact name of the group you want to display, you can use the find feature to access the information quickly.

- 1 From the Edit menu, choose **Find**.

The Find window opens.



- 2 Type the group name in the text box and click **Find**.


If OVPI finds the name you typed, the corresponding group is highlighted in the group management tree, and the associated forms and reports are listed in the respective areas.

- 3 Do one of the following:
 - If OVPI did not find the group name you typed, you can try typing another name.

- If the correct group name is highlighted in the tree, click **Close** in the Find window.
- 4 When the desired group name is highlighted, proceed to [Step 4 on page 172](#).

Creating Groups

You create a group by selecting one or more managed objects. When you right-click a group folder, a shortcut menu opens, which lets you create a new instance of any of the group types: enumerated, derived, property, and folder.



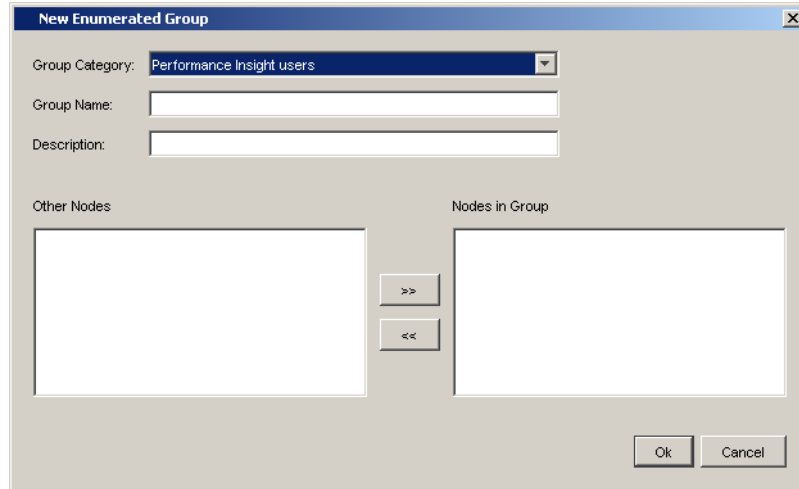
Creating an Enumerated List Group

An *enumerated list group* comprises a group of managed objects for a particular property table in a given category.


To create an enumerated list group:


- 1 Do either of the following:
 - If you know the group category that you want to use to create the group, right-click the group category's folder in the navigation pane, and choose **Create Enumerated** from the menu.
 - From the File menu, point to New, and choose **Enumerated Group**.

The New Enumerated Group or Create Group dialog box opens.



Depending on the method you used to open this dialog box, its appearance will vary as follows:

- When you use the File menu, the New Enumerated dialog box will open. It contains a Group Category list. You will have to specify the group category that you want to use to create the enumerated group.
 - When you right-click a group category, the Create Group dialog box opens. The Group Category list will not appear, because Group Manager assumes that you want to use the group category that you right-clicked on in the navigation pane.
- 2 *New Enumerated dialog box only:* Select a group category from the Group Category list.
 - 3 Type the name you want for the group in the Group Name box.
 - 4 Type a description for the group in the Description box.
 - 5 To add nodes to the group, do either of the following:
 - Double-click a node in the Other Nodes box to move it to the Nodes in Group box.
 - Select a node or an interface in the Other Nodes box, and then click  to add the node to the Nodes in Group box.

- 6 To remove any node from the group before you create it, do either of the following:
 - Double-click a node in the Nodes in Group box to move it to the Other Nodes box.
 - Select a node in the Nodes in Group box, and then click  to return the node to the Other Nodes box.
- 7 Click **OK** when you are finished.

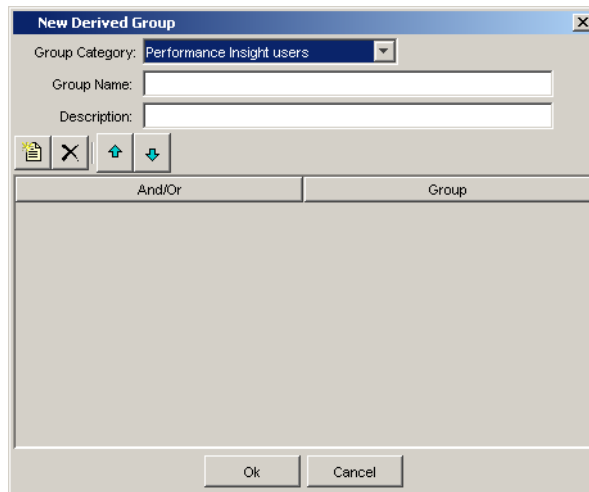
Creating a Derived Group

A *derived group* comprises a combination (for example, the union or intersection) of a group of enumerated list groups and property groups.



To create a derived group:


- 1 Do either of the following:
 - If you know the group category that you want to use to create the group, right-click the group category's folder in the navigation pane, and choose **Create Derived** from the menu.
 - From the File menu, point to New, and choose **Derived Group**.



The New Derived Group or Create Group dialog box opens.



Depending on the method you used to open this dialog box, its appearance will vary as follows:

- When you use the File menu, the New Derived dialog box will open. It contains a Group Category list. You will have to specify the group category that you want to use to create the derived group.
 - When you right-click a group category, the Create Group dialog box opens. The Group Category list will not appear, because Group Manager assumes that you want to use the group category that you right-clicked on in the navigation pane.
- 2 *New Derived dialog box only*: Select a group category from the Group Category list.
 - 3 Type the name you want for the group in the **Group Name** box.
 - 4 Type a description for the group in the **Description** box.
 - 5 To select a group, do the following:
 - a Click . A blank row appears below the **And/Or** and **Attributes** columns.
 - b Click the box below the **Attribute** column. A list appears from which you can select a group.
 - c Select a group from the list.
 - d Click . Another blank row appears.
 - e Click the box below the **And/Or** column. A list appears from which you can select how you want the groups compared.

The options are **AND** or **OR**. The **AND** option provides the intersection of the groups as a result; the **OR** option provides the union of the groups as the result.
 - f Click the box below the **Attribute** column. A list appears from which you can select another group.
 - g Select a group from the list.
 - 6 To change the order of a group in the list, select the group, and then do either of the following.
 - Click  to move the group up in the list.

- Click  to move the group down in the list.
- 7 To delete a group, select the group, and then click .
- 8 Click **OK** when you are finished.

Creating a Property Group

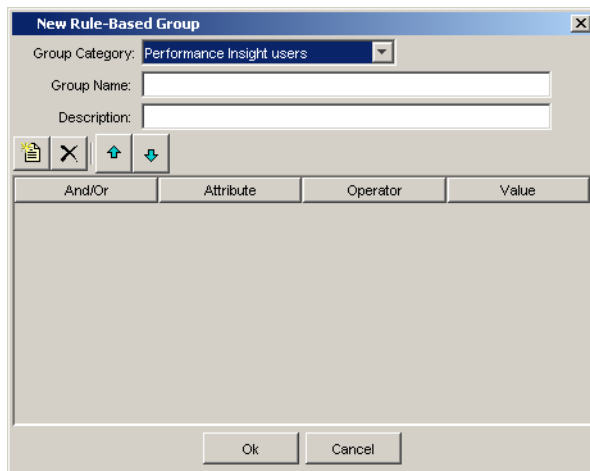
A *property group* comprises a group of managed objects that are based on certain properties in their property tables.

▶ A property group can only resolve two properties due to a limitations imposed by the polling system.

To create a property group:


- 1 Do either of the following:
 - If you know the group category that you want to use to create the group, right-click the group category's folder in the navigation pane, and choose **Create Property** from the menu.
 - From the File menu, point to New, and choose **Property Group**.

The New Rule-Based Group or Create Group dialog box opens.



Depending on the method you used to open this dialog box, its appearance will vary as follows:

- When you use the File menu, the New Rule-Based Group dialog box will open. It contains a Group Category list. You will have to specify the group category that you want to use to create the property group.
 - When you right-click a group category, the Create Group dialog box opens. The Group Category list will not appear, because Group Manager assumes that you want to use the group category that you right-clicked on in the navigation pane.
- 2 *New Rule-Based Group dialog box only:* Select a group category from the Group Category list.
 - 3 Type the name you want for the group in the Group Name box.
 - 4 Type a description for the group in the Description box.
 - 5 To select a group, do the following:


- a Click . A blank row appears below the And/Or, Attributes, Operator, and Value columns.
- b Click the box below the Attribute column. A list appears from which you can select a property.
- c Select a property from the list.
- d Select an operator from the Operator list.
- e Click the Value box, and then enter a value for the property.

This value can be an alphanumeric string.




When using the LIKE operator, its value must be a string and can include the following wildcard characters: underscore (`_`) and percent (`%`). The underscore character represents any single character. The percent character represents zero or more characters. For example:

`12%3` would match the following: 1233, 123, and 1244443.

`_ime` would match the following: time, lime and dime.

- f Click . Another row appears.
- g Click the box below the And/Or column. A list appears from which you can select how you want the groups compared.

The options are **AND** or **OR**. The **AND** option provides the intersection of the groups as a result; the **OR** option provides the union of the groups as the result.

- h** Repeat [Step b](#) through [Step e](#).
- 6** To change the order of the property in the list, select the property, and then do either of the following.
 - Click  to move the property up in the list.
 - Click  to move the property down in the list.
- 7** To delete a property, select the group, and then click .
- 8** Click **OK** when you are finished.

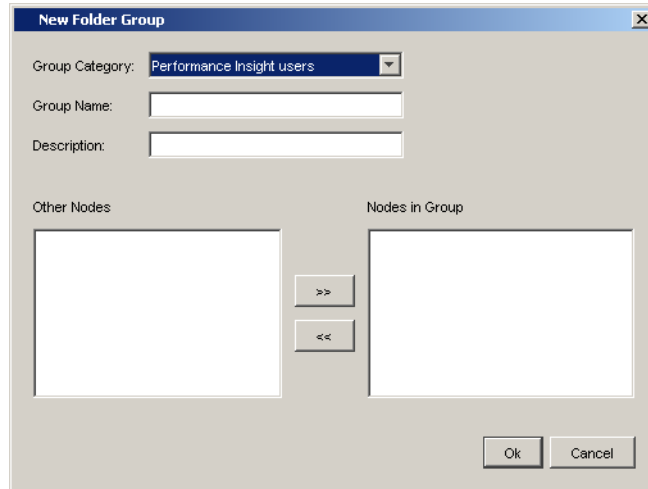
Creating a Folder Group

A *folder group* can contain an enumerated list of managed objects and other folder groups.


To create a folder group:


- 1** Do either of the following:
 - If you know the group category that you want to use to create the group, right-click the group category's folder in the navigation pane, and choose **Create Folder** from the menu.
 - From the File menu, point to New, and choose **Folder Group**.


The New Folder Group or Create Group dialog box opens.



Depending on the method you used to open this dialog box, its appearance will vary as follows:

- When you use the File menu, the dialog box will contain the Group Category list. You will have to specify the group category that you want to use to create the folder group.
 - When you right-click a group category, the Group Category list will not appear, because Group Manager assumes that you want to use the group category that you right-clicked on in the navigation pane.
- 2 *New Folder Group dialog box only*: Select a group category from the Group Category list.
 - 3 Type the name you want for the group in the Group Name box.
 - 4 Type a description for the group in the Description box.
 - 5 To add nodes to the group, do either of the following:
 - Double-click a node in the Other Nodes box to move it to the Nodes in Group box.
 - Select a node or an interface in the Other Nodes box, and then click  to add the node to the Nodes in Group box.

- 6 To remove any node from the group before you create it, do either of the following:
 - Double-click a node in the Nodes in Group box to move it to the Other Nodes box.
 - Select a node in the Nodes in Group box, and then click  to return the node to the Other Nodes box.
- 7 Click **OK** when you are finished.

 If you want to place subfolder groups in a folder group, you must first create all your folder groups, using the procedure above. Then use must edit the group to create subfolder groups by moving them into another folder group. See [Editing a Group on page 184](#).

Importing Groups

One way to add groups to the database is to import a file containing groups.

To import a file containing managed object groups:

- 1 From the File menu, choose **Import**.
The Open dialog box opens.
- 2 Navigate to the file that contains the groups you want to import.
- 3 Double-click the file name, or select the name and click **Open**.

A progress window displays while the file is imported into the database. You will now be able to see the imported groups in the group management tree.

Exporting Groups

You can create a file that contains one or more groups that you want to export to another database.

To create an export file containing a single managed object group:

- 1 In the group management tree, select the name of the group you want to export.

- 2 From the File menu, choose **Export** ⇒ **Selected**.

The Open dialog box opens.

- 3 Navigate to the location where you want to store the export file.

- 4 Type the file name in the **File name** box, and click **Open**.

A progress window displays while the file is being created.

To create an export file containing all managed object groups:

- 1 From the File menu, choose **Export** ⇒ **All Groups**.

The Open dialog box opens.


- 2 Navigate to the location where you want to store the export file.

- 3 Type the file name in the **File name** box, and click **Open**.

A progress window displays while the file is being created.

Managing Groups

When you right-click a *group leaf node*, a shortcut menu opens, which lets you open and view a group, edit a group (any group except an All group), or delete a group. You can also create a polling policy using the **New Collection** menu option; see [Creating a Group Polling Policy on page 188](#).




Viewing a Group

To view the contents of a group, do one of the following:

- Right-click a group name, then choose **Open Group** from the menu.
- Double-click a group name.

The contents of the group are displayed in an object browser window. (The format of the display depends on the display option of the browser window; see [Setting the Display Options on page 187.](#))



This window displays the managed objects that are members of this group. When you right-click a managed object, a shortcut menu appears, which lets you view the managed objects in different formats. You can view objects as large icons, small icons, a list, or as a list with all the object's property values displayed.

Editing a Group

You can add objects to a group or remove objects from it. If you are editing a folder group, you can also add other folder groups to the group.

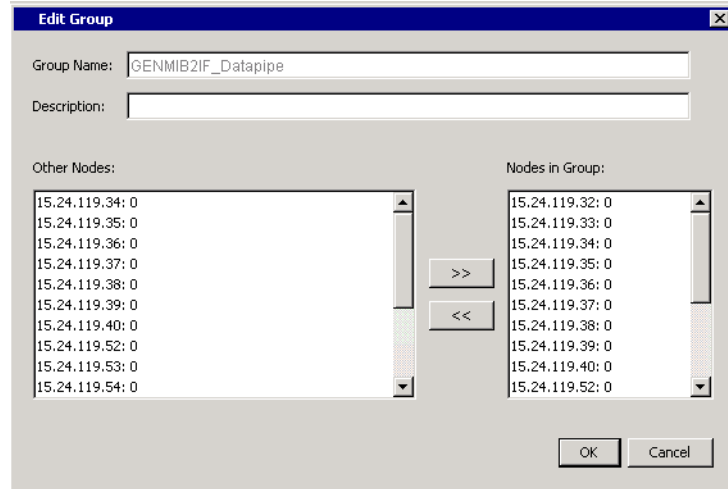




Group Manager does not allow you to edit All type groups.

To edit a group:

- 1 Right-click the group name, and choose **Edit *group_type*** (for example, **Edit enumerated**) from the menu.

The Edit Group dialog box opens.



- 2 To add nodes to the group, do either of the following:
 - Double-click a node in the Other Nodes box to move it to the Nodes in Group box.
 - Select a node or an interface in the Other Nodes box, and then click  to add the node to the Nodes in Group box.
- 3 To remove any node from the group before you create it, do either of the following:
 - Double-click a node in the Nodes in Group box to move it to the Other Nodes box.
 - Select a node in the Nodes in Group box, and then click  to return the node to the Other Nodes box.
- 4 Click **OK**.

Deleting a Group

When you delete a group, the managed objects are not removed; only the group is removed. If the group is referenced in a polling policy, you will not be able to delete the group until the polling policy is deleted.

To delete a group:

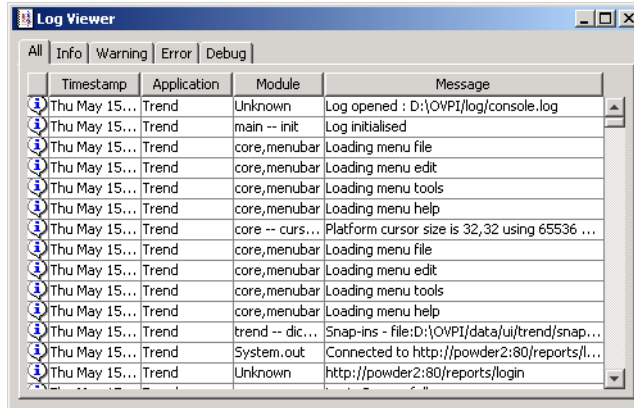
- 1 Right-click the group name, and choose **Delete Group** from the menu.
A deletion confirmation window displays.
- 2 Click **Yes**.

Viewing Log Information

Using the options on the View menu, you can view log file and debugging information.

To view log file information, choose **Log** from the View menu.

The Log Viewer window opens.



The Log Viewer window comprises five tabs, each displaying a different type of log information. There are four types of log messages: Info, Warning, Error, and Debug. The **All** tab displays all four types of messages.

To view output information, choose **Output** from the View menu.

The Output Viewer window displays the most current log messages as they are generated.

Setting the Display Options

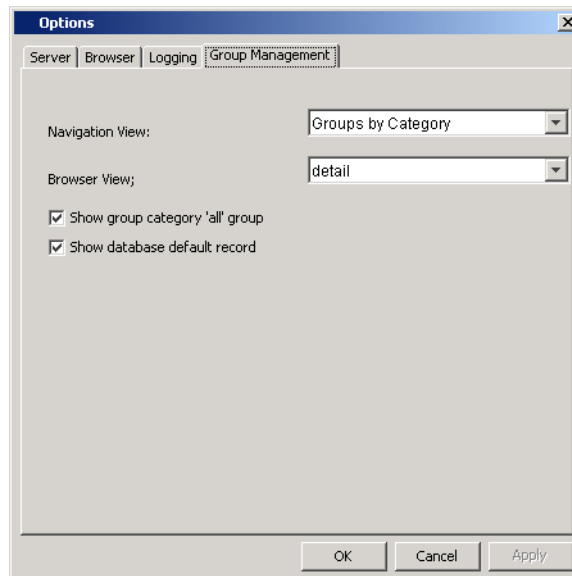
To set default values that affect the appearance of the group management displays, follow the steps below. If you change any of these default settings, the changes take effect the next time you log in to the Management Console.

- 1 From the Edit menu, choose **Options**.

The Options window opens.

- 2 Select the Group Management tab.

The Group Management page opens.



- 3 To set the group hierarchy structure in the navigation tree, click the Navigation View arrow and select one of the following options from the list: Group by Category, Group by Type, or Group by Name.
- 4 To set how group objects are displayed in a browser window, click the Browser View arrow and select one of the following options from the list: large, small, list, or detail.
- 5 Do one of the following:
 - To display groups of the All type, that is, groups that contain all the managed objects for the selected group category, select the **Show group category 'all' group** box.

- To suppress the display of All type groups, click to clear the **Show group category 'all' group** box.
- 6 Do one of the following:
 - To display default records during group management, select the **Show database default record** box. Default records are records that OVPI automatically inserts when a property table is created in order to maintain database integrity.
 - To suppress the display of default records, click to clear the **Show database default record** box.
 - 7 Click **OK**.

Managing Group Polling Policies

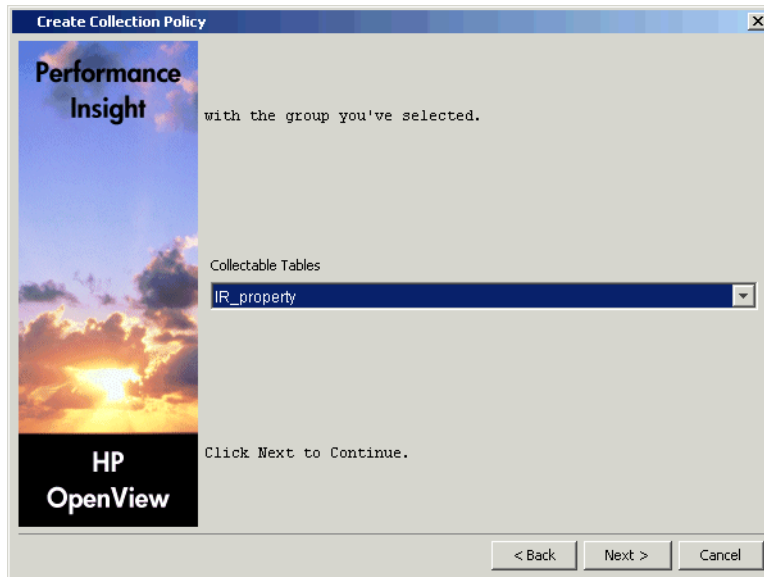
You can create a group polling policy using the Create Collection Policy wizard. You can also edit or delete a polling policy associated with any group by using the Polling Policy tab that is available when you view a group.

Creating a Group Polling Policy

To create a polling policy for a group:

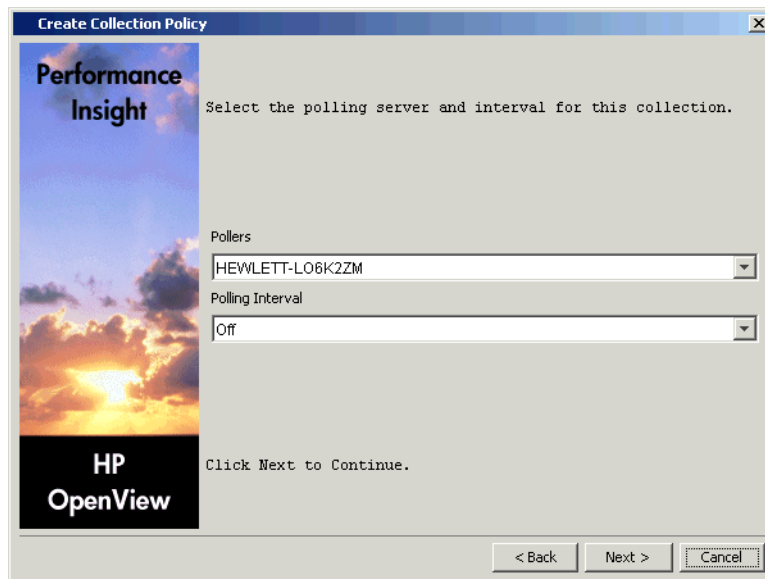
- 1 Right-click the group for which you want to create a polling policy, and choose **New Collection** from the shortcut menu.

The following page opens, which lists the tables associated with the group that are collectible (typically, these are raw and rate tables)



- 2 In the **Collectable Tables** list, click a table, and then click **Next**.

The following page opens, which asks for the system that will do the polling (the poller) and the interval at which you want to poll for data



- 3 Do the following:
 - a Select the server that will poll for the data from the **Pollers** list.
 - b Select the interval at which the poller should poll for data from the **Polling Interval** list.
 - c Click **Next**.

The following page opens, which asks for the name of the polling policy and a description

- 4 Do the following:
 - a Type the name of the polling policy in the **Collection Name** box.
 - b Type a description for the policy in the **Description** box. A description is optional.
 - c Click **Next**.

A summary page opens, which displays the features of the polling policy.

- 5 Click **Finish**.

Modifying a Group Polling Policy

To modify a group's polling policy:

- 1 Locate the group in the group management tree.

The right side of the window displays the associated polling policies in the **Polling Policies specific to the selected group** box.
- 2 Right-click the policy you want to modify, and choose **Properties** from the shortcut menu.

The Edit Polling Policy dialog box opens.

- 3 Modify the polling policy using the instructions in [Editing Polling Policies on page 123](#).

Deleting a Polling Policy

To delete a group's polling policy:

- 1 Locate the group you want to delete in the group management tree.
The right side of the window displays the associated polling policies in the **Polling Policies specific to the selected group** box.
- 2 Right-click the policy you want to delete, and choose **Delete** from the shortcut menu.
A dialog box opens, asking if you want to delete the selected polling policy.
- 3 Click **Yes**.

Changing the Polling Policy View

To change the way polling policies are displayed in the list box:

- 1 Locate a group in the group management tree.
The right side of the window displays the associated polling policies in the **Polling Policies specific to the selected group** box.
- 2 Right-click the policy, and choose **View** from the shortcut menu.
- 3 Select one of the following:
 - Large icons—Displays polling policy names horizontally, with large icons.
 - Small icons—Displays polling policy names horizontally, with small icons.
 - List—Displays polling policy names in column format, with small icons.
 - Details—Displays detailed information for each polling policy, including its name, description, and polling interval.

Using Groups in Reports

When opening a report with a deferred node selection using Report Viewer or the Web Access Server, you can select individual nodes and interfaces or a group of nodes and interfaces.

For example, suppose you create an enumerated list group called “HP Systems” and it contains the nodes Monty and Caesar. When you open the report, you can select this group for the report, and only the data for the Monty and Caesar systems will appear in the report.

Refer to the *Performance Insight Guide to Building and Viewing Reports* for more information about selecting deferred nodes for reports.

Managing Your System

This chapter explains the System and Network Administration application, which is available from the Management Console. This application lets you manage the following:

- **Web Access Servers.** You can add, delete, and modify the properties of these systems. Additionally, you can view their properties and change your current Web Access Server to another available server.
- **Databases.** You can add, delete, and modify the database by changing initialization parameters and the database password. You can also increase the size of the database and change your current database to another available database.

Starting the System and Network Administration Application

To start the System and Network Administration application:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.



You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).

- 2 Click the **Systems** icon in the navigation pane.

The System/Network Administration window appears in the Management Console. The left side of the window displays all of the Web Access Server systems that you can access; that is, you can log on to any of these to use OVPI.

When you expand a system, you will also be able to view the available databases.

Figure 1 shows an example of how this window may look.





Figure 1 System/Network Administration Window

When the  icon is blue, this indicates the system that you are currently are logged on to.

The System/Network Administration Window

The System/Network Administration window has the following panes:

- Viewing pane. The viewing pane is on the right side of the window.

- **Navigation pane.** The navigation pane is on the left side of the window. It shows a tree for the systems that you can access. The tree contains folders; the folders contain pages, which may include other pages or additional folders.

You expand the hierarchical levels of the system just as you do when you view the contents of a Windows system using Windows Explorer. You can expand the folders to see the additional pages or folders it may contain, or collapse those levels.

You open the folders in the navigation frame's tree as follows:

- To preview the contents of a folder, click the plus sign (+).
The folder's contents is displayed in the viewing frame.
- To close an open folder, click the minus sign (-).

Displaying System Information

To display and change information about the Web Access Server you are currently logged on to:

- 1 Click the name of the server in the navigation pane. The viewing pane displays the name of that server and also lists the other available servers.

You can right-click the server in the viewing pane to display a menu that lets you do any of the following:

- **Add System.** Add a new Web Access Server. See [Adding a Web Access Server on page 201](#) for details.
- **Properties.** Display and change the properties of the selected Web Access Server. See [Viewing and Modifying Web Access Server Properties on page 201](#) for details.
- **Delete System.** Delete the selected Web Access Server. See [Deleting a Web Access Server on page 206](#) for details.
- **View.** Change the appearance of the systems in the viewing pane. See [Changing Web Access Server Appearance in the Viewing Pane on page 199](#) for details.


If you select a Web Access Server that you are not connected to, the following dialog box opens:



To change the Web Access Server you are currently connected to, click **Yes**. Alternatively, you can also use the Change Server icon; see [Changing Your Web Access Server on page 205](#) for details.

- 2 Expand the plus sign next to your current Web Access Server system; the icon next to this system is highlighted blue.

The navigation pane displays the following for the Web Access Server system:

- Web Access Server Properties page. This page lets you view the hostname, HTTP port, and HTTPS port for the server.
- Databases folder. Expand this folder to view the databases that your current system can access. Your current database is denoted by the  icon. This folder contains the following options:
 - OVPI Databases. Lists the databases that contain the OVPI database schema. Each database folder contains the following page and folder:
 - Database Initialization Parameters page. This page lets you view all the database parameters and change the parameters that can be modified. See [Changing Database Initialization Parameters on page 210](#) for more information.
 - Tablespaces folder. Expand this folder to display the properties for the database's tables. See [Increasing Database Size on page 211](#) for more information.
 - External Reference Databases. Lists any generic databases; these database are designed for use with datapipes.

These database do not contain the OVPI database schema. The following generic databases are supported: Sybase, Oracle, SQLServer, and Solid. You can view the database's properties; however, you cannot change the size of the database or view initialization parameters for generic databases.

- Local Database References. Lists the local databases.

This folder is displayed only if you installed the OVPI Database Schema and Performance Manager components. Right-click this folder to add databases to the local system.

Changing Web Access Server Appearance in the Viewing Pane

When you click on the name of the server, the viewing pane displays the name of that server and also lists the other available servers.

To change how the servers are displayed in this pane:

- 1 Right-click a server in the viewing pane.
- 2 Select **View** from the menu, and choose one of the following options from the submenu: **Large Icons**, **Small Icons**, **List**, or **Details**.

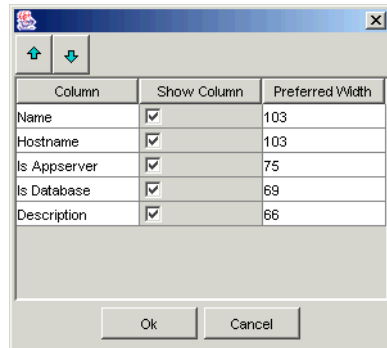
Changing the Details Display

If you selected **Details** when you followed the procedure in [Changing Web Access Server Appearance in the Viewing Pane](#), the servers are displayed listing the following details: the name of the system, its hostname, whether it is a Web Access Server, whether it is a database, and a description, if one was supplied when it was installed.

You can change the detail columns that will display or the order in which the details are displayed by doing the following:

- 1 Right-click a the name of a detail column (for example, Is Database). A menu displays the detail columns available. Do either of the following:
 - To prevent a column from being displayed, select the column name to remove the check mark next to it.

- Select **More** from the menu. The Modify Detail Columns dialog box opens.



- 2 Do any of the following:
 - To change the order in which the column details are listed, select the column name you want to reorder, and click the down arrow or the up arrow button to reorder the column, respectively.
 - To prevent a column name from displaying, click to clear its **Show Column** box *or* select its **Show Column** box to display it.
 - To change the width of a column, double-click the **Preferred Width** box next to the column name, delete the current value, and type a new value.
- 3 Click **OK** to make your changes take effect.

Managing Your Web Access Server Systems

Using the System/Network Administration component, you can add Web Access Server systems, edit their properties, remove them from the list of available systems, and synchronize the local system with the currently selected Web Access Server.

Adding a Web Access Server

To add a Web Access Server to the list of systems that you can access:

- 1 Do either of the following:
 - Right-click the System Information folder in the viewing pane of the System/Network Administration window, and select **Add OVPI Web Access Server** from the menu.
 - Right-click the Web Access Server in the viewing pane, and select **Add System** from the menu.

A wizard opens, which displays the Welcome window.

- 2 Click **Next** to continue. The Connection Information windows opens.
- 3 Do the following:
 - a In the Hostname box, type the hostname of the Web Access Server.
 - b In the HTTP Port box, type the port number for the Web Access Server if its port is not 80, which is the default.
 - c Select the **Use TSL/SSL** box if the Web Access Server is using SSL and you want secure communications with it.
 - d Click **Next**. The Summary window opens.
- 4 Do the following:
 - a Verify that the information is correct.


If necessary, click **Back** to return to a previous page to make a change. After you make a change, click **Next** until you return to the Summary window.
 - b Click **Finish** to add the Web Access Server.

The wizard closes, and the Web Access Server is added to the list of available systems.

Viewing and Modifying Web Access Server Properties

You can view the properties of the current Web Access Server and modify them, if needed.

To view general information for your current Web Access Server, expand the folder for your Web Access Server in the navigation pane, and click **Web Access Server Properties**. The properties are displayed in the viewing pane of the window, for example:



To view all the properties for your Web Access Server system:

- 1 Right-click the server in the navigation pane of the System/Network Administration window.

Alternatively, if the servers are displayed in the viewing pane, you can right-click the system there.

- 2 Select **Edit Properties** from the menu. The System Properties dialog box opens.



Depending on the status of the Web Access Server selected, the text in the System Properties dialog box will indicate where the changes you make will be saved. For example, if the server you selected is the current server, the changes take effect on it. If you selected a server that is not the current server, the changes are applied remotely to that server; your current server is not affected by the changes you make to the remote server.

This dialog box contains the following tabs:

- General. It contains the following information:
 - Name. Type the name of the database.
 - Hostname. Type the hostname of the database.
 - IP Address. Type the IP address of the database.
 - WAS Properties. It contains the following information:
 - Port. Type the port number of the system.
 - Authentication Realm. Type the realm. For more information about the realm, see [Step b on page 330](#).
 - Document Root. Type the document root. For more information about the realm, see [Step c on page 330](#).
 - Communication Protocol. Click **HTTP** or **HTTPS**.
 - Database Properties. It contains the following
 - Port. Type the port for the database.
 - Description. Type a description for the database; a description is optional.
 - Vendor. Click the Vendor arrow, and select the database type (**Sybase**, **Oracle**, **Solid**, and **SQLServer**).
 - Database Instance. Type the database instance. By default, they are as follows: DSQUERY (Sybase) or \$ORACLE_SID (Oracle).
 - Maximum Connections. Type the maximum number of database connections that the database pool will maintain for this database.
 - Database Username. Type the database username. By default, this is dsi_dpiped.
 - Password. Type the password for the database username.
-

- Default Database. Click this option if you want to make the database that the Web Access Server is currently accessing the default database.
 - Default Collection Database. Click this option if you want to make the database that the Web Access Server is currently accessing the default collection database.
- 3** Select a tab, make any changes, and click **Apply** to make your changes take effect. Click **OK** when you are finished.

If you would like to view the properties of a Web Access Server that is not the current server, you must make that server your current server. See [Changing Your Web Access Server on page 205](#) for details.

Synchronizing Web Access Servers

You can synchronize the local system with the current Web Access Server.

To synchronize a system with the current Web Access Server:

- 1** Right-click the system in the navigation pane of the System/Network Administration window.

If you select your current system, you will not be able to select any synchronization options. You cannot synchronize a system with itself; thus, choose another system.

- 2** Select **Synchronize** from the menu, and choose one of the following options:

- **Full Synchronization.** First, this option lets you synchronize your local systems with the Web Access Server you selected in [Step 1](#), where all information from the Web Access Server supersedes duplicate information on your local systems. Second, this option lets you synchronize all of your all of the systems that your system can access with the Web Access Server you selected in [Step 1](#) and adds only those systems that the Web Access Server does not currently access.

Basically, this synchronization option performs the synchronization done by the **Update localhost** option and the **Update Web Access Server** option.

For example, duplicate references on the local systems are overwritten according to what the Web Access Server contains, and then any references the local systems have that the Web Access Server does not are uploaded to the Web Access Server.

- **Update localhost.** Lets you synchronize your local systems with the Web Access Server you selected in [Step 1](#), where all information from the Web Access Server supersedes duplicate information on your local systems. (This update performs the opposite function of the **Update Web Access Server** option.)

For example, if your local system has an entry for a system named Thunderbolt and the Web Access Server has an entry for the same system, the Web Access Server's entry for Thunderbolt will overwrite the local systems entry for Thunderbolt, because the Web Access Server is the master in this case.


- **Update Web Access Server.** Lets you synchronize all of the systems that your system can access with the Web Access Server you selected in [Step 1](#) and adds only those systems that the Web Access Server does not currently access.


For example, if a system you compare the Web Access Server to has a database that you want added to the Web Access Server, the synchronization process will add the database to the Web Access Server.

After you select a synchronization option, the synchronization dialog box opens and starts the synchronization process; the dialog box will indicate if this process was successful. When the synchronization process completes, the dialog box closes.

Changing Your Web Access Server

To change your current Web Access Server system:

- 1 Click the Change Server icon  in the toolbar **or** choose **Change Server** from the File menu. The Change Server dialog box opens.



- 2 Select the Web Access Server you want, and click **Connect**. After the connection is made, the icon next to this system should be highlighted blue.


Deleting a Web Access Server

To delete a Web Access Server:

- 1 Right-click the server in the viewing pane of the System/Network Administration window.

Alternatively, if the servers are displayed in the viewing pane, you can right-click the system there.

- 2 Select **Delete**. A confirmation dialog box opens.



- 3 Click **Yes**.

Managing Your Databases

You can add a database, change the parameters associated with a database, change tablespace properties, and change the database server you are using.

Adding a Database

To add a database server:

- 1 Expand the folder for the server to which you want to add the database.
- 2 Right-click the Databases folder, and select **Add Database** from the menu.

Alternatively, if the servers are displayed in the viewing pane, you can right-click the system there and select the **Add Database** option.

A wizard opens, displaying the Welcome window.

- 3 Click **Next** to continue. The Type of Database window opens.

- 4 Select one of the following:
 - **OVPI Database.** This option lets you add a database that contains the OVPI database schema.
 - **Generic Supported Database.** This option lets you add a generic database. A generic database does not contain the OVPI database schema. Supported generic database types are as follows: Sybase, Oracle, SQLServer, and Solid.

If you selected an OVPI database, the Connection Information window opens. Go to [Step 5](#).

If you selected a generic supported database, the Database Connection Information window opens. Go to [Step 6](#).
 - 5 Do the following:
 - a In the Hostname box, type the hostname of the database.
 - b In the Port box, type the port number for the database if its port is not 80, the default.
 - c Select the **Use TLS/SSL** box if the database is using SSL and you want secure communications with it.
 - d Click **Next**. The Summary page opens. Go to [Step 7](#).
 - 6 Do the following:
 - a In the Hostname box, type the hostname of the database.
 - b In the Port box, type the port number for the database if its port is not 80, the default.
 - c In the Description box, type a description. A description is optional.
 - d Select a vendor from the Vendor list.
 - e In the Database Instance, type the database instance.
 - f In the password box, type the password for the `dsi_dpipes` account.
 - g Click **Next**. The Summary page opens.
 - 7 Do the following:
 - a Verify that the information is correct.

If necessary, click **Back** to return to a previous page to make a change. After you make a change, click **Next** until you return to the Summary page.
-

- b** Click **Finish** to add the database.

The wizard closes, and the database is added to the list of available database for that server.

Viewing Database Properties

You can view the properties for your database and modify the database password, if needed.



If you want to change the properties for your database, use the tools (Sybase Central or Oracle Enterprise Manager) your database provides.

To view the database properties:

- 1 Expand the server folder that contains the database you want to view.
- 2 Expand the Databases folder.
- 3 Click the database you want to view. The database's properties display in the viewing pane.

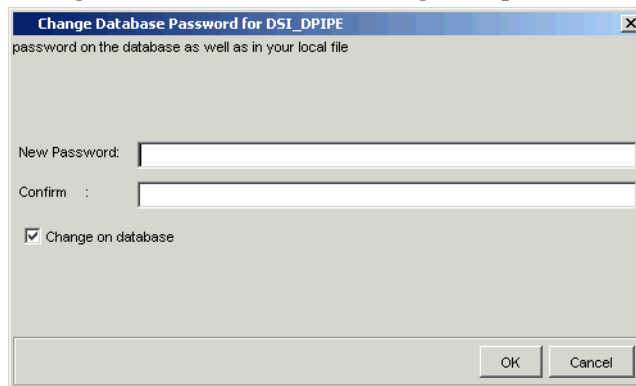
The screenshot shows a dialog box titled "Database Properties for gambit". It contains several fields and a button:

- Database Vendor: Oracle
- Database Instance: ovpi
- Connection Pool Size: 10
- JDBC Connection String: jdbc:oracle:thin:@gambit:1521:ovpi
- ODBC Connection String: DSN=OVPI_ORACLE;SID=ovpi;PORTNUMBER=1521;HOSTNAME=gambit
- Change Database Password... button

The properties shown are as follows:

- Database Vendor. This is the database type (Sybase or Oracle).

- Database Instance. By default, these names are as follows: DSQUERY (Sybase) or \$ORACLE_SID (Oracle).
 - Connection Pool Size. This is the number of connections that you have made available to OVPI.
 - JDBC Connection String. This is the string that OVPI's Java components use to connect to the database.
 - ODBC Connection String. This is string that OVPI's C-based components use to connect to the database.
- 4 To change the database password, click **Change Database Password**. The Change Database Password dialog box opens.



Do the following:

- a Type the new password in the New Password box.
- b Retype the new password in the Confirm box.
- c Select **Change on Database** to change the password on the database as well as in the `systems.xml` file.

For example, suppose you have changed the database password on system A, and the Web Access Server is located on system B. If you select the **Change on Database** option, the Web Access Server will be able to access the database on its next attempt, because it will get the updated password from the `systems.xml` file. If you did not select this option, the Web Access Server will not be able to access the database on System A.

- d Click **OK**.

- e Restart the Web Access Server and any client applications that are open. For information about restarting the Web Access Server, see the *Performance Insight Reference Guide*.

Changing Database Initialization Parameters

To change database parameters:

- 1 Expand the server folder that contains the database you want to view.
- 2 Expand the Databases folder.
- 3 Click **Initialization Parameters**.

The database initialization parameters display in the viewing pane.

Modifi...	Requires R...	Parameter	Value	Description
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	SQL Perfmon Integration	0	SQL Perfmon Integration
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	additional network mem...	5013504	additional netmem
<input checked="" type="checkbox"/>	<input type="checkbox"/>	allow backward scans	1	allow backward scans
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	allow nested triggers	1	Allow triggers to be invoked within triggers
<input checked="" type="checkbox"/>	<input type="checkbox"/>	allow procedure grouping	1	allow procedure grouping
<input checked="" type="checkbox"/>	<input type="checkbox"/>	allow remote access	1	Allow remote access
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	allow resource limits	0	allow resource limits
<input checked="" type="checkbox"/>	<input type="checkbox"/>	allow sendmsg	0	allow sendmsg
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	allow sql server async i/o	1	allow sql server async i/o
<input checked="" type="checkbox"/>	<input type="checkbox"/>	allow updates to syste...	0	Allow updates to system tables
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	audit queue size	100	Audit Queue Size
<input checked="" type="checkbox"/>	<input type="checkbox"/>	auditing	0	auditing
<input checked="" type="checkbox"/>	<input type="checkbox"/>	cis bulk insert batch size	0	cis bulk insert batch size
<input checked="" type="checkbox"/>	<input type="checkbox"/>	cis connect timeout	0	cis connect timeout
<input checked="" type="checkbox"/>	<input type="checkbox"/>	cis cursor rows	50	cis cursor rows
<input checked="" type="checkbox"/>	<input type="checkbox"/>	cis packet size	512	cis packet size
<input checked="" type="checkbox"/>	<input type="checkbox"/>	cis rpc handling	0	cis rpc handling
<input checked="" type="checkbox"/>	<input type="checkbox"/>	configuration file	0	configuration file
<input checked="" type="checkbox"/>	<input type="checkbox"/>	cpu accounting flush int...	200	CPU Accounting Flush Interval
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	cpu grace time	500	cpu grace time
<input checked="" type="checkbox"/>	<input type="checkbox"/>	curread change w/ ope...	1	curread change w/ open cursors
<input checked="" type="checkbox"/>	<input type="checkbox"/>	current audit table	1	current audit table
<input checked="" type="checkbox"/>	<input type="checkbox"/>	deadlock checking period	500	deadlock checking period
<input checked="" type="checkbox"/>	<input type="checkbox"/>	deadlock retries	5	deadlock retries
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	default character set id	2	default character set ID

The Parameters window shows the following information about each parameter:

- Modifiable. Indicates if you can modify the parameter

- Requires Restart. Indicates if you must restart the database after you modify the parameter
- Parameter. The name of the parameter
- Value. The current value of the parameter
- Description. A description of the parameter

Increasing Database Size

A *tablespace* is a place where database tables are stored; it is similar to a catalog. You can increase the size of the tablespaces, which essentially increases the size of the database.

When using Sybase, you can expand tablespaces by adding datafiles. Generally, Oracle tablespaces are set to automatically extend; therefore, you should not need to add additional space. In a situation where the database administrator has put limits on the amount of space used, then it may be necessary for you to extend the database.

Database Component Sizing Guidelines

Table 1 lists the recommended sizes for database components.

Database Component	Guideline
Temporary database (tempdb)	50% of your OVPI database (minimum), up to 1 GB
Transaction log	50% of your OVPI database (minimum), up to 2 GB

Each database consists of data and log relating to Sybase devices where each device maps to a file or raw partition. The installation process designates an initial amount of disk space for data and log. Since each OVPI system has its own unique characteristics, you may not need to increase your database size. If you are unsure whether or not to change the size of your database, contact HP Technical Support.

Increasing the Size of the Database

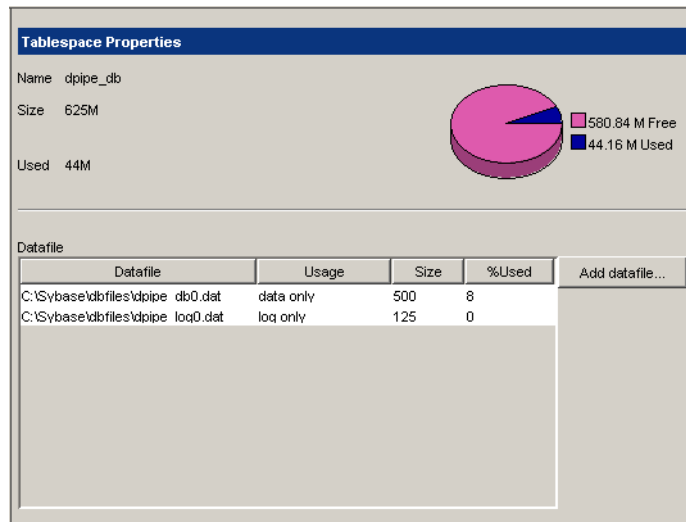
To expand the database:

- 1 Expand the server folder that contains the database you want to change.
- 2 Expand the Databases folder. The available databases are listed.
- 3 Expand the folder for the database you want to change.
- 4 Expand the `Tablespaces` folder.

The following tablespaces will appear depending on the type of your database:

- Sybase. The tablespaces for Sybase are as follows:
 - `dpipe_db`. Sybase recognizes the OVPI database as `dpipe_db`.
 - `tempdb`. This is where Sybase creates temporary tables as necessary. For example, if you run a report that is based on a query that does grouping or sorting, Sybase will create a temporary table in `tempdb` in which to store the data returned by the view. Sybase installation allocates a small amount of space to `tempdb` (approximately 2MB).
- Oracle. The tablespaces for Oracle are `dpipe_default_seg`, `dpipe_property_ind_seg`, `dpipe_property_seg`, `dpipe_upload_ind_seg`, `dpipe_upload_seg`, `dpipe_rate_ind_seg`, `dpipe_rate_seg`, `dpipe_summary_ind_seg`, `dpipe_summary_seg`, and `dpipe_overflow_seg`.


- 5 Click on a tablespace. The Tablespace Properties window opens in the viewing pane. It will look similar to the following:



The Tablespace Properties window shows the following information:

- Name. The name of the datafile
- Size. The size of the datafile
- Used. The percentage of the datafile that is used
- Datafile. The filename of the datafile.
- Usage. Lists what the file is used for
- Size. The size of the file
- %Used. The percentage of file that is used. Generally, if this value shows more than 60% of the database as used, it is recommended that you enlarge the database. Alternatively, you can delete unnecessary data or shorten the length of time you store your data.

- 6 To add a datafile, click **Add Datafile**. The Extend dialog box opens.



- 7 Do the following:

- a Type the name of the data file in the Datafile box **or** click **Browse** to locate the file.

The Browse option lets you select a file on the local system or navigate to a directory on the database system where you want to create the file. For example, if you are on a Windows system and the database is located on a UNIX system, the Browse option will let you browse the directory hierarchy of the UNIX system.

- b Type the size of the datafile in the Size (M) box.

- c Sybase only: Do the following:

- Click **Data Storage** to add more space for OVPI data in the Sybase database.
- Click **Transaction Log** to add more space to the database logsegment.

The *logsegment* is the area of the database where changes to the database are stored until a commit is issued. It is recommended that you keep the log segment at about 50% of the total size of all non-log data files.

- d Click **Extend**.


The new data file will appear in the window.




If Sybase is unable to allocate space for a database, you may see error #1105. See [Troubleshooting Sybase Space Allocation Problems on page 432](#) for information about how to free up space in the log.

Changing the Database Server

To change the database server you are using:

- 1 Click  on the toolbar *or* choose **Change Server** from the File menu. The Change Server dialog box opens.



- 2 Select the database you want, and click **Connect**. After the connection is made, the database icon next to this system should be highlighted green.

Managing Tables

This chapter explains how to monitor the use and growth of the database, view and create data and property tables, and set defaults for tables, both individual and global, using the HP OpenView Performance Insight (OVPI) Table Manager application.

Specifically, this chapter describes how to:

- View database information and statistics. See [page 219](#).
- View table information. See [page 221](#).
- View table details. See [page 227](#).
- View data in a table. See [page 235](#).
- Create tables. See [page 259](#).
- Set table defaults. See [page 277](#).

For more information about tables, refer to the *Performance Insight TEEL Reference Guide*. This guide provides information about the attributes of tables and creating TEEL files.

Starting Table Manager

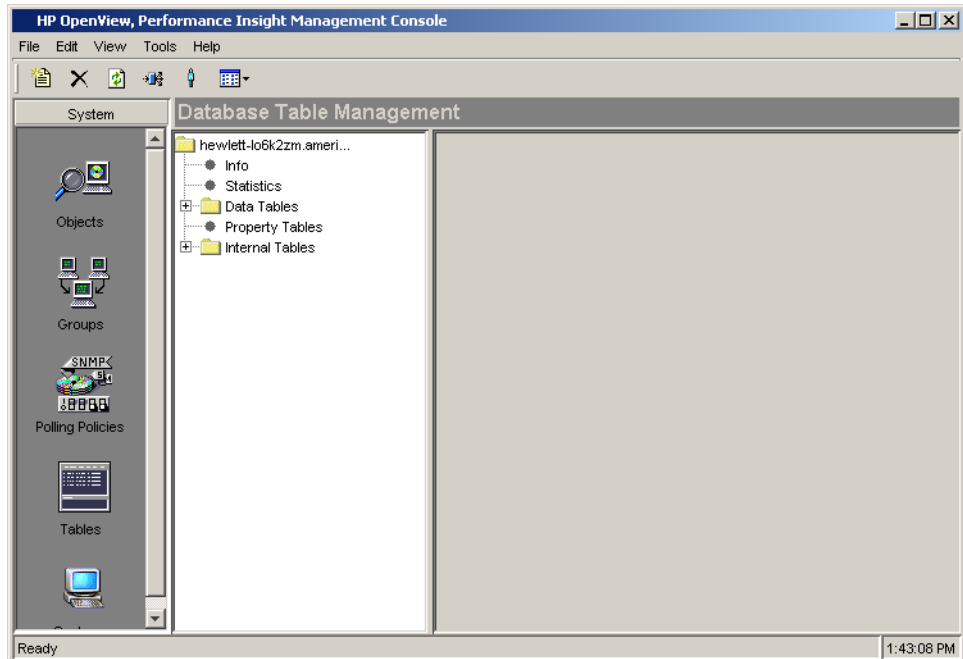
To start Table Manager:

- 1 Start the Management Console. See [Starting the Management Console on page 76](#) for details.

▶ You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts on page 295](#).

- 2 Click the **Tables** icon in the navigation pane of the Management Console.

The Database Table Management window opens.



The Database Table Management window has the following panes:

- **Viewing pane.** The viewing pane is on the right side of the window.

- **Navigation pane.** The navigation pane is on the left side of the window. It shows a tree for the database that you are currently viewing. The tree contains folders; the folders contain items, which may include other items or additional folders.

You expand the hierarchical levels of the database just as you do when you view the contents of a Windows system using Windows Explorer. You can expand the folders to see the additional items or folders it may contain, or collapse those levels.

You open the folders in the navigation frame's tree as follows:

- To preview the contents of a folder, click the plus sign (+).
The contents of the folder show in the viewing frame.
- To close an open folder, click the minus sign (-).

Viewing Database Information and Statistics

Table Manager provides information and statistics about the database.

Viewing Database Information

To view database information, click **Info** in the navigation pane. The following database information displays in the viewing pane:

Server Name	Port Number	Hostname	Host Address	Vendor	Db Release
POWDER2	5,000	powder2	134.70.64.172	sybase	Adaptive Server Enterprise/11.9.2.1/1108/IP/SWR 87...

The information that displays is described as follows:

Server Name	The name of the server where the database resides
Port Number	The number of the port that connects to the database server
Hostname	The name of the host where the database resides
Host Address	The address of the host where the database resides
Vendor	The name of the vendor for the database software
Db Release	The name, version number, and release number for the database software

Viewing Database Statistics

To view database statistics, click **Statistics** in the navigation pane. The following database statistics display in the viewing pane:

Device	Usage	Size (MB)	Free (MB)	% Used
dpipe_db0	data only	800	649	19
dpipe_log0	log only	1,200	1,049	13

The information that displays is described as follows:

Device	The name of the device in the database
Usage	The type of information that the device contains
Size (MB)	The size allotted for the device in the database in megabytes
Free (MB)	The amount of free space available in the device in megabytes
% Used	The amount of space used in the device

OVPI will stop collecting data when the database reaches its *high water mark*. By default, the high water mark is 90, which is 90% of its maximum storage capacity.

If the database reaches the high water mark, you must do one of the following to continue collecting data:

- Reduce the amount of data in the database.
- Increase the size of the database.
- Increase the high water mark.

The SNMP collector, the `mw_collect` command, has an option to override the database high water mark value. The option is `-w`. Refer to the *Performance Insight Reference Guide* for more information about the `mw_collect` command.

Monitoring Database Usage

By checking the % Used statistic, you can monitor the database space use. To see this statistic, click **Statistics** in Table Manager's navigation pane. In addition to keeping track of the percentage of available space that is left, you can monitor how quickly that space is used.

You can also monitor your database usage by watching the size of your data tables. If you are using up space quickly, observe which tables are the largest and how fast they are growing by sorting them. You can control their growth by modifying the number of days the data is kept. To do this, decide which data you really need to keep for an extended period of time versus the data you can discard after a day or two.

For example, you probably do not need to keep raw data longer than the time it takes to aggregate it, such as a day. Because this is the most voluminous data, it makes sense to manage it more carefully. In extreme cases, you may need to truncate a large table to keep the database from filling up.



When you truncate a table, you cannot recover the data that it contained.

To keep your database from filling up, you may want to make adjustments when the % Used value exceeds 70%. You may also want to delete tables you no longer need.

Viewing Table Information

You can view information about the various types of tables. The columns for data and internal tables differ from the columns for property tables. This section describes the columns that appear for the various types of tables along with the functions that let you modify how the information is displayed.

Viewing Functions

Table Manager lets you modify or locate the information that is displayed. These functions are as follows: sort in ascending or descending order, change the size of the columns, refresh, and find.

Sorting Information

To sort the information in the column in ascending order, place the cursor in the heading of the column you want to sort, then click.

Similarly, to sort the information in the column in descending order, place the cursor in the heading of the column you want to sort, hold down the **Shift** key, and click.


Changing Column Size

To change the size of a column:

- 1 Place the cursor on the border of the right side of the column.
- 2 When the pointer becomes a \leftrightarrow , press and hold the left mouse button, and move it until the column is the width you want.

Refreshing Information

You can refresh Table Manager's display. When Table Manager refreshes its display, it sorts the columns in the original order, sets the sizes of the columns back to the original sizes, and removes the highlighting from any highlighted rows.

To refresh the screen, click , or choose **Refresh** from the View menu.


Finding Alias Names

Table Manager lets you locate an alias name from the list of tables that are currently active.

To find an alias name from the current list:

- 1 From the Edit menu, choose **Find Alias Name**.

The Find Alias Name dialog box opens.



- 2 Type the alias name that you want to locate in the Alias Name box.
- 3 Click **Match Whole Word Only** to match a complete word.
- 4 Click **Match Case** to find or ignore a name with specific capitalization.
- 5 Click **Find** to locate the alias name.

Viewing Data Tables by Type or Category

You can view data tables by type or by category by opening the Data Tables folder, which is located in the navigation pane of Table Manager. The Data Tables folder contains the following folders:

- Type

This folder contains the data tables organized by the type of data. The sum type data is in the Summary folder with the data tables organized by the time type of the data.

- Category

This folder contains the data tables organized by the category field for the data. The categories appear in the order that you add them to the database.

For example, the following shows the results that are displayed in the viewing pane when you select router from the Category folder.

Category	Alias Name	SQL Name	Type	Time Type	Time Zone	Size (rows)	Size (KB)	Keep data for (days)
router	rate_router_system	Router_system_	rate	5-minute	local	1,294	538	2
router	rate_router_ipaddress	Router_ipaddress	rate	as-pollled	local	20	46	2
router	test_tbl_1	test_tbl_1	raw	hourly	local	30	64	30
router	router_system	router_system_	raw	5-minute	local	1,289	448	2
router	router_ipaddress	router_ipaddress	raw	as-pollled	local	30	32	2
router	hour_router_SHroutersla	SHroutersla	sum	hourly	local	3,615	1,132	30
router	day_router_SDroutersla	SDroutersla	sum	daily	local	250	80	90



The columns that you see may be different than the columns in this example. You can select which columns you want to view when the data table is displayed. See [Changing the Table Column Display on page 225](#) for details.

The columns that are displayed are described as follows:

- | | |
|------------|---|
| Category | This is the category for the table. The category is a standard code assigned to a group of tables by the OVPI reporting solution. It allows tables to be grouped together by the type of reporting the table supports. The last category added appears at the bottom of the list. |
| Alias Name | This is a logical, unique, descriptive name of the table in the database. |
| SQL Name | This is the name the SQL Server recognizes for the table. Use this name for any SQL queries. |

Is View	If selected, this option indicates that the selected table is actually a view created by querying one or more database tables.
Type	This is the type of data stored in this table. The available types are as follows: raw, rate, sum, forecast, asset, rank, baseline, lkeys, and upload. For a description of the types, see Task 3 on page 262 .
Time Type	This specifies the time type of the data. The available time types are as follows: as-pollled, sub-minute, 5-minute, hourly, daily, weekly, monthly, quarterly, and yearly.
Time Zone	This is the time zone for the table.
Size (rows)	This is the number of rows in the table.
Size (KB)	This is the size of the table in kilobytes.
Keep data for (days)	This is the number of retention days for the data in the table. Data is discarded when it is age + 1 days old. The available retention options are as follows: <ul style="list-style-type: none">• permanently—never deletes the data• <i>n</i>—keeps data for a specific number of days• default—keeps data for the default time period, which is the retention period for the time type. This value is set using the Days to Keep Data box in the Set Table Defaults dialog box. To change default data storage times, see Setting Data Table Storage Times on page 228.

Viewing Property Tables

To view existing property tables in the database, select **Property Tables** from Table Manager's viewing pane. The following example shows the results in the viewing pane when you select Property Tables.

Alias Name	SQL Name	Is View	Type	Is Standard	Size (rows)	Size (KB)
kst_user	kst_user	<input type="checkbox"/>	property	<input checked="" type="checkbox"/>	2	48
kst_managed_node	kst_managed_node	<input type="checkbox"/>	property	<input checked="" type="checkbox"/>	61	48
kst_dbstats_	kst_dbstats_	<input type="checkbox"/>	property	<input checked="" type="checkbox"/>	1	48
kst_comm_string	kst_comm_string	<input type="checkbox"/>	property	<input checked="" type="checkbox"/>	2	46
kst_snmp_profile	kst_snmp_profile	<input type="checkbox"/>	property	<input checked="" type="checkbox"/>	1	48
KReportPackProcLog	KReportPackProcLog	<input type="checkbox"/>	property	<input type="checkbox"/>	2	48
Customer_Property_Table	K_Customer	<input type="checkbox"/>	property	<input type="checkbox"/>	2	46
Location_Property_Table	K_Location	<input type="checkbox"/>	property	<input type="checkbox"/>	2	48
Node_Property_Table	K_Node	<input type="checkbox"/>	property	<input type="checkbox"/>	22	48
propview_Customer	KV_Customer	<input checked="" type="checkbox"/>	property	<input type="checkbox"/>	0	0
propview_Location	KV_Location	<input checked="" type="checkbox"/>	property	<input type="checkbox"/>	0	0
propview_Node	KV_Node	<input checked="" type="checkbox"/>	property	<input type="checkbox"/>	0	0
prop_DeviceResource_Card	K_DeveRes_Card	<input type="checkbox"/>	property	<input type="checkbox"/>	21	48
prop_DeviceResource_Device	K_DeveRes_Device	<input type="checkbox"/>	property	<input type="checkbox"/>	21	48
prop_DeviceResource_DeviceMemory	K_DeveRes_DeviceMem	<input type="checkbox"/>	property	<input type="checkbox"/>	41	48
prop_Property_CustomerLocationAggregate	K_DeveRes_CustLocAgg	<input type="checkbox"/>	property	<input type="checkbox"/>	4	48

The columns that are displayed are described as follows:

Alias Name	This is a logical, unique, descriptive name of the table in the database.
SQL Name	This is the name the SQL Server recognizes for the table. Use this name for any SQL queries.
Is View	If selected, this indicates that the selected table is actually a view created by querying one or more database tables.
Type	This is the type of data stored in the table. In this case, the value is always property .
Is Standard	If selected, this identifies a well-known table to the OVPI database.
Size (Rows)	This is the number of rows in the table.
Size (KB)	This is the size of the table in kilobytes.


Changing the Table Column Display

You can change whether a table column displays and change the order in which table columns are displayed.

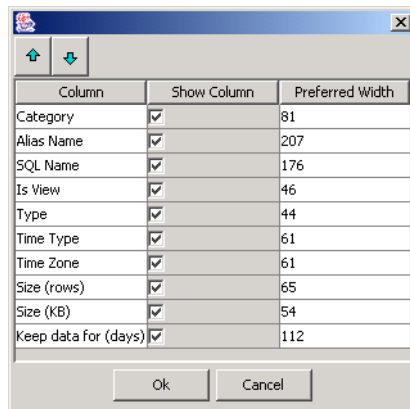
To change how table columns are displayed:



- 1 Right-click on any table column name.

A shortcut menu displays some of the available column headings, similar to the menu shown below.



- 2 Click **More** to display the entire list of headings, as in the following example.



- 3 To change the order of the columns, select a column name and click  or  until the name is in the desired position.
- 4 To prevent a column from displaying, click to clear the **Show Column** option next to the column name.
- 5 To display a column that is not currently displayed, select the **Show Column** option next to the column name.

- 6 To change the width of a column, do the following:
 - a Double-click in a column's Preferred Width box.
 - b Delete the current value, and type a new value.
 - c Type the new value.
- 7 Click **OK**.

Viewing Data and Internal Table Details

You can view detailed information about any data or internal table by double-clicking on the table in Table Manager's viewing pane.

For example, to view detailed information about the `rate_IR_property` table, you would do the following:

- 1 In Table Manager's navigation pane, expand the Data Tables folder, expand the Category folder, and then click on **Interface_Reporting**.

The viewing page displays all the tables associated with the `Interface_Reporting` category.

- 2 Double-click the `rate_IR_property` table in the viewing pane.

The Data Table Details dialog box opens as shown in [Figure 2](#).

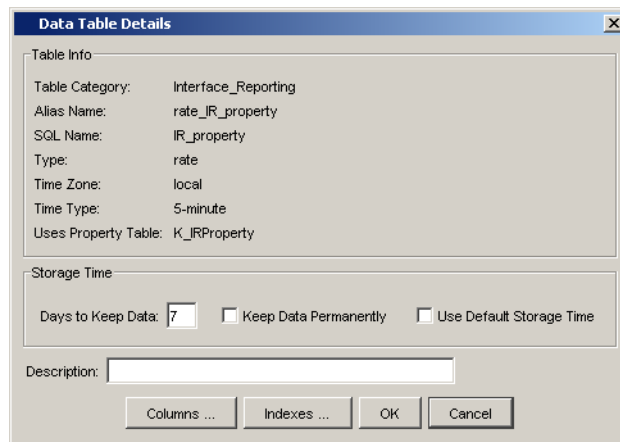


Figure 2 Data Table Details Dialog Box

The Data Table Details dialog box lets you do the following:

- View information about the table's columns.
- Set the storage time.
- Add a description.
- View information about the table's columns.
- View information about the corresponding indexes.

Viewing Data Table Column Information

The Table Info box in the Data Table Details dialog box ([Figure 2 on page 227](#)) shows the table information for the selected table. For data tables, the Table Info box also shows the name of the property table that this data or internal table uses.

Setting Data Table Storage Times

The Storage Time box in the Data Table Details dialog box ([Figure 2 on page 227](#)) shows the current number of days that the OVPI database keeps the data; however, you can change the number of days to keep the data.

To change the storage time for data, do one of the following:

- Type the desired number of days in the Days to Keep Data box.
OVPI will age out the data from the table after the specified number of days has elapsed.
- Select **Keep Data Permanently**.
OVPI will not age out the data from the table.
- Select **Use Default Storage Time**.
OVPI keeps the data for the number of days that corresponds to the default retention time for the time type of the data.

To view or modify the default retention time for a particular table time type, see [Setting Table Defaults on page 277](#).

Adding a Data Table Description

The Description box in the Data Table Details dialog box (Figure 2 on page 227) lets you add a description of the table. Enter text in the Description box, and click **OK**. The description cannot exceed 255 characters.

Viewing Data Table Column Information

To view the details about the columns in a table, click **Columns** in the Data Table Details dialog box (Figure 2 on page 227). The Columns window opens.

Alias Name	SQL Name	Data Type	SQL Type	Asset Column	Object Identifier	Is Nullable	Default Value
target_name	target_name	snmp_char_string	varchar(64)	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
received_ts	received_ts	unix_time	datetime	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
request_ts	request_ts	unix_time	datetime	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
delta_time	delta_time	snmp_timeticks	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
user_name	user_name	snmp_char_string	varchar(8)	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
table_key	table_key	snmp_char_string	varchar(128)	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
report_type	report_type	snmp_char_string	varchar(1)	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
received_usec	received_usec	snmp_gauge	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
cpmCPUTotalIndex	cpmCPUTotalIndex	snmp_integer	float	<input type="checkbox"/>	1.3.6.1.4.1.9.9.109.1.1.1.1.1	<input checked="" type="checkbox"/>	
cpmCPUTotalPhysicalIndex	cpmCPUTotalPhysicalIndex	snmp_integer	float	<input type="checkbox"/>	1.3.6.1.4.1.9.9.109.1.1.1.1.2	<input checked="" type="checkbox"/>	
cpmCPUTotal5sec	cpmCPUTotal5sec	snmp_gauge	float	<input type="checkbox"/>	1.3.6.1.4.1.9.9.109.1.1.1.1.3	<input checked="" type="checkbox"/>	
cpmCPUTotal1min	cpmCPUTotal1min	snmp_gauge	float	<input type="checkbox"/>	1.3.6.1.4.1.9.9.109.1.1.1.1.4	<input checked="" type="checkbox"/>	
cpmCPUTotal5min	cpmCPUTotal5min	snmp_gauge	float	<input type="checkbox"/>	1.3.6.1.4.1.9.9.109.1.1.1.1.5	<input checked="" type="checkbox"/>	
dsi_key_id	dsi_key_id_	numeric	numeric(10,0)	<input type="checkbox"/>		<input type="checkbox"/>	
dsi_agg_type	dsi_agg_type	seq_key	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
ta_sysuptime	ta_sysuptime	snmp_timeticks	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
ta_period	ta_period	unix_time	datetime	<input type="checkbox"/>		<input type="checkbox"/>	
ta_samples	ta_samples	snmp_integer	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	


The Columns window shows the following information for each column in the table.

- Alias Name** This is a logical, unique, descriptive name of the column in the table.
- SQL Name** This is the SQL name of the column in the table.
- Data Type** This is the OVPI data type for the column. For a description of these types, see [Task 4 on page 264](#).
- SQL Type** This is the database type of the data for the column. For a description of these types, see [Task 4 on page 264](#).
- Asset Column** If selected, this specifies that the column contains an asset property for use in an asset summary.

- Object Identifier This specifies the path through the MIB tree to the selected object represented as a string of numbers separated by periods. Only columns associated with SNMP collectors will have a MIB OID value; otherwise, this column can be blank.
- Is Nullable If selected, this specifies that the column can be NULL.
- Default Value This identifies the default value for the column.

Viewing Data Table Index Information

To view the details about the indexes for the table, click **Indexes** in the Data Table Details dialog box (Figure 2 on page 227). The Indexes window opens.



The Indexes window shows the following information about each index for the table.

- Index Name This is the name of the index for the table.
- Unique If selected, this specifies that the index is unique.
- Clustered If selected, this specifies that the index is clustered.
- Column 1 The first SQL Name of the column used for building the index.
- Column 2 The second SQL Name of the column used for building the index.



A column appears for every element in the index.

Viewing Property Table Details

You can view detailed information about any property table by double-clicking on the table in Table Manager's viewing pane.

For example, to view detailed information about the Customer Key table, you would do the following:

- 1 In Table Manager's navigation pane, click **Property Table**.

The viewing page displays all the property tables.

- 2 Double-click the **Customer Key Table** table in the viewing pane.

The Property Table Details dialog box opens as shown in [Figure 3](#).

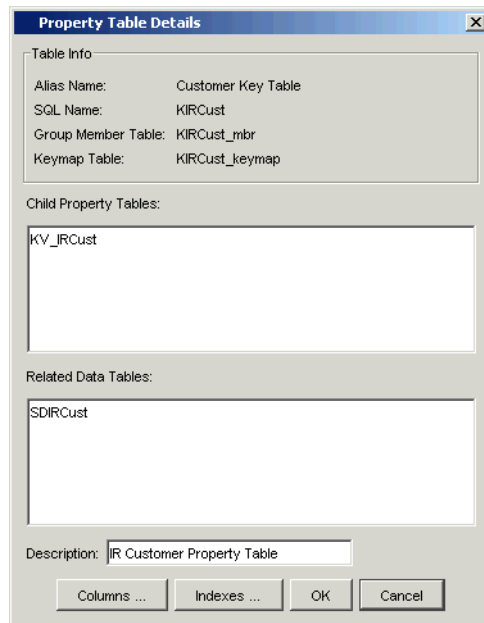


Figure 3 Property Table Details Dialog Box

The Property Table Details dialog box lets you do the following:

- View information about the table.
- Add a description.
- View information about the table's columns.

- View information about the corresponding indexes.

Viewing Property Table Information

The Table Info box in the Property Table Details dialog box ([Figure 3 on page 231](#)) contains the following information for property tables.

Alias Name	This is a logical, unique, descriptive name of the table in the database.
SQL Name	This is the SQL name of the table.
Group Member Table	This is the name of the table that contains the lists of managed objects that are members of enumerated list groups for this property table.
Keymap Table	This is the name of the table that contains a cross-reference of the property table contents for the local server to the property table contents of a remote server.

The dialog box also lists the following information:

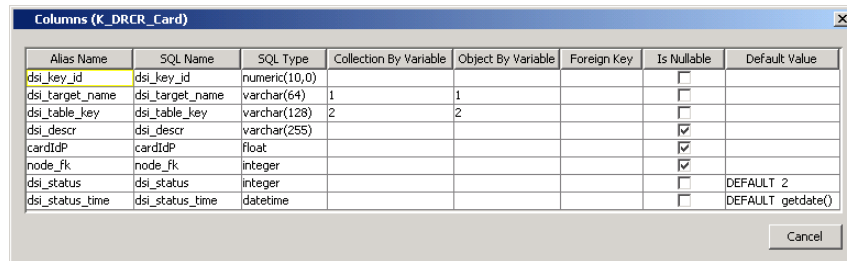
Child Property Tables	This is a list of property tables that reference this property table.
Related Data Tables	This is a list of data tables that use this property table.

Adding a Property Table Description

The Description box in the Property Table Details dialog box ([Figure 3 on page 231](#)) lets you add a description of the table. Enter text in the Description box, and click **OK**. The description cannot exceed 255 characters.

Viewing Property Table Column Information

To view the details about the columns in a table, click **Columns** in the Data Table Details dialog box (Figure 3 on page 231). The Columns window opens.




Alias Name	SQL Name	SQL Type	Collection By Variable	Object By Variable	Foreign Key	Is Nullable	Default Value
dsi_key_id	dsi_key_id	numeric(10,0)				<input type="checkbox"/>	
dsi_target_name	dsi_target_name	varchar(64)	1	1		<input type="checkbox"/>	
dsi_table_key	dsi_table_key	varchar(128)	2	2		<input type="checkbox"/>	
dsi_descr	dsi_descr	varchar(255)				<input checked="" type="checkbox"/>	
cardIdP	cardIdP	float				<input checked="" type="checkbox"/>	
node_fk	node_fk	integer				<input checked="" type="checkbox"/>	
dsi_status	dsi_status	integer				<input type="checkbox"/>	DEFAULT 2
dsi_status_time	dsi_status_time	datetime				<input type="checkbox"/>	DEFAULT getdate()

The Columns window shows the following information for each column in the table.

Alias Name	This is a logical, unique, descriptive name of the column in the table.
SQL Name	This is the SQL name of the column in the table.
SQL Type	This is the database type of the data for the column. For a description of these types, see Task 4 on page 264 .
Collection By Variable	This specifies the order of the column in the unique key for the collector in the table. If this column is blank, it indicates that this column is not part of the collector key for the table.
Object By Variable	This specifies the order of the column in the unique key for a row in the table. If this column is blank, it indicates that this column is not part of the object key for the table.
Foreign Key	This identifies the parent table and the column name that contain the data if the column is a foreign key. The format is <i>parent_table.column_name</i> .
Is Nullable	If selected, this specifies that the column can be NULL. If there is a value in the Collection By Variable column or the Object By Variable column for the table, then the Is Nullable column cannot be selected.
Default Value	This identifies the default value for the column.

Viewing Property Table Index Information

To view the details about the indexes for a table, click **Indexes** in the Property Table Details dialog box (Figure 3 on page 231). The Indexes window opens.



The Indexes window shows the following information about each index for the table.

- Index Name** This is the name of the index for the table.
- Unique** If selected, this indicates that the index is unique.
- Clustered** If selected, this indicates that the index is clustered.
- Column 1** The first SQL Name of the column used for building the index.
- Column 2** The second SQL Name of the column used for building the index.



Every property table has at least two unique indexes. It builds the indexes for the columns referenced by both the collection by-variable and the object by-variable. If the collection by-variable and the object by-variable have the same columns in the index, then only one index will appear in the list. If the by-variables have different columns in the indexes, then at least two indexes will appear in the list. Also, a column appears for every element in an index.

Displaying Table Data

Table Manager's **Display Table Contents** option lets you examine the data in the OVPI database; you can view the data for data, property, and internal tables in a tabular format. These reports are useful for reviewing lists of values displayed in different orders or grouped for comparison. They are especially useful for exception reporting, for example, by showing exceeded thresholds.

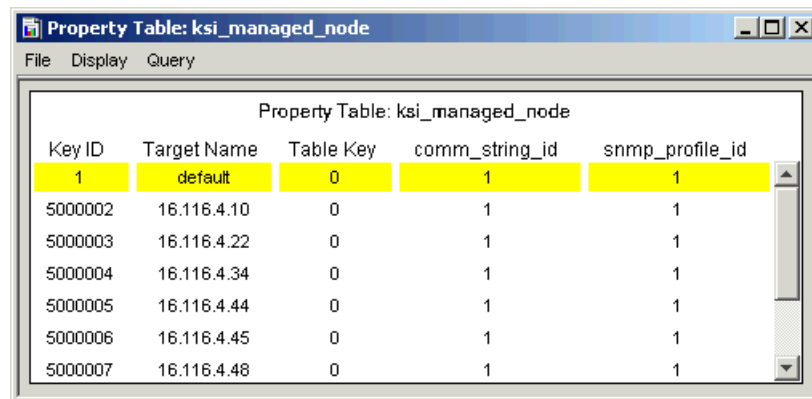
You can change the type of data displayed and its appearance, and then save it to the OVPI report formats (.rep or .srep). You can view the saved reports using any of the following OVPI client applications: Web Access Server, Report Builder, and Report Viewer. Additionally, you can print these reports, and export the data to different formats (CSV and ASCII) for transfer to spreadsheets or other applications.

Displaying Reports

To the select table data and generate a report from it:

- 1 Select the table by highlighting the appropriate row.
- 2 From Table Manager's View menu, select **Display Table Contents**.

A window opens, which displays the data for the selected table in tabular format. The name of the table appears in the window's title bar.



Key ID	Target Name	Table Key	comm_string_id	snmp_profile_id
1	default	0	1	1
5000002	16.116.4.10	0	1	1
5000003	16.116.4.22	0	1	1
5000004	16.116.4.34	0	1	1
5000005	16.116.4.44	0	1	1
5000006	16.116.4.45	0	1	1
5000007	16.116.4.48	0	1	1

Using the Menus

The Data Table window also contains a menu bar with the follow menus: File, Display, and Query. Using the menu options, you can modify how the data will appear in the report.

Table 2 lists the options on the File menu.

Table 2 File Menu Options

Option	Description
Print (CTRL+P)	Opens the Print dialog box, which lets you print a report. You can select a printer, the portion of the report to print, and the number of copies to print.
Print Setup	Opens the Page Setup dialog box, which lets you specify paper size, the location of the paper in the printer, the report's position on the page, and the printing area of the page.
Save (CTRL+S)	Saves the table data to its current filename. (When you save a report for the first time, the Save dialog box will open.)
Save As	Opens the Save dialog box, which lets you save the table data as a .rep or .srep report.
Export	Opens the Export Data Table dialog box, which lets you export data from a report to a CSV (.csv) or ASCII (.txt) file.
Close	Closes the window in which the data table is displayed.

Table 3 lists the options on the Display menu.

Table 3 Display Menu Options

Option	Description
Statistic Properties	Displays a submenu from which you can select a statistic. Once you select a statistic, the Statistic Properties dialog box opens, which lets you change how the statistic's properties display in the report. You can change the font, font size, bold the text, italicize the text, change the text color, change the alignment for the statistic's column, and change the alignment for the statistic data.
Cell Properties	Displays the Cell Properties dialog box, which lets you change various properties for the table cells.
Column Label Properties	Displays the Column Label Properties dialog box, which lets you change various properties for the table columns.
Row Label Properties	Displays the Row Label Properties dialog box, which lets you change various properties for the table rows.

Table 4 lists the options on the Query menu.

Table 4 Query Menu Options

Option	Description
Select Statistics	Displays the Select Statistics dialog box, which lets you select the statistics for the data displayed.
Order Statistics	Displays the Order Statistics dialog box, which lets you change the order in which the statistics will appear in the table.
Sort Order	Displays the Sort Order dialog box, which lets you select the sort order of the data returned from the database.

Table 4 Query Menu Options (cont'd)

Option	Description
Change Time Period	Displays the Set Time Period dialog box, which lets you select the time period for the data that will be displayed. This option is only available for data tables.
Add/Remove Constraint Values	Displays the Change Constraints dialog box, which lets you add, modify, and remove constraints and clauses.
Change Constraint Values	Displays the Change Constraint Values dialog box, which lets you change only the value for a constraint.
Select Nodes/Interfaces	Displays the Select Nodes/Interfaces wizard, which lets you specify the nodes and interfaces from which the data is taken to generate the report.
Change Maximum Rows	Displays the Change Maximum Number of Rows dialog box, which lets you constrain the amount of data that is returned and subsequently displayed. The default number of rows returned is 50.

Managing Data

You can print, save, and export table data.

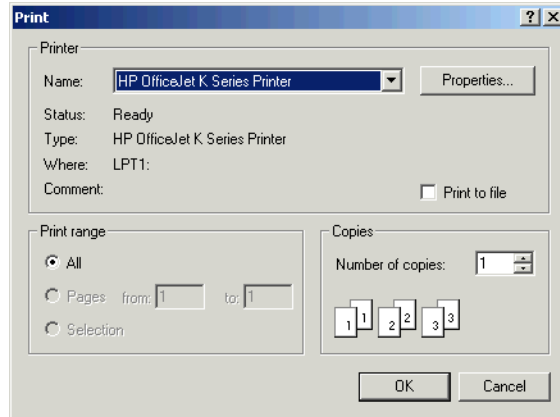
Printing Table Data

Using the Print or Print Setup options on the File menu, you can print table data or set your report printing preferences.

To print data:

- 1 Do either of the following:
 - From the File menu, click **Print**.
 - Press **CTRL+P**.

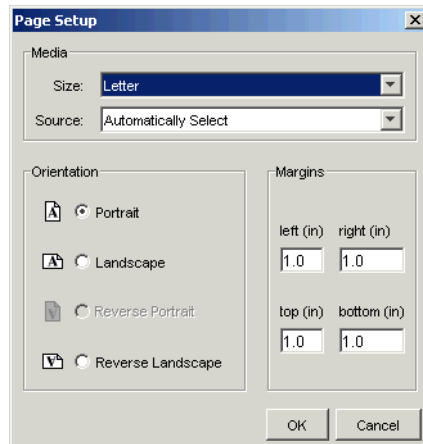
The Print dialog box opens.



- 2 Make any changes necessary, and click **OK**.

To set data printing preferences:

- 1 From the File menu, click **Print Setup**. The Page Setup dialog box opens.



- 2 Make any changes necessary, and click **OK**.

Saving Table Data

After you display table data, you can save it to one of the following OVPI report formats:


- Dataset (.srep)
- Report Definition (.rep)

You can view the saved report using any of the following OVPI client applications: Web Access Server, Report Builder, and Report Viewer.

To save a table data:

- 1 Do either of the following:
 - To save a report with its current file name, choose **Save** from the File menu *or* press **CTRL+S**.
 - To save a report with a different file name or format, choose **Save As** from the File menu.

If you chose **Save As** from the File menu, or if you have never saved the report, the Save dialog box opens.



- 2 Do one or both of the following:
 - Enter the file name in the File Name box.
 - From the Files of Type list, select a report format.
- 3 Click **Save**.

Exporting Table Data

You can export the data shown in the Data Table window to a comma separated value (CSV) file or a text (tab delimited) file.

To export the data table:

- 1 *Recommended but not required:* Select **Change Maximum Rows** from the Query menu, type 0 in the Max Rows box, and click **OK**. This way you will make sure that the entire dataset that you have constrained will be utilized.
- 2 From the File menu, select the **Export** option. [Figure 4](#) shows the Export dialog box.

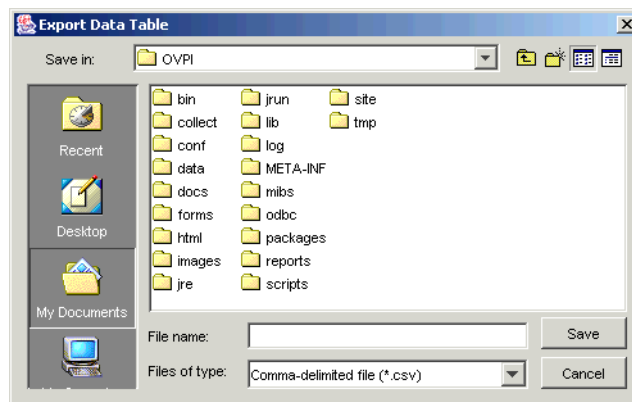


Figure 4 Export Data Table Dialog Box

- 3 Navigate to the location where you want to save the file.
- 4 Select a file type from the Files of type list.
- 5 In the Filename box, type a filename.
- 6 Click **Save** to export the data.

Formatting Data

You can change how table data displays by using the options on the Display menu to change the format of the data's statistics, cells, columns, and rows.

Formatting Statistics

To change the format of how the statistic values will appear in a table column:

- 1 From the Display menu, select **Statistic Properties**.
- 2 Select the statistic that you want to change. **Figure 5** shows the **Statistic Properties** dialog box.

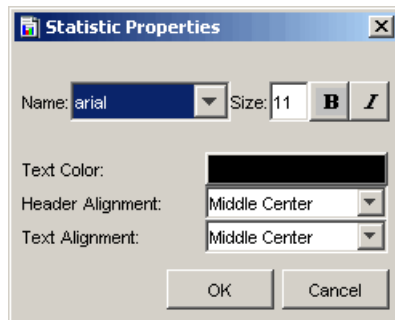


Figure 5 **Statistic Properties Dialog Box**

- 3 Do the following:
 - d Select a font from the Name list to change the font used for the statistic values.
 - e Type a size for the font used for the statistic values in the Size box.
 - f Click the *Bold* button to bold the statistic values.
 - g Click the *Italics* button to italicize the statistic values.
 - h Click the Text Color box to display the Select a Color dialog box, which lets you choose a color for the statistic values.
 - i To position the statistic's name in the heading of the table column, select an alignment option from the Header Alignment list.
 - j To position the statistic text that appears in the column below the statistic's heading, select an alignment option from the Text Alignment list.
- 4 Click **OK** to make your changes take effect.

Formatting Cells

To change the properties of table cells:

- 1 From the Display menu, select **Cell Properties**. [Figure 7](#) shows the Cell Properties dialog box.

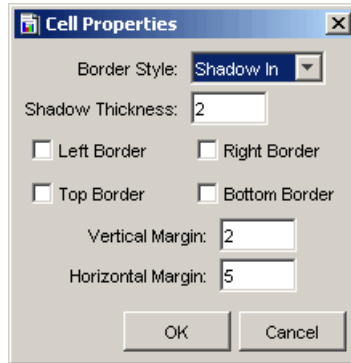


Figure 6 Cell Properties Dialog Box

- 2 Do any of the following:
 - a Select a value from the Border Style list to outline the table cells. By default, this is set to **Shadow In**.
 - In order for the outline to appear, you must select at least one of the following options: Left Border, Right Border, Top Border, or Bottom Border.
 - b Type a value in the Shadow Thickness box to control the width of the border style that you selected in [Step a](#).
 - c Select any of the following:
 - Left Border. Displays the border style, which you selected in [Step a](#), on the left side of each cell.
 - Right Border. Displays the border style, which you selected in [Step a](#), on the right side of each cell.
 - Top Border. Displays the border style, which you selected in [Step a](#), on the top of each cell.
 - Bottom Border. Displays the border style, which you selected in [Step a](#), on the bottom of each cell.

For example, if you selected all of the border options with the Border Style option set to Shadow Out, the table's cells would appear similar to the following:

Key ID	Target Name	Table Key	comm_string_id	snmp_profile_id
1	default	0	1	1
5000002	16.116.4.10	0	1	1
5000003	16.116.4.22	0	1	1
5000004	16.116.4.34	0	1	1
5000005	16.116.4.44	0	1	1
5000006	16.116.4.45	0	1	1
5000007	16.116.4.48	0	1	1

- d In the Vertical Margin box, type a value to increase or decrease the amount of space between the rows of cells. For example, type 10 to increase the amount of space, for example:

Key ID	Target Name	Table Key	comm_string_id	snmp_profile_id
1	default	0	1	1
5000002	16.116.4.10	0	1	1
5000003	16.116.4.22	0	1	1
5000004	16.116.4.34	0	1	1
5000005	16.116.4.44	0	1	1

- e In the Horizontal Margin box, type a value to increase or decrease the amount of space between the table columns.

- 3 Click **OK** to make your changes take effect.

Formatting Columns

To make a report more readable, you can change the order in which the columns appear in a report, the width of each column, and the properties for the column.

Changing Column Order

You can make your report much easier to understand if you put the left-hand columns in the report in the same sequence as your sort order. For example, if you are sorting a report on line utilization by utilization level and time of day (Received Time), put the Utilization column in the left-hand column of the report and Received Time in the column to the right of Utilization by dragging each column to the left of the fixed Time Period column.

To change the order of table columns:

- 1 Place the pointer on the column heading you want to move.
- 2 Press the left mouse button, and drag the pointer over a column in either the left or right direction.

As you move the pointer, the column that you drag the current column over moves (to the left or to the right, depending on the direction in which you are dragging the pointer) past the column you are dragging.

- 3 Release the left mouse button.
- 4 Repeat steps 1 through 3 to move another column.

Changing Column Width

To change column width:

- 1 Place the pointer between the column you want to change the width of and the column next to it. The pointer becomes a double-ended arrow.
- 2 Drag the double-ended arrow to adjust the column to the size you want.

Changing Column Properties

To change the properties of table columns:

- 1 From the Display menu, select **Column Label Properties**. Figure 7 shows the Column Label Properties dialog box.

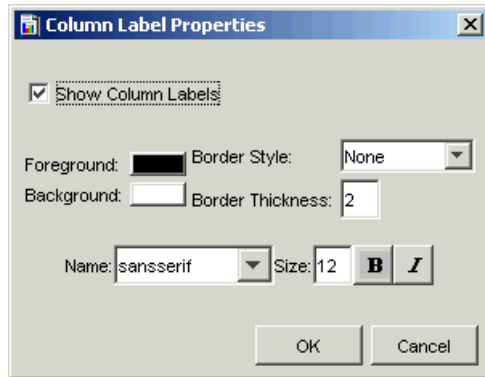


Figure 7 Column Label Properties Dialog Box

- 2 Do any of the following:
 - a Click the Foreground box to display the Select a Color dialog box, which lets you change the color for the column names. By default, the foreground color is black.
 - b Click the Background box to display the Select a Color dialog box, which lets you change the color for the boxes that contain the column names. By default, the background color is white.
 - c Select a value from the Border Style list to outline the boxes around the column names.

By default, this is set to **None**, which shows no border for the column heading boxes. If you choose **Shadow Out**, for example, the column heading box looks like the following:

Target Name

- d If you selected a border style in Step c, the Border Thickness option controls the width of the border. By default, this is 2.
- e Select a font from the Name list to change the font used for the column heading.
- f Type a size for the font used for the column heading in the Size box.

- g Click the *Bold* button to bold the column heading text.
 - h Click the *Italics* button to italicize the column heading text.
 - i Click to clear the **Show Column Labels** box to prevent the column headings from displaying or select the **Show Column Labels** box to display the column headings.
- 3 Click **OK** to make your changes take effect.

Formatting Rows

- 1 From the Display menu, select **Row Label Properties**. Figure 8 shows the Row Label Properties dialog box.

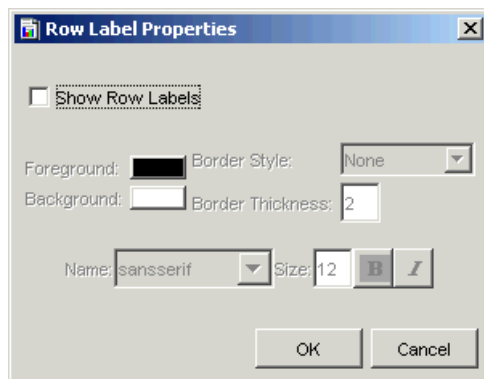


Figure 8 Row Label Properties Dialog Box

- ▶ The Row Label Properties dialog box is almost identical to the Column Properties dialog box shown in Figure 7.
- 2 See the steps for the procedure in [Changing Column Properties on page 246](#) to use all options except the Show Row Labels option.
 - 3 Select the **Show Row Labels** box to display the row headings or click to clear the **Show Row Labels** box to prevent the row headings from displaying.
 - 4 Click **OK** to make your changes take affect.

Changing Table Queries

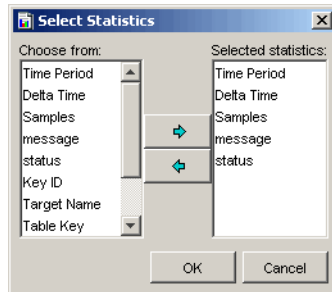
You can change the information that a report contains by specifying different database query information. Using the Data Table window's Query menu, you can do the following: select statistics, change the time period, add and remove constraints, change constraint values, select nodes, interfaces, or both, and change the maximum number of rows returned.

Selecting Statistics

To change the statistics that are displayed for certain table data:

- 1 Select Properties from the Query menu.

The Select Statistics dialog box opens, which lists the statistics currently displayed in the report.



- 2 To select a statistic, double-click a statistic in the Choose from box **or** select the statistic in the Choose from box and click .
- 3 To remove a statistic from the Select Statistics box, double-click the statistic in the Selected statistics box or select the statistic in the Selected statistics box, and click .
- 4 Click **OK** when you have selected the statistics that you want.

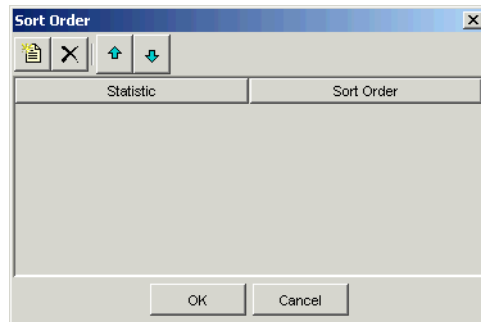
Sorting the Data in the Report





You can select the sort order of the data returned from the database. You may want to sort your reports using a different field, or in ascending rather than descending order.

OVPI's data is related to time; therefore, the default sort order for the data is by target name table keys for a particular target name, and received timestamp, newest to oldest, for a particular table key. Because most reports will be run for a particular device and a specific key for that device, the timestamp is the significant sort field.

To sort statistics:

- 1 From the Query menu, select **Sort Order**. The Set Sort dialog box opens.



- 2 To add a statistic to the list, do the following:
 - a Click the  button. A Statistic and Sort Order box appear.
 - b Click the box, and select the statistic you want to add from the list of statistics.
- 3 To change the order in which the data is sorted, click the Sort Order box next to each statistic, and select one of the following from the list:
 - Click ASC to sort the statistic in ascending order.
 - Click DESC to sort the statistic in descending order.
- 4 To delete a value from the list, click the  button.
- 5 To rearrange the order of the statistics in the list, do *one* of the following:
 - Click the  button to move a statistic up in the list.
 - Click the  button to move a statistic down in the list.
- 6 Click **OK**.

Changing the Time Period

You can change the time period for data displayed in a report. You can select an absolute period (specific dates) or a relative time period (a period of time relative to the time you view the data).

Selecting an Absolute Time


An absolute time period comprises a specific start and end time.

To select an absolute time period:

- 1 To change the time period, select **Change Time Period** from the Query menu. The Set Time Period dialog box (Figure 9) opens.
- 2 Click **Use Absolute Time**.

The initial start and end times in the Set Time Period dialog box represent the range of available data in the database.

- 3 Click the Browse button next to the Start Time box or the End Time box. A window similar to this one opens:



- 4 Click the down arrow to display a calendar, similar to the one below:



- 5 Click the Month and Year arrows to select the month and year you want.
- 6 Click the *Hour*, *Minutes*, *Seconds*, and *Time of Day* arrows to select the time.
- 7 Click a date in the calendar.

After you click a date, the calendar window closes, and the selected date and time display in the Start Time (or End Time) window. You can click the down arrow to reopen the calendar if you want to make changes.

- 8 Click **OK**.

Selecting a Relative Time Period

A relative time period is a period of time relative to the time you view the data.

To select a relative time period:

- 1 To change the time period, select **Change Time Period** from the Query menu. The Set Time Period dialog box (Figure 9) opens.

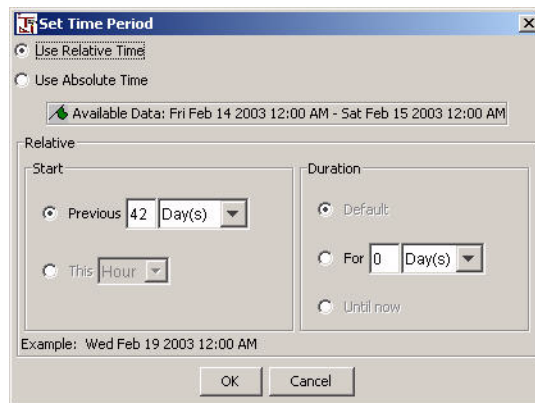


Figure 9 Set Time Period Dialog Box

- 2 Click **Use Relative Time**.
- 3 In the Start box, do *one* of the following:
 - To select a current time period, click the **This *time period*** button. Click the ***time period*** arrow and select an option. For example, if the time period is This Week, the report displays data for the current week.
 - To select a previous time period, click the **Previous *n time period*** button. Enter a number in the box, and then click the ***time period*** arrow, and select an option. For example, if the time period is Last 1 Week(s), the report displays data for the last full week (Sunday to Sunday) prior to the current date.
- 4 In the Duration box, do *one* the following:

- Click **Default** to use the default duration period.
 - Click **For *n* time period** to set the duration. Enter a number in the box, and then click the **time period** arrow, and select an option.
 - Click **Until Now** to display data for selected time period up until the current time. For example, if it is Thursday, and you select a time period of 1week, the report will display data for the prior week (Sunday to Sunday) through Thursday.
- 5 Click **OK**.

Adding Constraints

A *constraint* specifies how the data for a graph or table is retrieved from the database. You can create a basic or a custom constraint. A *basic constraint* comprises a statistic, an operator, and a value. A *custom constraint* can contain any valid SQL expression, and there are no limitations placed on the structure of the constraint as with basic constraints.

Add a Basic Constraint

To add a basic constraint:

- 1 From the Query menu, select **Add/Remove Constraint**. [Figure 10](#) shows the Change Constraints dialog box.




Figure 10 Change Constraints Dialog Box

- 2 Right-click **Constraint List**, and click **Add Constraint** on the shortcut menu. [Figure 11](#) shows the Add a Constraint dialog box.

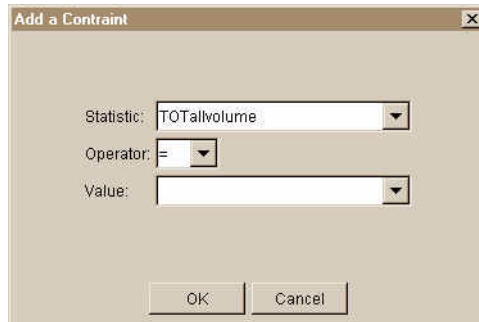


Figure 11 Add a Constraint Dialog Box

- 3 In the Add a Constraint dialog box, do the following:
 - a Click the Statistic arrow, and select a statistic.
 - b Click the Operator arrow, and select an operator.
 - c In the Value box, type a value, or you can select a parameter, if available. (Refer to *Performance Insight Guide to Building and Viewing Reports* for more information about creating a parameter.)
- 4 Click **OK**.

Adding a Custom Constraint

A custom constraint can contain any valid SQL expression; there are no limitations placed on the structure of the custom constraint.

To add a custom constraint:

- 1 From the Query menu, select **Add/Remove Constraint**. The Change Constraints dialog box opens. See [Figure 10](#).
- 2 Right-click **Constraint List**, and click **Add Constraint** on the shortcut menu. The Add a Constraint dialog box opens. See [Figure 11](#).

- 3 Right-click **Constraint List**, and click **Add Custom Constraint** on the shortcut menu. **Figure 12** shows the Add a Custom Constraint dialog box.




Figure 12 Add Custom Constraint Dialog Box

- 4 In the Label box, type the constraint's label.
- 5 To create the constraint's expression in the *Expression* box, do the following:
 - a Click the *Statistics* arrow, select a currently-defined statistic, and click **Insert**.
 - b Click any operator.
 - c Enter text.
- 6 Click **OK**.

Adding an AND or OR Clause

The difference between an AND clause and an OR clause is how the constraint will be processed. The database will process the list of constraints from top to bottom, applying boolean logic to return a list of matching rows.

To add an AND or OR clause:

- 1 From the **Query** menu, select **Add/Remove Constraint**. The Change Constraints dialog box opens. See **Figure 10**.
- 2 Right-click **Constraint List**, and click **Add And** or **Add Or** on the shortcut menu.

An AND or OR clause is added to the list of constraints.

Changing Constraint Values

To change only the value of a constraint:

- 1 From the Query menu, select **Change Constraint Values**. The Change Constraint Values dialog box opens. See [Figure 13](#). The Change Constraint Values dialog box opens.

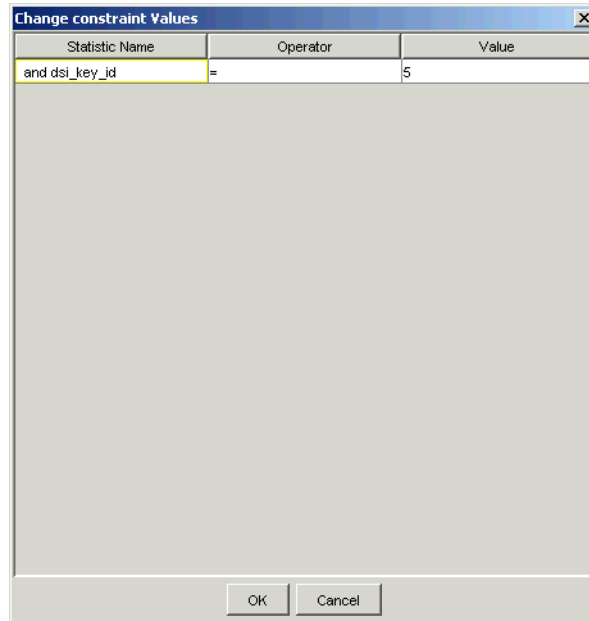


Figure 13 Change Constraint Values Dialog Box

➤ When the value of a constraint is a parameter, the parameter name appears in italics in the Value box. You must edit the parameter. To modify a parameter, refer to the *Performance Insight Guide to Building and Viewing Reports*.

- 2 Double-click the **Value** box, and type a new value.
- 3 Click **OK**.

Modifying Constraints

To modify all constraint properties:

- 1 From the Query menu, select **Add/Remove Constraint**. The Change Constraints dialog box opens. See [Figure 10](#).
- 2 Click the constraint you want to modify.
Its properties and values appear in the lower-half of the dialog box.
- 3 Change the values you want by clicking in any property's Value box.
- 4 Type or select a new value.
- 5 Click **OK** to make your changes take effect.

Deleting Constraints and Clauses

To delete a constraint or clause:

- 1 From the Query menu, select **Add/Remove Constraint**. The Change Constraints dialog box opens. See [Figure 10](#).
- 2 Do one of the following:
 - Right-click the clause you want to delete, and click **Delete Clause** on the shortcut menu.
 - Right-click the constraint you want to delete, and click **Delete Constraint** on the shortcut menu.

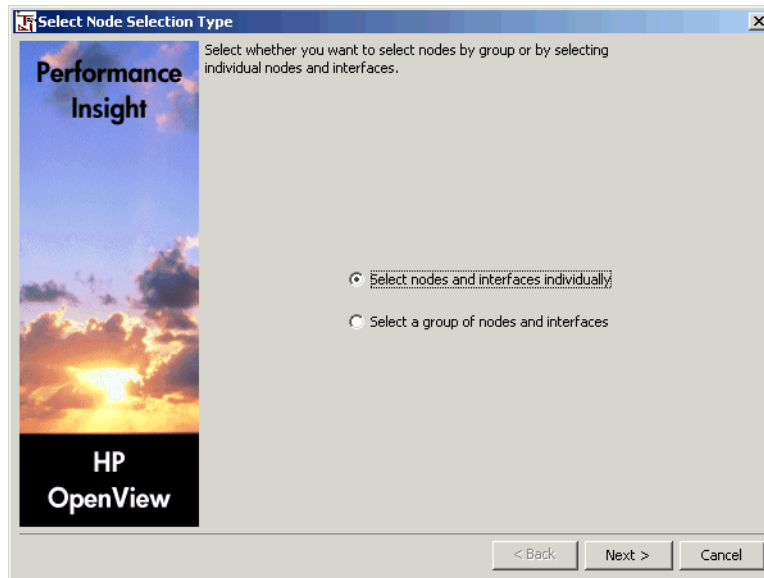
Selecting Nodes and Interfaces

When you select data for your report, you can specify the nodes and interfaces from which the data should be selected. You can change the selected nodes and interfaces by using Select Nodes/Interfaces wizard.

To change selected nodes and interfaces:

Task 1: Start the Select Nodes/Interfaces Wizard.

- 1 From the Query menu, select **Select Nodes/Interfaces**. The Select Nodes Selection Type page opens.





- 2 Select *one* of the following options:
 - **Select nodes and interfaces individually.** The Select Nodes/Interfaces page opens. Go to [Task 2](#).
 - **Select a group of nodes and interfaces.** The Select Group page opens. Go to [Task 3 on page 258](#).
- 3 Click **Next**.

Task 2: Select nodes and interfaces individually.

Do *one* of the following:

- To collect data from all nodes, select **Use All Nodes/Interfaces**, and then click **Next**.
- To collect data from only certain nodes, select **Use Selected Nodes/Interfaces**. In the Choose from box, each node is represented by a folder.

To select nodes and interfaces, do the following:

- a** Expand a node's folder to display its interfaces.
- b** To select a node or an interface, do *one* of the following:
 - Double-click a node or an interface in the Choose from box.
 - Select a node or an interface in the Choose from box, and click .
- c** To remove a node or an interface, do *one* of the following:
 - Double-click a node or an interface in the Selected Nodes/Interfaces box.
 - Select a node or an interface in the Selected Nodes/Interfaces box, and click .
- d** When you have selected the nodes and interfaces you want, click **OK**
- e** Click **Next**.

Task 3: Select a group of nodes and interfaces.

- 1** Select a group from the Group box. The Managed Objects For Group box displays the members of the selected group.
- 2** Click **Next**.

Task 4: Review the Nodes/Interfaces Selection Summary.

- 1** Review the summary.
- 2** If necessary, click **Back** to return to make a change. After you make any changes, click **Next** until you return to the summary.
- 3** Click **Finish**.



You can defer node selection until a report is viewed. For more information about how to defer node selection, refer to the *Performance Insight Guide to Building and Viewing Reports*.

Changing the Maximum Number of Rows Returned

The database can contain thousands of data entries, which span gigabytes of disk space. Thus, you can constrain the amount of data that is returned and subsequently displayed. The default number of rows returned is 50.

The data returned depends on the nature of the query. For example, if you specify in the query to sort the data by time in descending order, OVPI displays the most recent 50 rows of data.

Depending on the type of report you are viewing and how its data is distributed, 50 rows may not be enough data. You can change the number of rows displayed using the shortcut menu or a report's Components tree.

To change the number of rows using the shortcut menu:

- 1 From the Query menu, select **Change Maximum Rows**. [Figure 14](#) shows the Change Maximum Number of Rows dialog box.




Figure 14 Change Maximum Number of Rows Dialog Box

- 2 In the Max Rows box, enter the number of rows you want, and then click **OK**.

Creating Tables

Using Table Manager, you can create custom data tables and property tables. You can create collectible data tables; however, if you want to create an SNMP collectable table, use the MIB Browser to create it. See [Creating a Collection Table on page 400](#) for more information.

A *collectable data table* is a table that stores the data collected from a datapipe. If you want to create other tables such as summary, baseline, or forecast tables, see [Creating Data Tables on page 260](#). If you want to create a property table, see [Creating Property Tables on page 275](#).

Tables have the following requirements:

- The maximum width is 1,962 characters, which is the total character count for all of its columns.
- The maximum number of columns is 250. Note that there are 13 default columns that you must allow for when you select the objects for the data table. If you choose to define a property table, there are seven (7) default columns that you must allow for when you select the objects for it.
- A data table must have a corresponding property table. You can select an existing property table from a list or create a new property table. If you create a new property table, you can use the default by-variables or select different by-variables.


This section contains the basic information for creating a table. If you want additional information about creating the tables, refer to the *Performance Insight TEEL Reference Guide*. This document provides information about the attributes of tables and creating TEEL files.

Creating Data Tables

To create a custom data table, complete the following tasks:

Task 1: Start the Table Creation wizard.

From the Management Console, do one of the following to start the Table Creation wizard.

- Click , and select **Data Table** from the menu.
- From the File menu, select **New**, and then select **Data Table**.

Task 2: Enter a name for the collection table.

The Enter Collection Name and Alias page opens (Figure 15).

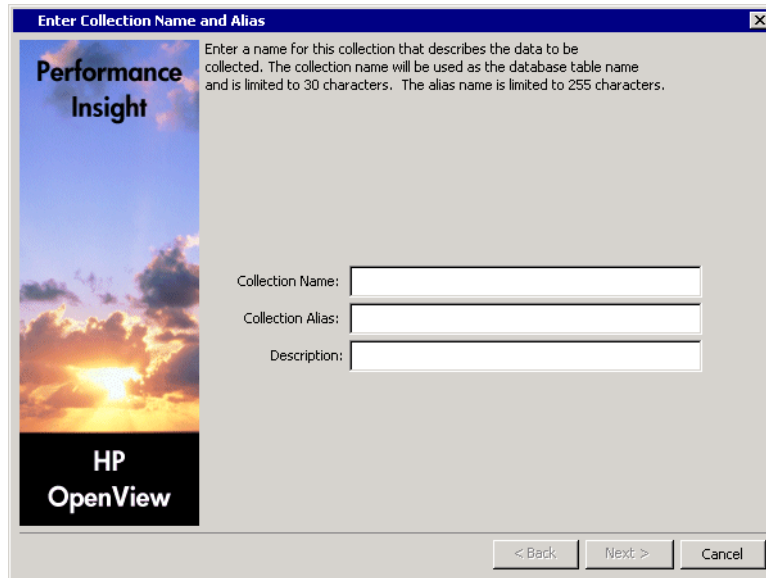


Figure 15 Enter Collection Name and Alias Page

- 1 In the Collection Name box, type the SQL name of the table.
The name must be unique and can be up to 30 characters. Valid entries include uppercase alphabetic characters, lowercase alphabetic characters, or both, all numeric digits, and the special character underscore (_).
- 2 In the Collection Alias box, type the alias name of the table.
The name must be unique and can be up to 255 characters. Valid entries include uppercase, lowercase alphabetic characters, or both, all numeric digits, and the special characters hyphen (-) and underscore (_).

- 3 In the Description box, type a description for the table. A description is optional. If you specify a description, it cannot exceed 255 characters.
- 4 Click **Next**.

Task 3: Select the table category, data pipe, time type, table type, and table retention attributes of the table.

The Define the Category page opens (Figure 16).

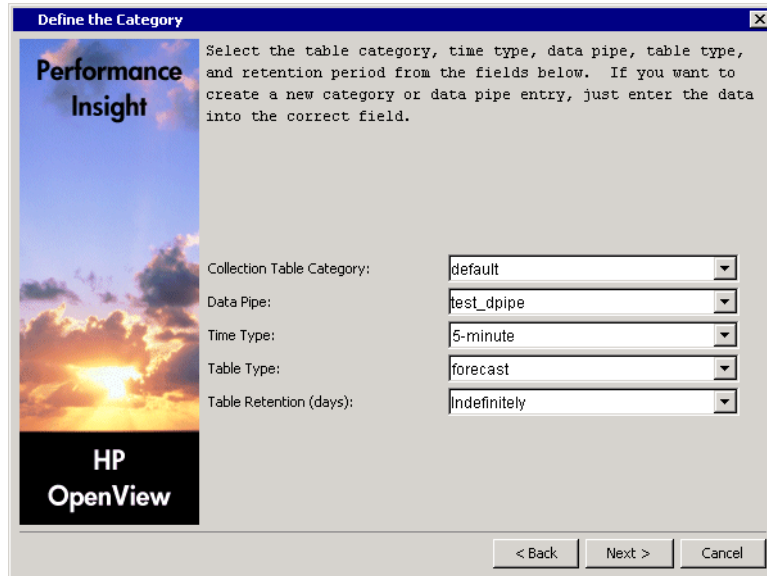


Figure 16 Define the Category Page

- 1 In the Collection Table Category box, type or select a category name.
The category name groups sets of tables together. If you do not know which name to select, select **default**. This name cannot exceed 30 characters.
- 2 In the Data Pipe box, type or select a datapipe name.
The datapipe name identifies the collectible property of the table. This name cannot exceed 30 characters.

If you want to use the `dpipe_snmp` datapipe, you should create the table using the MIB Browser. See [Creating a Collection Table on page 400](#) for details.

- 3 In the Time Type box, select a time characteristic for the data in the table.
- 4 In the Table Type box, select a table type.

The table type specifies the class of data contained in the table. Following is the list of available table type options:

baseline	Specifies the table contains baseline data.
event	Specifies the table contains event data, which is non-time-series oriented data with independent events.
forecast	Specifies the table contains forecast data.
rate	Specifies the table contains rate data, which is delta data computed from a pair of rows from the raw input data.
raw	Specifies the table contains raw (as-pollled) data.
sum	Specifies the table contains sum data.

- 5 In the Table Retention box, type or select the number of days you want to keep the data in the table.
If you never want to age the data out from the table, select **Indefinitely**.
- 6 Click **Next**.

Task 4: Define the statistics for the collection.

The Manually Define Statistics for the Collection page opens (Figure 17).

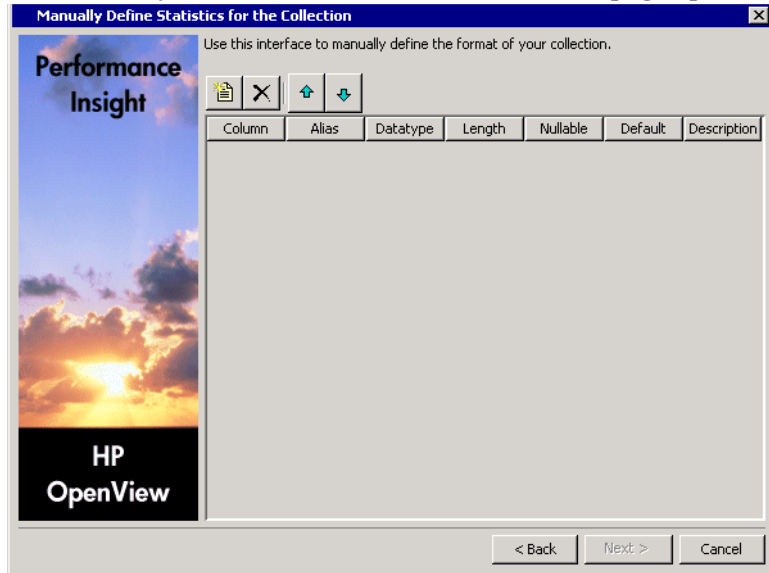






Figure 17 Manually Define Statistics for the Collection Page

- 1 Click  to create a column to add to the table.
- 2 In the Column box, type a name for the statistic.
The statistic name cannot exceed 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.
- 3 In the Alias box, type an alias name for the statistic.
The name cannot exceed 255 characters and contain the following characters: alphabetic, numeric, an underscore, or a hyphen.
- 4 In the Datatype box, select a value.
Following is the list of available datatype options and their corresponding database storage type:

OVPI Data Type	Sybase Data Type
char_string	varchar
counter	float

counter_high	float
counter_low	float
counter64	float
double	float
enum	integer
float	float
gauge	float
hex_string	varchar
integer	float
ip_address	varchar(15)
numeric	numeric
object_id	varchar(128)
octet_string	varchar(128)
seq_key	integer
smalldatetime	smalldatetime
snm_float	float
snmp_char_string	varchar
snmp_counter	float
snmp_double	float
snmp_enum	integer
snmp_float	float
snmp_gauge	float
snmp_hex_string	varchar
snmp_integer	float
snmp_ip_address	varchar(15)
snmp_object_id	varchar(128)
snmp_octet_string	varchar(128)
snmp_snm_float	float

snmp_timeticks	float
timeticks	float
tinyint	tinyint
unix_time	datetime

- 5 In the Length box, type a numeric value if needed.
Valid data types for length values are as follows: char_string, hex_string, numeric, octet_string, snmp_char_string, snmp_hex_string, or snmp_octet_string.
- 6 Click **Nullable** if the statistic can be NULL.
- 7 In the Default box, type the default value for the statistic.
The length of the value cannot exceed 255 characters.
- 8 In the Description box, type a description of the statistic.
A description is optional. The length of the description cannot exceed 255 characters.
- 9 If you want to change the order of the columns in the data table, select the column that you want to change and do the following:
 - To move the column up, click .
 - To move the column down, click .
 - To delete a column from the property table, click .
- 10 Click **Next**.

Task 5: Specify the property table associated with the data table.

The Associate with a Property Table page opens (Figure 18).

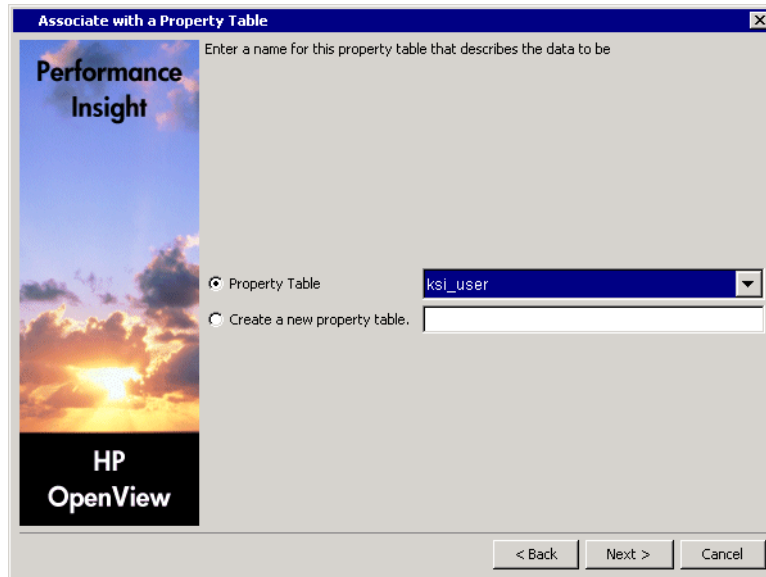


Figure 18 Associate with a Property Table Page

- 1 Do one of the following:
 - Select **Create a new property table**, and type the SQL name for the new property table. The name must be unique and cannot be greater than 30 characters. Valid characters are as follows: alphabetic characters, numeric digits, and underscores (_). Go to [Task 6](#).
 - Select a property table from the Property Table list. Go to [Task 7](#).
- 2 Click **Next**.

Task 6: Define the property attributes.

The Define Property Attributes page opens (Figure 19).

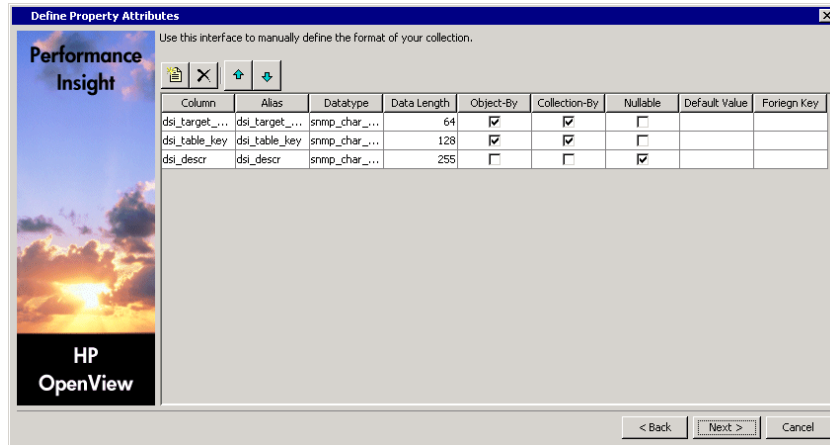



Figure 19 Define Property Attributes Page

You can create a property table with the existing columns or you can add more columns. In either case, you can modify the attributes of the columns.

To add or modify columns in the property table:

- 1 To add a column, click .

A column line appears at the end of the existing list. A default name appears in the Column and Alias boxes, which is *Column_n*, where *n* is a number that increments every time you add a column.
- 2 To specify the SQL name of the column, select the name in the Column box and type the new name.

The name cannot be greater than 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.

- 3 To specify the alias name of the column, select the name in the Alias box and type the new name.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Click the Datatype box, and select a value.

For the available datatype options and their corresponding database storage types, see [Step 4 on page 264](#).

- 5 In the Data Length box, type a numeric value if needed.

Valid data types for length values are as follows: `char_string`, `hex_string`, `numeric`, `octet_string`, `snmp_char_string`, `snmp_hex_string`, or `snmp_octet_string`.

- 6 Select or click to clear the **Object-By** box.

When you select this option, it indicates that the column is part of a unique combination of properties that define elements in the property table; in this case, the Nullable and Default Value boxes should not be selected.

If you specify any other column other than `dsi_target_name` as an object by-variable, you may not specify `dsi_table_key` as an object by-variable. However, you may specify `dsi_table_key` as an object by-variable only if you specify `dsi_target_name` as the first object by-variable.

- 7 Select or click to clear the **Collection-By** box.

When you select this option, it indicates that the column is part of a unique combination of columns that define elements to collect; in this case, the Nullable and Default Value boxes should not be selected.

If you specify any other column other than `dsi_target_name` as a collection by-variable, then you may not specify `dsi_table_key` as a collection by-variable. However, you may specify `dsi_table_key` as a collection by-variable only if you specify `dsi_target_name` as the first collection by-variable.

- 8 Select or click to clear the **Nullable** box.

When you select this option, it indicates that the column can contain NULL values. If the Object-By or Collection-By boxes are selected, this option should not be selected.

- 9 If the column should have a default value, type a value in the Default Value box.

The length of the value can be up to 255 characters. This field should be blank if this column is an object by-variable or a collection by-variable.


- 10 Do one of the following:


- a If this column contains a foreign-key reference, select the property table that has the key from the Foreign Key list.


- b If this column should not contain a foreign-key reference, select **no fk** from the Foreign Key list; the default is no foreign key.


- 11 If you designate a column as a by-variable, you must put the columns in order of the by-variables such that the first column in the list is the first by-variable and the second column is the second by-variable, and so on.

If you want to change the order of the columns in the property table, select the column that you want to change and do the following:

— To move the column up, click .

— To move the column down, click .

— To delete a column from the property table, click .

 The property table has only six default columns if the object by-variables and collection by-variables are the same.

Task 7: Specify an output file for the collection.

The Select the Output File page opens (Figure 20).

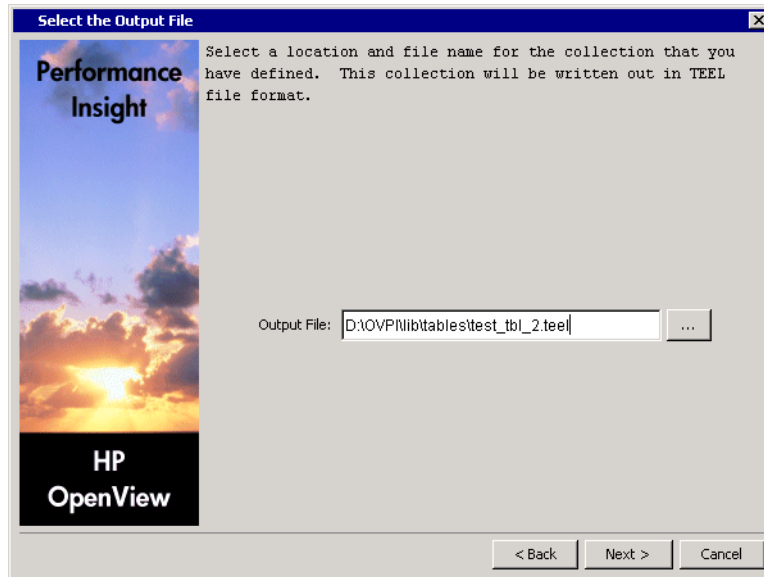



Figure 20 Select the Output File

- 1 Verify that the name and location of the output file are correct.
- 2 If you need to change the name or location of the file, do one of the following:
 - Type the new information in the Output File box.
 - Click . The Open dialog box opens.
Locate the file, double-click it, and then click **Open**. The Output File box is updated with the new file information.
- 3 Click **Next**.

Task 8: Register the collection with the database.

The Add Collection to Database page opens (Figure 21).

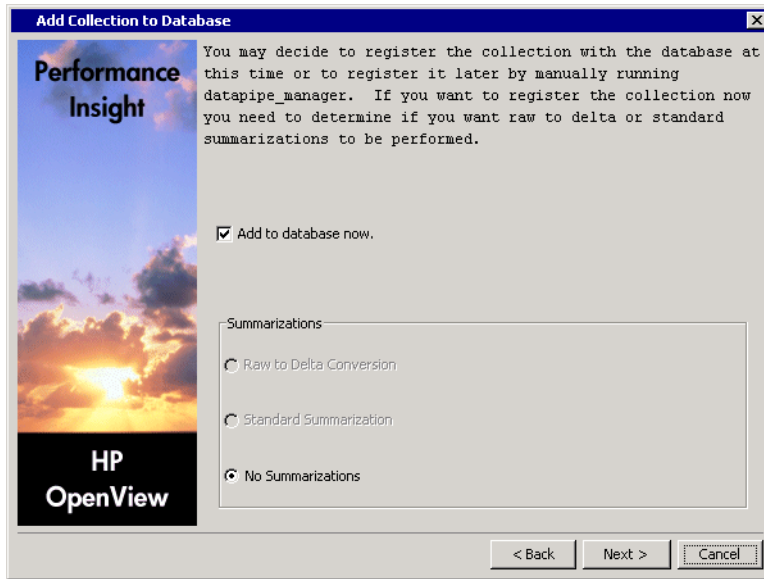


Figure 21 Add Collection to Database Page

- 1 If you want to generate the table now and add it to the database, select **Add to database now.**

If you click to clear this option, you will need to enter the following command at the command line:

```
datapipe_manager -p create -a teel_file
```

In this instance, *teel_file* is the name of the output file from [Task 7 on page 271](#).

If you need additional information about the options for the `datapipe_manager` command, refer to the *Performance Insight TEEL Reference Guide* and the *Performance Insight Reference Guide*.

- 2 In the Summarizations box, select the type of summarization that the system will perform on the data in the table.

Click one of the following options:

— **Raw to Delta Conversion**

The system will generate the required database objects for raw-to-delta processing.

For manual processing, add `-r r2d` to the command in [Step 1](#) of this task.

— **Standard Summarization**

The system will generate the required database objects for raw-to-delta processing and the default summaries. The database objects include the rate table, the hourly, daily, weekly, and monthly summary tables, and the supporting summary procedures. It also generates the **.sum** files and stores them in the `DPIPE_HOME/scripts` directory.

For manual processing, add `-r default` to the command in [Step 1](#) of this task.

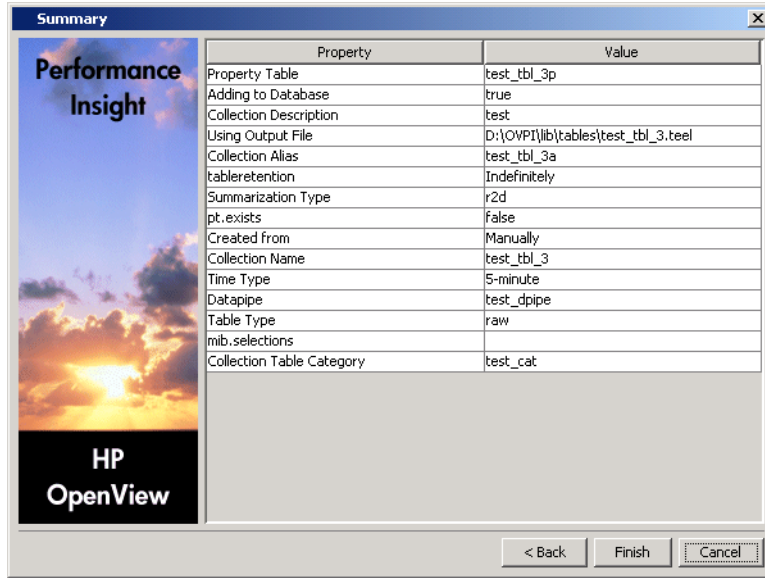
— **No Summarizations**

The system will not generate any additional database objects.

- 3 Click **Next**.

Task 9: Review the Summary page.

The Summary dialog box opens.



- 1 Verify that all the collection property values are correct.

If necessary, click **Back** to return to a previous page to make a change. After you make a change, click **Next** until you return to the Summary page.

- 2 Click **Finish** to create the data table.


The data table creation wizard closes, and OVPI creates the TEEL file and stores it in the designated location.

Creating Property Tables

To create a property table, complete the following tasks:

Task 1: Start the Table Creation wizard.

From the Management Console, do one of the following to start the Table Creation wizard.

- Click , and select **Property Table** from the menu.
- From the File menu, select **New**, and then select **Property Table**.

Task 2: Enter a name for the property table.

The Create Property Table page opens.

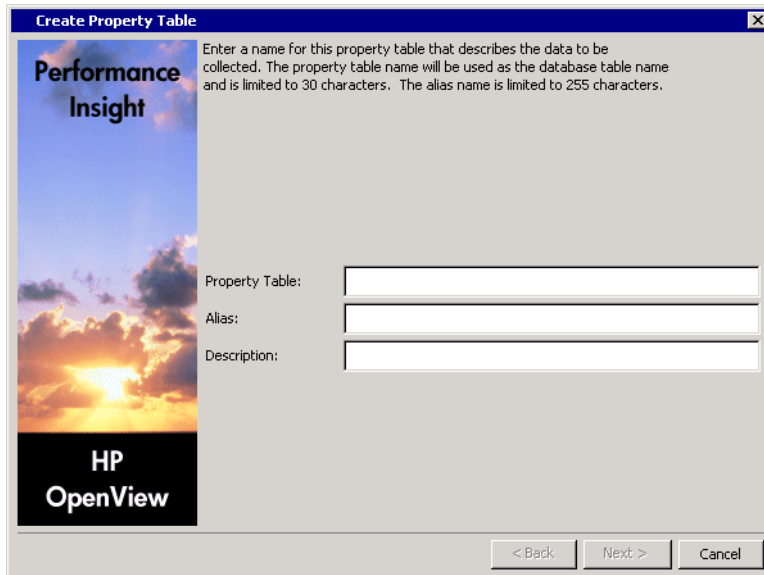


Figure 22 Create Property Table Page

- 1 In the Property Table box, type the SQL name of the table.

The name must be unique and can be up to 30 characters. Valid entries include uppercase alphabetic characters, lowercase alphabetic characters, or both, all numeric digits, and the special character underscore (_).

- 2 In the Alias box, type the alias name of the table.

The name must be unique and can be up to 255 characters. Valid entries include uppercase alphabetic characters, lowercase alphabetic characters, all numeric digits, and the special characters hyphen (-) and underscore (_).

- 3 In the Description box, type a description for the table.

A description is required; it cannot exceed 255 characters. If you enter only spaces, an error message will display after you create the table.

- 4 Click **Next**.

Task 3: Define the property attributes.

The Define Property Attributes page opens. See [Figure 19 on page 268](#).

To complete this task, see [Task 6 on page 268](#).

Task 4: Specify an output file for the collection.

The Select the Output File page opens. See [Figure 20 on page 271](#).

To complete this task, see [Task 7 on page 271](#).

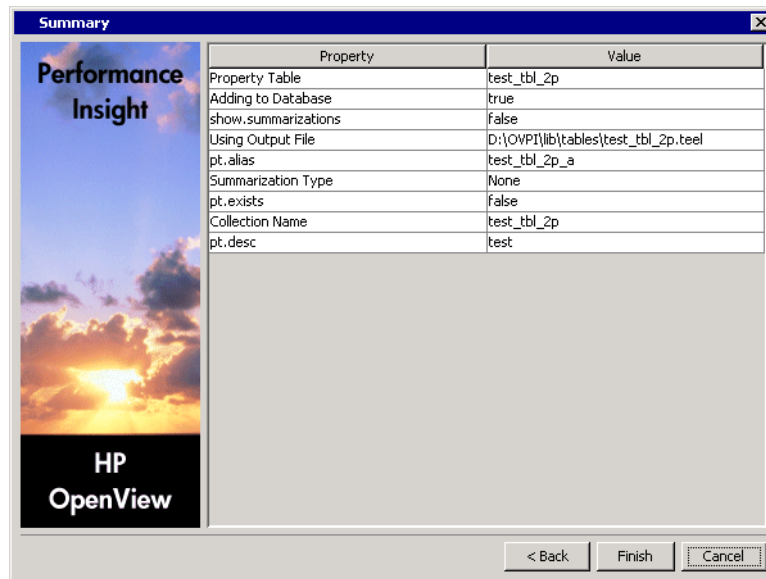
Task 5: Register the collection with the database.

The Add Collection to Database page opens ([Figure 21 on page 272](#)).

To complete this task, see [Task 7 on page 271](#).

Task 6: Review the Summary page.

The Summary dialog box opens.



- 1 Verify that all the values are correct.

If necessary, click **Back** to return to a previous page to make a change. After you make a change, click **Next** until you return to the Summary page.

- 2 Click **Finish** to create the property table.

The data table creation wizard closes, and OVPI creates the TEEL file and stores it in the designated location.

Setting Table Defaults

Table Manager lets you set defaults for two types of tables: data tables and time-period tables. You can set the defaults for the data tables by table type, which apply to all data tables with that type. The default for the time-period tables applies only to the selected time-period table, which may affect many applications.

Setting Data Table Defaults

There are various default options that you can set for a particular time type associated with a table. These options are retention time, time lag value, and hysteresis units.




If you change the default values, it may affect the performance and the data in the installed OVPI reporting solutions.

To change defaults:

- 1 From the File menu, select **Set Table Defaults**.

The Set Table Defaults dialog box opens.



- 2 From the Time Type list, select the time type of the table you want to modify.



When you change the default values for a particular table type, the change affects all tables with that time type.

- 3 If you want to change the retention time, type a number of days in the Days to Keep Data box.

The *retention time* is the number of days that OVPI will keep the data for a particular table. It allows OVPI to age out data after the specified number of days.

- 4 If you want to change the lag time, type a number of minutes in the Time Lag box.

The *time lag* value is the number of minutes that OVPI will delay processing the data.

- 5 If you want to change the hysteresis units, type the number of units in the Hysteresis Units box.

The *hysteresis units* specify the amount of data to reprocess. The value of the units is equal to the time type of the table. For example, the hysteresis units for a daily table are days; similarly, the hysteresis units for an hourly table are hours.

- ▶ The dialog box also displays the setting for execution frequency; however, you cannot modify its value. The *execution frequency* specifies how often OVPI normally processes the data in the table.

- 6 Click **OK**.

Time-Period Table Default

There is only one default option that you can set for time-period tables; it is retention time. Time-period tables contain lists of time periods for a specified number of days, which is the retention time for the table. There are time-period tables for various time categories, such as day, hour, 5-minute, and 30-second. These tables let OVPI optimize its searches of the database. Refer to the *Performance Insight Reference Guide* for more information about tpmaint.




If you change the retention time, it may affect the performance and the data in the installed OVPI reporting solutions.

To change the retention time for a time-period table:

- 1 From the File menu, select **Set Time Period Table Retention**.

The Set Time Period Table Retention dialog box opens.



- 2 From the Time Period Tables list, select the appropriate time-period table.
This list contains a set of time-period tables for each time zone available in the database. Every database should have access to a local time zone set and a GMT time zone set.
 - ▶ If you change the retention days for the time-period tables, you may affect all the applications that use these tables for optimization.
- 3 Type a number of days in the Retention Days box. This will modify the number of time-periods that the table contains.
 - ▶ If you set this value incorrectly, it can degrade performance.
- 4 Click **OK**.

Setting the Display Options

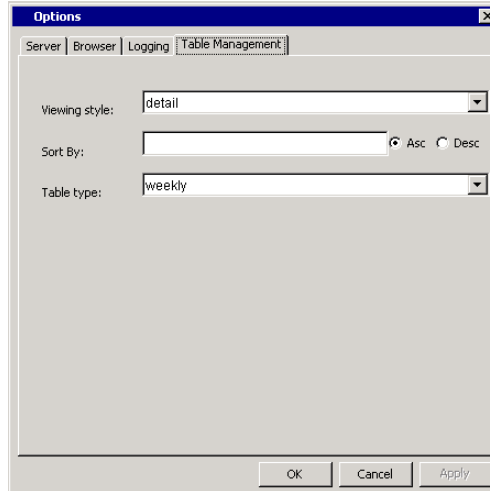
You can change how tables will display in Table Manager.

- ▶ When you change the default display settings, the changes that you make will take effect the next time you log on to the Management Console.

To change how tables display:

- 1 From the Edit menu, select **Options**.
The Options dialog box opens.

2 Select the Table Management tab.



3 Do any of the following:

- To specify how the tables will appear in the viewing pane, select an option from the Viewing Style list.

You can choose to display tables as large icons, as small icons, in a simple list, or in a detailed list.

- To sort the rows in a table based on the values in a particular table column, do the following:
 - Type the column name in the Sort By box. You must match the exact spelling and capitalization of the actual table column.
 - Select **Asc** (ascending) or **Desc** (descending) to determine the sort order. The rows will be in numerical or alphabetical order, depending on the type of data the column contains.

If a table does not contain the column name specified in the Sort By box, Table Manager ignores these settings.

- To determine the default time type for new tables, select the desired time period from the Table Type list. You can also set default values for tables of each time type; see [Setting Table Defaults on page 277](#).

Using the Web Access Server

This chapter describes the administrative features of the Web Access Server; these features are available by using the Administration and Catalog links. To access these links, you must be logged on to the Web Access Server with an account that has Administrator privileges.

Specifically, this chapter describes how to do the following:

- Maintain the user accounts required for the following client applications: Web Access Server, Report Viewer, Report Builder, and the Management Console.
- Configure the Web Access Server, OVPI system attributes, mail server, and Secure Socket Layer (SSL) communications.
- View client and Web Access Server log files and control their appearance.
- Control Java settings for the Web Access Server, Report Builder, and Report Viewer.
- Set reporting attributes.
- Create report views.
- Control the appearance of the Web Access Server.
- Deploy reports to the `System` and `Users` folders.
- Change settings for the services.

Accessing the Web Access Server

Before you attempt to access the Web Access Server, ensure that you have met the following prerequisites:

- The Web Access Server is running. If not, see [Chapter 16](#), .
- You have one of the following Web browsers:
 - Netscape Navigator 6.x or later
 - Internet Explorer 5.x or later

To access the Web Access Server:

- 1 Start your Web browser.
- 2 Enter the following URL in the Address or Location field of your browser, and then press **ENTER**.

http://server_name.organization.type:port_number

server_name is the name of the Web Access Server.

organization is the name of your organization.

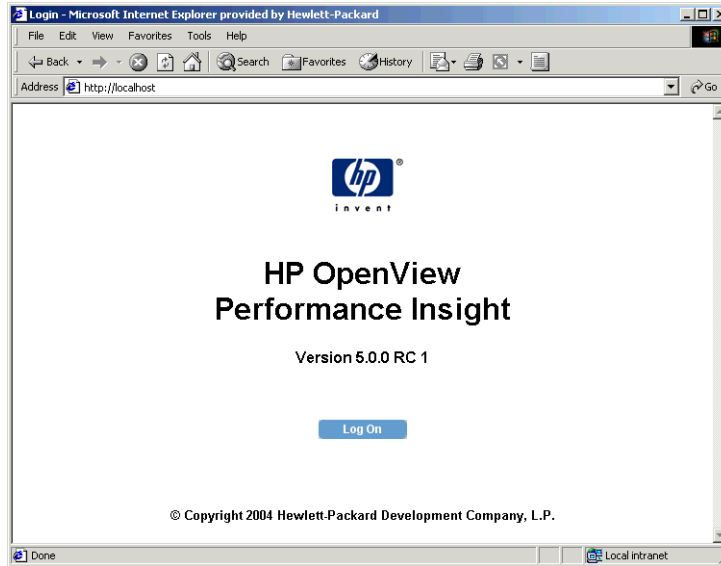
type is the domain type (for example, com, org, or net).

port_number is the port number of the default Web Access Server.



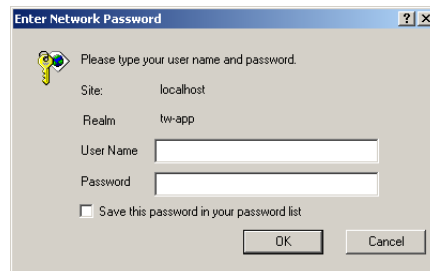
You do not need to enter a port number if your Web Access Server is using port 80.

The Log On page opens.



- 3 Click **Log on**.

The Enter Network Password dialog box opens.



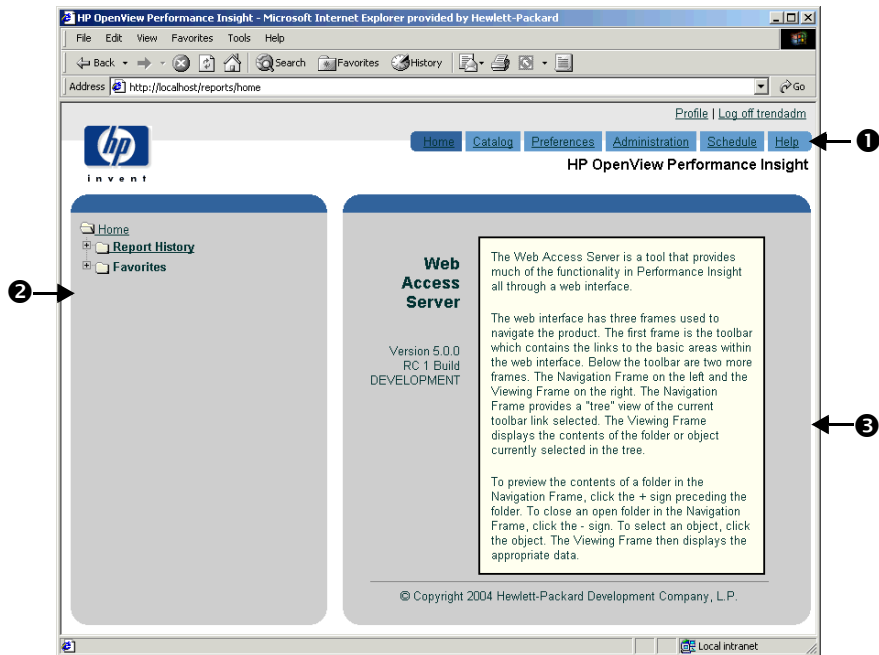
- 4 Enter your Administrator username and password, and then click **OK**.

OVPI creates a default Administrator account called `trendadm` when it is installed. Use this account or use another account that you may have created that has Administrator privileges.

The Web Access Server opens.

Web Access Server Overview

The following figure shows the Web Access Server.



The Web Access Server comprises the following components:

- Links bar (1)

The links bar contains the following links: Home, Catalog, Preferences, Administration, Schedule, and Help. The links that allow you to perform administrative features are the following:

- The Catalog link lets you deploy reports to the Web server in both the System and Users folders and view the reports that have been deployed to Web server.
- The Administration link lets you maintain user and user group accounts, configure services, view log files, modify Java and report settings, and control the appearance of the Web Access Server.

For information about the links that are available to non-Administrator users, refer to the *Performance Insight Guide to Building and Viewing Reports*.

- Navigation frame (2)

Lets you access the folders associated with each link using a standard tree structure.

▶ If the navigation frame does not display, refer to the Troubleshooting appendix in the *Performance Insight Installation Guide*.

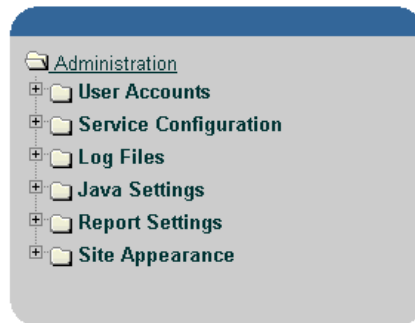
- Results frame (3)

Displays the contents of the folder or the report currently selected in the navigation frame.

Administration Link

Use the Administration link to maintain client application accounts and modify services, log file, Java, and report settings.

The navigation tree for the Administration link appears as follows:

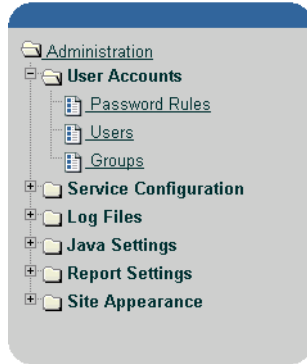


It contains the following folders: User Accounts, Service Configuration, Log Files, Java Settings, Report Settings, and Site Appearance.

User Accounts

Use this folder to create, modify, and delete user accounts and set password rules. These accounts are used to log on to the following client applications: Web Access Server, the Management Console, Report Viewer, and Report Builder.

The navigation tree for the User Accounts folder appears as follows:



The user accounts that you create using the User Accounts folder are the accounts that let users access OVPI's client applications (Management Console, Web Access Server, Report Builder, and Report Viewer).

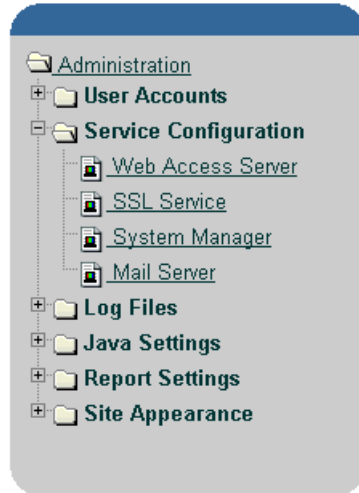
The User Accounts folder contains the following:

- Password Rules page
Use this page to set password validation rules for the passwords that you create for user accounts. See [Setting Password Validation Rules on page 295](#) for details.
- Users folder
Use this folder to create, modify, or delete user accounts including the default user account, `trendadm`. See [Maintaining User Accounts on page 295](#) for details.
- Groups folder
Use this folder to create, modify, or delete groups of user accounts. See [Maintaining User Groups on page 303](#) for details.

Service Configuration

Use this folder to modify settings for the Web Access Server and the system on which you installed OVPI, set up Secure Socket Layer communications, and specify the attributes for mailing a report link.

The navigation tree for the Service Configuration folder appears as follows:



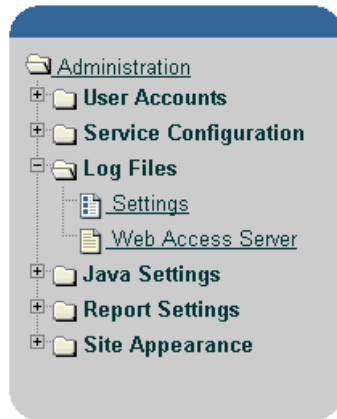
The Service Configuration folder contains the following:

- **Web Access Server page**
Use this page to change the current settings for the Web Access Server. The *Web Access Server* controls the OVPI client's access to the OVPI database. See [Changing the Web Access Server Settings on page 329](#) for details.
- **SSL Service page**
Use this page to enable Secure Socket Layer (SSL) communication between the Web Access Server and any supported Web browser. See [Changing the SSL Service Settings on page 333](#) for details.
- **System Manager page**
Use this page to change the current settings for the System Manager. The *System Manager* maintains the list of all Web Access Servers and databases. See [Changing System Manager Settings on page 337](#) for details.
- **Mail Server page**
Use this page to specify the attributes for mailing a report link. See [Changing Mail Server Settings on page 345](#) for details.

Log Files

Use this folder to access the log file for the Web Access Server.

The navigation tree for the Log Files folder appears as follows:



The Log Files folder contains the following:

- Settings page

Lets you set the number of log file entries that appear on a page and the background color of the log file's frame. See [Changing Log File Settings on page 311](#) for details.

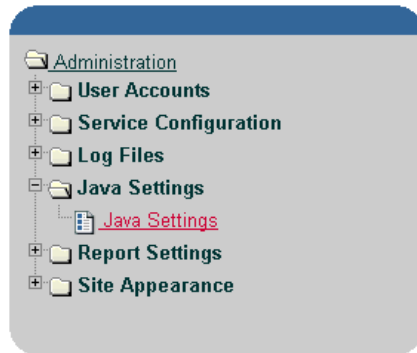
- Web Access Server page

Lets you view the log file for the Web Access Server. This log file contains all the information, warning, error, and debug messages for the server. See [Viewing Log Files on page 312](#) for details.

Java Settings

Use this folder to change the Java settings for Report Builder and Report Viewer and the Web Access Server services.

The navigation tree for the Java Settings folder appears as follows:

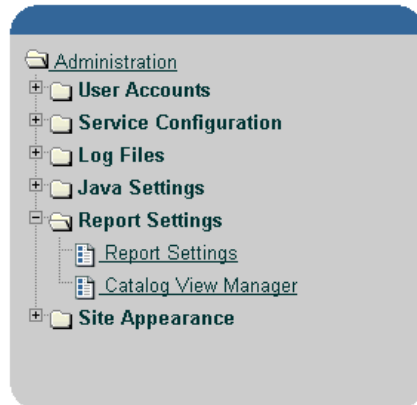


It contains the following: Java Settings page. This page lets you change heap size and arguments. See [Changing Java Settings on page 314](#) for details.

Report Settings

Use this folder to change report settings and set up catalog views.

The navigation tree for the Report Settings folder appears as follows:

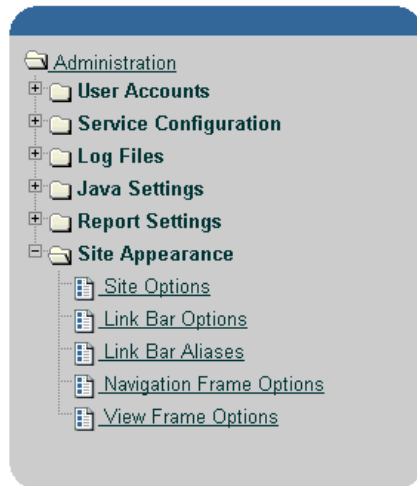


It contains the following:

- Report Settings page
This page lets you set attributes for reports. See [Changing Report Settings on page 315](#) for details.
- Catalog View Manager page

This page lets you create views. A *view* allows users to see a certain list of deployed reports in the `System` folder. See [Creating Views on page 318](#) for details.

The navigation tree for the Site Appearance folder appears as follows:



The Site Appearance folder contains the following:

- **Site Options page**
Lets you change the title that appears in the title bar of your Web browser. See [Changing Site Options on page 321](#) for details.
- **Link Bar Options page**
Lets you change numerous link bar options, some of which include the link bar's title, alignment, height, color, image size and location, and focus color. See [Changing Link Bar Options on page 322](#) for details.
- **Link Bar Aliases page**
Lets you show or hide and change the name of any of the links. See [Changing Link Bar Aliases on page 324](#) for details.
- **Navigation Frame Options page**
Lets you change the width, background color, and background image for the navigation frame. See [Changing Navigation Frame Options on page 325](#) for details.
- **Results Frame Options page**

Lets you change the margin color and the background color of the results frame. See [Changing Results Frame Options on page 326](#) for details.

Catalog Link

Use this link to deploy reports to the Web server and view the reports that have been deployed in the `System` and `Users` folders.

The navigation tree for the Catalog link appears as follows:



The Catalog link contains the following page and folders:

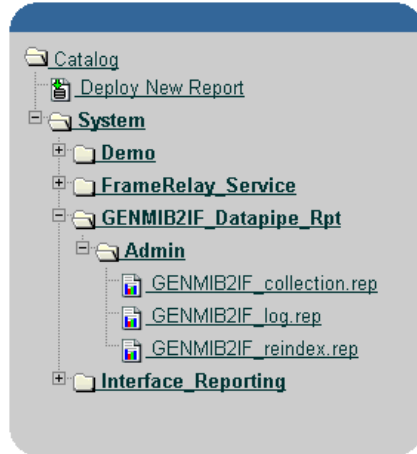
- Deploy New Report page

Lets you deploy reports to the Web Access Server in order to view them using the Web Access Server; you must deploy reports before you can view them. You can also view these reports using the Browse option from Report Builder's or Report Viewer's File menu. See [Deploying Reports on page 327](#) for details.

- System folder

All users can access the reports deployed to the `System` folder; however, if the **Show System Directory to Users** option is not selected or the reports in this folder are not in a user's catalog view, the reports cannot be viewed. Only the OVPI Administrator can deploy reports to this folder.

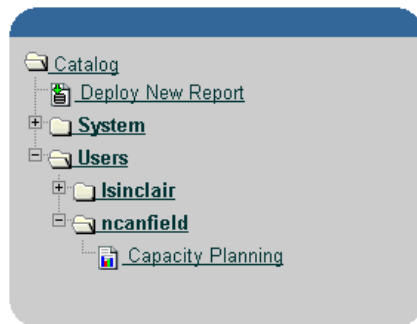
Following is an example of how the navigation pane might appear when reports are deployed to the `System` folder:



- Users folder

This folder appears only after reports have been deployed to the Web Access Server by a user or by the OVPI Administrator. Once the user or Administrator deploys a report, the report appears in that user's folder. Only users (and the Administrator) can access the reports deployed to their folders.

Following is an example of how the navigation pane might appear when a Capacity Planning report is deployed to a user folder that belongs to Nathan Canfield:



- Generated Results folder

This folder displays the results of generated reports. Its contents is identical to that of the Results folder, which is displayed when you click the Schedule link.

Maintaining User Accounts

All OVPI clients (Report Builder, Report Viewer, Web Access Server, and the Management Console) require users to log on.

By default, OVPI creates one user account when you install it. This is the `trendadm` account with a password of `trendadm`; ideally, you should create a user account for each person that will use the OVPI client applications.

This section explains how to do the following:

- Set password validation rules for user accounts.
- Create, modify, and delete user accounts.
- Assign a user the role of user or Administrator.

Alternately, you can perform most of the same tasks described in this section using the **userimport** and **groupimport** commands. For more information about these utilities, refer to the *Performance Insight Reference Guide*.



The procedures contained in this section assume that you are logged on to the Web Access Server using the default OVPI Administrator user account or a user account that has OVPI Administrator privileges.

Setting Password Validation Rules

You can set validation rules for user account passwords.

To set password validation rules:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts and then click **Password Rules**.

The Password Rules page opens.

Password Rules

Enable or disable the following rules to modify the behavior that account passwords must follow.

<input type="checkbox"/>	Password must have minimum length.	4
<input type="checkbox"/>	Password cannot be circular shift of account name.	
<input type="checkbox"/>	Password must have minimum number of numeric characters.	1
<input type="checkbox"/>	Password must have minimum number of punctuation characters.	1
<input type="checkbox"/>	Password cannot be the same as the last <i>n</i> values.	3
<input type="checkbox"/>	User can continue to login for <i>n</i> days after password expires.	10

Apply

3 Do any of the following:

- a** To set a minimum password length, select the **Minimum length** check box, and then type a value. (The maximum value that you can enter is 9999.)

To disable this rule, click to clear the **Minimum length** check box.

- b** To prevent the password from being a circular shift of the letters used in the account name, select the **Circular shift of account name** check box.

For example, if the account name is `caesar`, the password for the account cannot be any of the following: `caesar`, `aesarc`, `esarca`, `sarcae`, `arcae`, or `rcaesa`.

To disable this rule, click to clear the **Circular shift of account name** check box.

- c** To require that the password have a minimum number of numeric characters, select the **Minimum number of numeric characters** check box, and then type a value. (The maximum value that you can enter is 9999.)

To disable this rule, click to clear the **Minimum number of numeric characters** check box.

- d** To require that the password have a minimum number of punctuation characters, select the **Minimum number of punctuation characters** check box, and then type a value. (The maximum value that you can enter is 9999.)

To disable this rule, click to clear the **Minimum number of punctuation characters** check box.

- e To specify that a password cannot be reused until a specified number of passwords have been used after it, select the **Allow reuse of a password after n passwords** check box, and then type a value. (The maximum value that you can enter is 9999.)

For example, suppose the current password is `caesar` and the value for this option is set to 3. When the password expires, the user must wait until she has used three other passwords before she can use `caesar` again for her password.

To disable this rule, click to clear the **Allow reuse of a password after n passwords** check box.

- f To specify the number of days a user can use a password after it has expired, select the **Expired password valid for n days** check box, and then type a number of days. (The maximum value that you can enter is 9999.)

To disable this rule, click to clear the **Expired password valid for n days** check box.

- 4 Click **Apply** to make any changes.



Changes you make will not apply to current passwords; password rules are enforced whenever the user is asked to make a password change.

Creating a User Account

To create a user account:

- 1 From the Links bar, select **Administration**.

- In the Administration tree, expand User Accounts, and then click **Users**. The User Administration page opens, which lists the current user accounts.



- Click **Add New User**. The Add a New User page opens.



- 4 Enter the requested information into all boxes, noting the constraints that follow.

Ensure that you do **not** use any of the following characters in any of the fields on this page: ampersand (&), backslash (\), comma (,), double quotation mark (“), forward slash (/), question mark (?), percent (%), and single quotation mark (‘).

The valid entries are as follows:

- All alphabetic characters (uppercase and lowercase), all numeric symbols, and all special characters except those in the previous list.
 - The Username must be a minimum of one (1) character, can be a maximum of 255 characters, and can only contain alphanumeric characters, spaces, underscores (_), and hyphens (-).
 - The user’s real name and company/department have no minimum length (that is, they can be blank) and can be a maximum of 255 characters.
 - The password can be a minimum of four (4) characters, a maximum of 255 characters, and cannot contain a space (blank) character.
 - The telephone number can be in any format but cannot exceed 255 characters.
 - The email address can be in any format but cannot exceed 255 characters.
- 5 Select the **Is Administrator** check box only if the user will have an Administrator role. If you do not select this option, the account is automatically assigned a user role.
 - 6 To choose a user group status for the user, do one of the following:
 - To assign the user to a user group, select the user group name in the Group box.
 - To assign the user to multiple user groups, hold down the **Ctrl** key, and then select the user group names you want.
 - To leave the user out of any user groups, select **No Group**.
 - 7 In the Database list, select the database that the user will access. For information about how to add another database to the current list, see [Changing System Manager Settings on page 337](#).

- 8 In the View list, select the view that the user will use. For information about how to create a view, see [Creating a View on page 318](#).
- 9 To choose the password expiration setting for the user account, do *one* of the following:
 - Click the **Never Expires** option, which will keep the password for the account indefinitely.
 - Click the **Password Expires** option, and then enter a number of days to set an expiration period for the password.

When a password expires, a page will display informing the user that the password for the account has expired and will give the user the option to change the password for the account.

In addition to this option, you can do the following:

- Set up password validation rules. For example, you can set a grace period for a password. This option lets the user to log on to an account a specific number of times after the password for the account has expired. When this period expires, the account is locked. See [Setting Password Validation Rules on page 295](#) for details.
 - Specify the number of unsuccessful log on attempts that a user can have before an account is locked. See [Changing the Web Access Server Settings on page 329](#) for details.
- 10 To lock or unlock a user account, do *one* of the following:
 - To lock a user account, select the **Do not allow user to login** check box.
 - To unlock a user account, click to clear the **Do not allow user to login** check box.

Account locking is enforced when a user logs on to an application; thus, if you lock an account while the user is logged on, the user will not know the account has been locked until the user logs on again.

When an account is locked and the user attempts to log on, a page will display notifying the user that the account has been locked and that the user must contact the Administrator, the only one who can unlock the account.

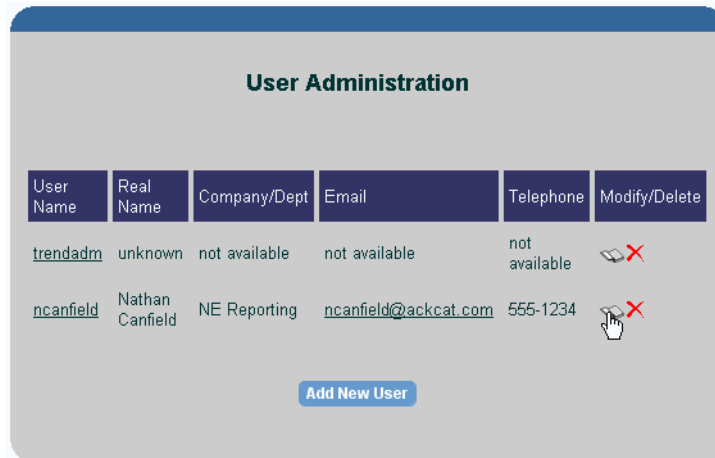
- 11 Click **Apply**.
- 12 Click **OK** to confirm the new user account.


After you confirm the new user account, the User Administration page lists the new user account information. You can also view the user account in the navigation tree.

Modifying a User Account

To modify a user account:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand **User Accounts**, and then click **Users**. The User Administration page opens, which lists the user accounts.



- 3 Do one of the following:
 - Click the user account on the User Administration page.
 - Click the **Edit** icon () next to the user account that you want to modify.

The Edit *Username* page appears.

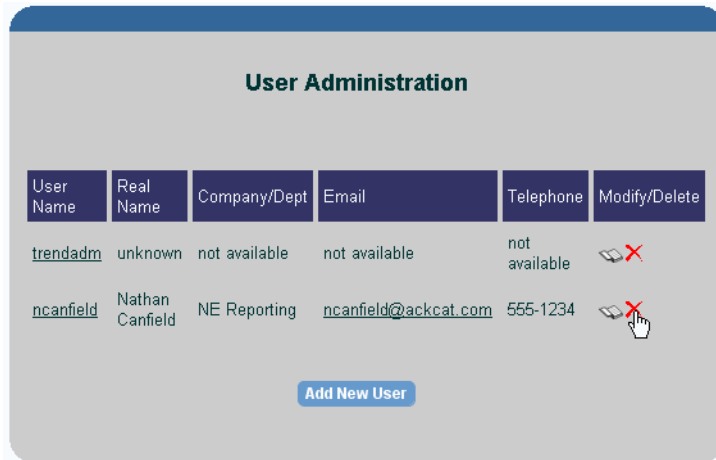
- 4 Make your changes, and then click **Apply**.
- 5 Click **OK** to confirm your changes.

After you confirm your changes, the User Administration page lists the updated user account's information.

Deleting a User Account

To delete a user account:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and then click **Users**. The User Administration page opens, which displays the list of user accounts.



- 3 Click (the Delete icon) next to the user that you want to delete. A confirmation dialog box appears.
- 4 Click **OK**.

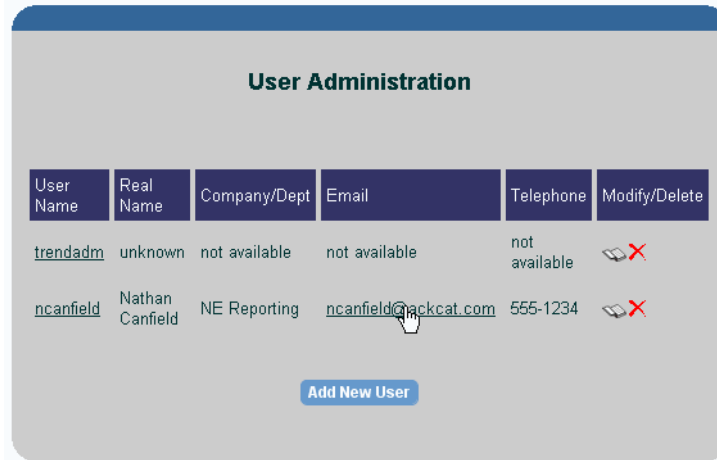
Sending Email to a User Account

You can send an email to a user from within your Web browser. You must have an email program installed on your system to do so.

To send email to a user account:

- 1 From the Links bar, select **Administration**.

- In the Administration tree, expand User Accounts, and then click **Users**. The User Administration page opens, which lists the user accounts.



- Click the email address for the user account. Your email application appears with a mail message addressed to the specified user.
- Complete the mail message, and then send it.

Maintaining User Groups

This section explains how to do the following using the Web Access Server:

- Create, modify, and delete user groups.
- Create, modify, and delete user group filters.



The procedures contained in this section assume that you are logged on to the Web Access Server as the OVPI Administrator.

You can perform most of the same tasks described in this section using the **userimport** and **groupimport** utilities. For more information, refer to the chapter about Account Administration Utilities in the *Performance Insight Reference Guide*.

Understanding How User Groups Work

A *user group* is:

- A collection of user accounts

For example, you can create a user group called *Thunderbolt* for all users in a company called Thunderbolt, Inc. This is useful if you are a service provider who wants to limit the data that certain companies can view.

- A collection of user accounts and user groups

For example, you can create a user group called *All*, which contains some user accounts and four user groups called *North*, *East*, *South*, and *West*.

The constraint placed on the users in the *North*, *East*, *South*, and *West* user groups lets those users view only the data for their specific region.

The users in the *All* user group can view all the data that the users in the *North*, *East*, *South*, and *West* user groups can view because the top-level user group inherits the constraints of the user groups below it in the hierarchy.

When you create user groups, it is recommended that you try to create effective constraints for the user groups to limit the number of user groups you might have. For example, instead of using `interface` as a constraint and having 1000 or more user groups for each interface, use `cust-id` as a constraint instead. A large number of user groups can affect the efficiency and size of the query sent to the database.



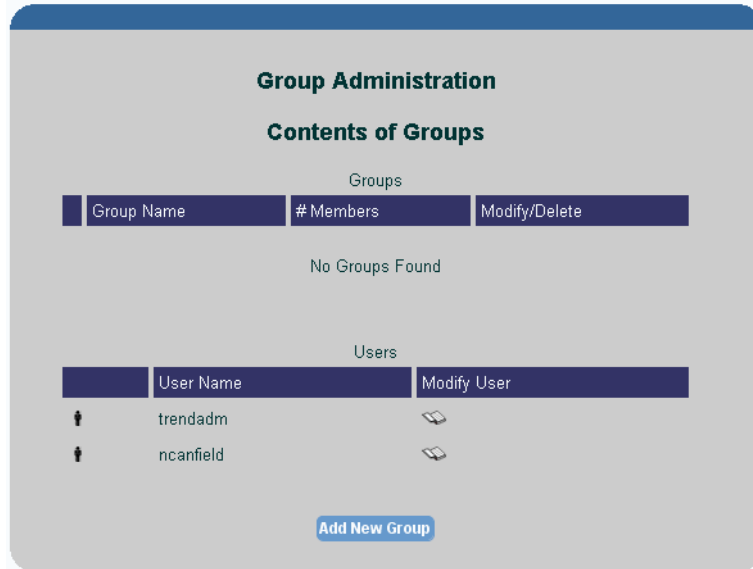
If there is a report that only one user should see, deploy that report to the user's report directory under the users folder. See [Deploying Reports on page 327](#) for details.

Creating a User Group

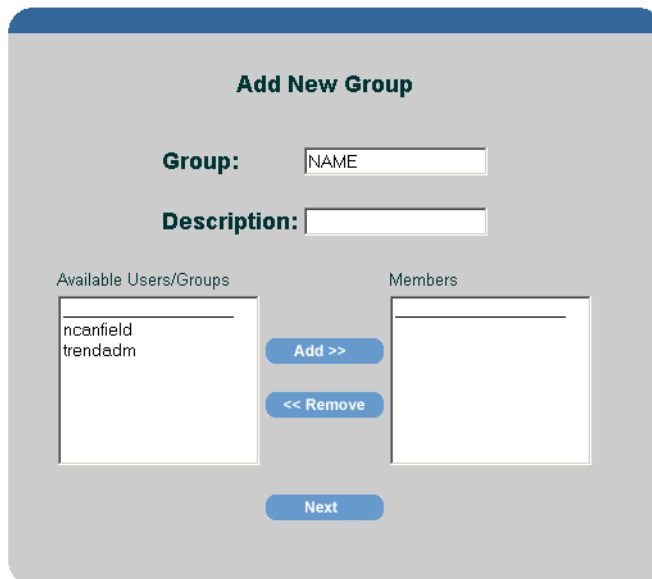
To create a group:


- 1 From the Links bar, select **Administration**.

- In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens.



- Click **Add New Group**. The Add New Group page opens.



- 4 To create a new user group, do the following:
 - a In the Group box, enter the name of the user group you want to create.
 Do **not** enter a user group name that already exists or use the name “Group”.

The user group name must adhere to the following constraints:

Ensure that you do **not** use any of the following characters in the user group name:

ampersand (&), asterisk (*), backslash (\), colon (:), comma (,), double quotation mark (“), exclamation point (!), forward slash (/), greater than (>), left bracket ([), less than (<), period (.), pound (#), question mark (?), right bracket (]), single quotation mark (‘), percent (%), and a blank space.

The valid entries for the user group name are as follows:

- The name can contain all alphabetic characters (upper case and lower case), all numeric symbols, and all special characters except those cited above.
 - Minimum of one (1) character and a maximum of 255 characters.
- b Select the user or users and the user group or user groups from the Available Users/Groups list that will belong to this new user group, and then click **Add**. The user names, user group names, or both are moved to the Members list.

- c Click **Next**. The Group Filters page opens.



A *filter* is a comparison that involves a table column name, a numerical operator, and a value for the statistic.

- 5 To create a filter:
- a In the Column field, enter the table column name to filter on.
 - b Click the Operator arrow, and then select the operator.
 - c In the Value field, enter the value.

For example, you are a service provider who has the ABC Company as a client. Create the ABC Sales user group with a filter that allows this user group only to see ABC's reports. To do this, you would create a filter similar to the following:

```
cust_id = 05
```

`cust_id` is the table column name.

`=` is the numerical operator.

`05` is ABC's customer identification number.


- d Click **Add Filter**.

- 6 Click **Finish**.

The Contents of Groups page lists the new user group; the user group also appears in the navigation frame under the Groups folder.

Modifying a User Group

To modify a user group:


- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Do one of the following:
 - Click the user group on the Contents of Groups page.
 - Click  (the Edit icon) next to the user group. If you choose this option, go to step 5 of this procedure.

The Contents of Group_Name page opens, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Group**. The Modify Group page appears, which is identical to the Add New Group page.
- 5 Select the user or user group from the Available Users/Groups list that you want to add to the user group, and then click **Add**. The user or user group is added to the Members list.
- 6 Select the user or the user group from the Members list that you want to remove from the user group, and then click **Remove**. The user or user group is removed from the Available Users/Groups list.
- 7 Click **Next**, and then click **Finish**.

Deleting a User Group

To delete a group:


- 1 From the Links bar, select Administration.
- 2 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Click  (the Delete icon) next to the user group that you want to delete. A confirmation dialog box appears.

- 4 Click **OK**.

Creating a User Group Filter

To create a user group filter:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Do one of the following to select the user group to which you want to add a filter:

- Click the user group on the Contents of Groups page.
- Click  (the Edit icon) next to the user group.


The Contents of Group_Name page opens, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Filter**. The Group Filters page opens.
- 5 Do the following:
 - a In the Column field, enter the table column name to filter on.
 - b Click the Operator arrow and select the operator.
 - c In the Value field, enter the value.
 - d Click **Add Filter**.
- 6 Click **Finish**.

Modifying a User Group Filter

To modify a user group filter:

- 1 From the Links bar, select Administration.
- 2 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.


- 3 Do one of the following:
 - Click the user group on the Contents of Groups page.
 - Click  (the Edit icon) next to the user group.

The Contents of Group_Name page appears, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Filter**. The Group Filters page opens.
- 5 Select the filter, and click **Edit**. The values for the filter appear in the Column, Operator, and Value boxes.
- 6 Edit the values you want, and then click **Add Filter**.
- 7 Click **Finish**.

Deleting a User Group Filter

To delete a user group filter:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and then click **Groups**. The Contents of Groups page opens, which displays the list of user groups and users.
- 3 Do one of the following:
 - Select the user group in the navigation frame.
 - Click the user group on the Contents of Groups page.
 - Click  (the Edit icon) next to the user group.

The Contents of Group_Name page appears, which displays the list of user groups and users contained in the user group.

- 4 Click **Edit Filter**. The Group Filters page opens.
- 5 In the Group Filters box, select the filter you want to delete, and then click **Delete**.
- 6 Click **Finish**.

Changing Log File Settings

To change log file settings:

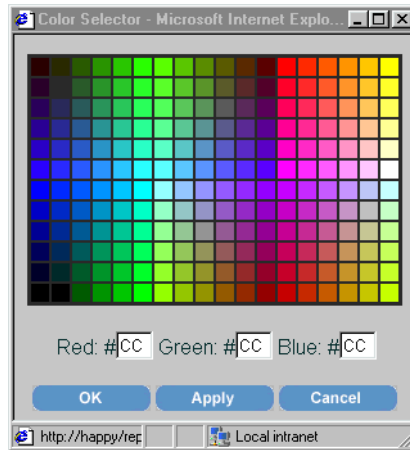
- 1 On the Links bar, click Administration.
- 2 In the Administration tree, expand **Log Files**.
- 3 Click **Settings**. The Log File Settings page opens.

Attribute	Value
Viewable Entries Per Page	<input type="text" value="500"/>
Log Frame Background Color	<input type="text" value="#CCCCCC"/> Select Color

[Apply](#) [Reset](#)

- 4 Do either of the following:
 - In the Viewable Entries Per Page box, type the number of log entries that you want to appear on each page of the log file. (The default is 500.)
 - In the Log Frame Background Color box, do one of the following:
 - Enter an HTML standard color name (aqua, black, blue, fuchsia, gray, green, lime, maroon, navy, olive, purple, red, silver, teal, yellow, or white) or a hexadecimal color value.

- Click **Select Color**. The Color Selector dialog box opens. Use it to select a color from the palette or type specific color values. After you select a color, click **Apply** then click **OK**.



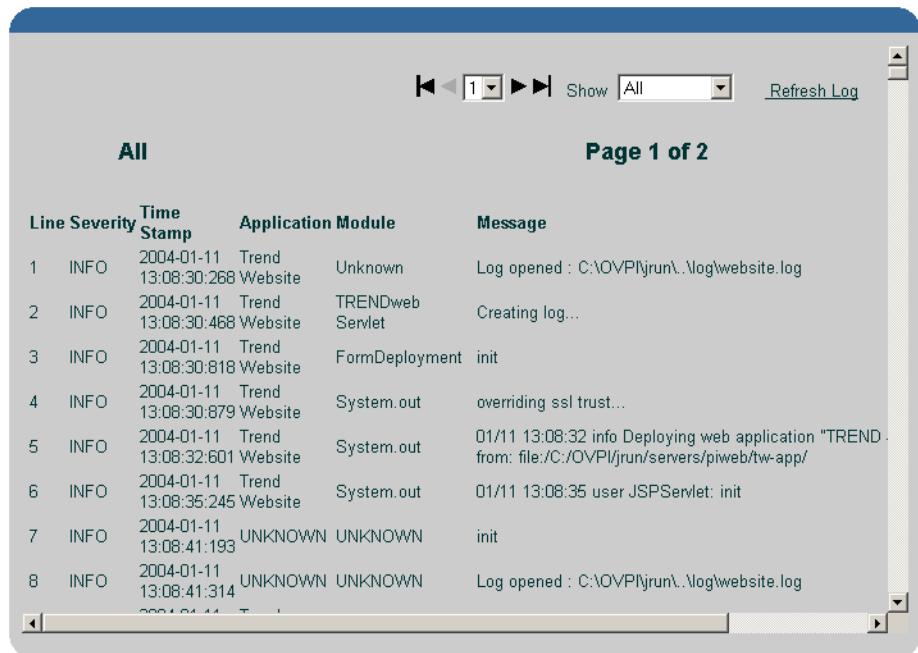
- 5 Click **Apply** again.

Viewing Log Files

To view a log file:

- 1 On the Links bar, click **Administration**.
- 2 In the Administration tree, expand **Log Files**.
- 3 Click **Web Access Server**.

The Viewing area, shown below, displays an example of the Web Access Server log file.



- 4 Do either of the following:
 - To page through the log file, click ► to go to the next page and ◀ to go to the previous page; click ◀◀ to go to the first page and ▶▶ to go to the last page.
 - To go to a specific page, click the pull-down arrow next to the page number, and then select the page of the log you want to view.
- 5 To select the type of messages you want to view, select the message type (All, Information, Warning, Error, and Debug) from the Show list.
- 6 To refresh the log file, click **Refresh Log**. This refreshes the copy of the log file you are viewing, displaying any new messages that were generated since you first displayed the log file.

Changing Java Settings

To change one or more of the Java settings:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Java Settings, and then click **Java Settings**. The Java Settings page opens.

Attribute	Value
Report Builder Heap Size	256 MB
Report Builder Java Arguments	-Djava.library.path="{ovpi.home}/bin
Report Viewer Heap Size	256 MB
Report Viewer Java Arguments	-Djava.library.path="{ovpi.home}/bin
Web Access Server Heap Size	256 MB
Web Access Server Java Arguments	-Djava.awt.headless=true -Dtrend.home={application.l

- 3 Do any of the following:
 - a In any of the Heap Size boxes, enter the maximum amount of memory that the clients and services are allowed to use.



The client or service may not necessarily use all the memory that is allocated.

- b In any of the Java Arguments boxes, enter the Java command-line arguments that you want to apply when the clients, services, or both are started.

For more information about these command-line arguments, refer to your Java documentation.

- 4 After you make all of your changes, click **Apply**.
- 5 If you changed options for any of the services, stop and restart those services.


Changing Report Settings

To change report settings:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Report Settings, and then click **Report Settings**. The Report Settings page opens.

Attribute	Value
Generated Reports Location	C:\OVPI\data\schedules
Deployed Reports Location	C:\usr\OVPI\reports\deploy
Image Cache Location	C:\OVPI\run\servers\piweb\tr-app\images
System Directory Alias	System
Default 'Download As' Drilldown Depth	5
Show System Directory to Users	<input checked="" type="checkbox"/>
Show Deploy Link to Users	<input checked="" type="checkbox"/>
Show Generated Results	<input checked="" type="checkbox"/>
Non-Admin Users Can Schedule Reports	<input checked="" type="checkbox"/>
Show All User Leaf Nodes	<input checked="" type="checkbox"/>
Show User Leaf Nodes in Results Frame Only	<input type="checkbox"/>
Deploy Reports With All Views	<input type="checkbox"/>
Enable Launchpad Filtering	<input checked="" type="checkbox"/>

- 3 Do one or more of the following:
 - a In the Generate Reports Location box, type the name of the folder in which you want to generate your schedules.
 - b In the Deployed Reports Location box, type the name of the folder in which you want to keep deployed reports.
 - c In the Image Cache Location box, type the name of the folder in which you want to keep images (labels and so on) when you view a report.
 - d In the System Directory Alias box, type a name for the directory that contains the reports that the OVPI Administrator has deployed to the Web server. (By default, this is System.)


- e To change the default value for the drilldown depth displayed in the Select Drilldown Depth dialog box, type the new value in the Default 'Download As' Drilldown Depth box.
-  A value of -1 will generate all possibilities; this may result in a very large file set and is not recommended.
- f To show or hide the system directory, do one of the following:
 - To show the directory, select the **Show System Directory to Users** check box.
 - To hide the directory, click to clear the **Show System Directory to Users** check box.
- g To show or hide the Deploy link, do one of the following:
 - To show the Deploy link, select the **Show Deploy Link to Users** check box.
 - To hide the Deploy link, click to clear the **Show Deploy Link to Users** check box.
- h To show or hide the Generated Results folder that is accessed using the Catalog link, do one of the following:
 - To show the Generated Results folder, select the **Show Generated Results** check box.
 - To hide the Generated Results folder, click to clear the **Show Generated Results** check box.
- i To show or hide the Schedule link, do one of the following:
 - To show the Schedule link, select the **Non-Admin Users Can Schedule Reports** check box.
 - To hide the Schedule link, click to clear the **Non-Admin Users Can Schedule Reports** check box.
- j Do one of the following:
 - If you want users with administrative privileges to be able to view all users' folders, select the **Show All User Leaf Nodes** check box.
 - If you want users with administrative privileges to be able to view only their own folders, click to clear the **Show All User Leaf Nodes** check box.

A *leaf node* is any report or empty folder. When leaf nodes are hidden, they do not appear in the Catalog tree; however, you can access them from the results frame of the Web Access Server.

This option has no effect on users who do not have administrative privileges.

- k** To show the leaf nodes only in the results frame of the Web Access Server window, do one of the following:
 - To show the leaf nodes in the results frame, select the **Show Leaf Nodes in Results Frame Only** check box.
 - To prevent the leaf nodes from displaying in the results frame, click to clear the **Show Leaf Nodes in Results Frame Only** check box.
- l** Do one of the following:
 - If you want to make user folders display only in the results frame, select the **Show User Leaf Nodes in Results Frame Only** check box.

If you have many users, selecting this option can save time, because the browser does not have to load all the user records before displaying the Catalog navigation tree.
 - If you want to make user folders display in the navigation tree, click to clear the **Show User Leaf Nodes in Results Frame Only** check box.
- m** Do one of the following:
 - If you want to add all deployed reports to all existing views, select the **Deploy Reports With All Views** check box.
 - If you do not want to add any deployed reports to existing views, click to clear the **Deploy Reports With All Views** check box.
- n** Select the **Enable Launchpad Filtering** box if you want the Web Access Server's launchpad to check the database when it builds a list of reports to make sure that the values specified in the reports are located in the database.

 This option is only used for NNM integration.

- 4** After you make your changes, click **Apply**.

- 5 Stop and then restart the Web Access Server. See [Chapter 16, Managing Services](#) for details.


Creating Views

A *view* allows users to see a certain list of deployed reports in the System folder. When you create a user account, you assign a view that determines which reports that user can view. The view applies to any database that the user may access.

Creating a View

To create a view:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Report Settings and then click **Catalog View Manager**. The Catalog Views page opens.



- 3 Click **New**. The View page opens, displaying all the reports that are currently deployed to the System folder.

View

Name:


Description:


📁 Demo 🔍 📄

📁 FrameRelay_Service 🔍 📄

📁 PVC 🔍 📄

- pvc_NRTquickview.srep | Default ▾
- pvc_NRTsnapshot.srep | Default ▾
- pvc_availability.srep | Default ▾
- pvc_capacity.srep | Default ▾
- pvc_dailyquickview.srep | Default ▾
- pvc_dailyquickviewtoend.srep | Default ▾
- pvc_dailysnapshot.srep | Default ▾
- pvc_dailysnapshottoend.srep | Default ▾
- pvc_execsum.srep | Default ▾
- pvc_execsumbyregion.srep | Default ▾
- pvc_forecast.srep | Default ▾
- pvc_hotspots.srep | Default ▾
- pvc_monthlyquickview.srep | Default ▾
- pvc_monthlysnapshot.srep | Default ▾
- pvc_roundtripdelay.srep | Default ▾
- pvc_topten.srep | Default ▾

- 4 In the Name box, type a name for the view.
- Ensure that you do **not** use any of the following characters in any of the fields on this page: ampersand (&), backslash (\), comma (,), double quotation mark (“), forward slash (/), question mark (?), and single quotation mark (‘).
- 5 In the Description box, you can enter a description for the view. (A description is optional.)
- 6 To select a folder of reports for the view, do *one* of the following:
- Select the box next to any folder to select all the reports in the folder.
 - Click  (the Select All button) to select all the reports in the folder.

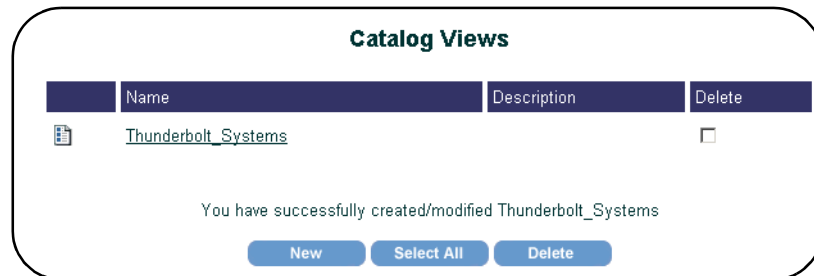
To deselect all the reports in a folder, click  (the Deselect All button).

- 7 To select a single report for the view, select the box next to the report or reports you want to include.

▶ If you select all the reports in a folder, the selection applies only to that folder and not to any of its subfolders. For example, if you selected the check box next to the ATM folder in the previous figure, the view will show only the reports in the ATM folder. Currently, no reports are deployed in the ATM folder. To include the reports in the PVC folder, you must select the check box next to it.

- 8 To specify the database from which the report should get its data, select a database from the *Database* list; there is one located next to each report. By default, each report will retrieve its data from the default database.
- 9 To select all the reports that are deployed to the *System* folder, click the **Select All** button at the bottom of the View page.
- 10 When you are finished selecting reports, click **OK** to create the view.


The Catalog Views page opens, indicating that your view was successfully created.



Modifying a View

To edit or delete a view:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Report Settings and then click **Catalog View Manager**. The Catalog Views page opens.
- 3 To edit a view, do the following:
 - a Select the view you want to edit by doing either of the following:

- Click  (the Edit button) next to the view you want to edit.
- Click the name of the view on the Catalog Views page.

The View page appears.

- b** Make your changes.
 - c** Click **OK** to apply your changes.
- 4** To delete a view, do the following:
 - a** Select the **Delete** check box next to the view you want to delete **or** click **Select All** to select all the views.
 - b** Click **Delete** to delete the view or views. A confirmation dialog box appears.
 - c** Click **OK**.

Changing the Appearance of the Web Access Server

Using the Site Appearance folder, you can change various options for the overall site, the link bar, the navigation frame, and the results frame.

Changing Site Options

To change the title of the Web Access Server:

- 1** From the Links bar, select **Administration**.
- 2** In the Administration tree, expand Site Appearance, and then click **Site Options**. The Site Options page opens.

Site Options

Attribute	Value
Site Title	<input style="width: 90%;" type="text" value="HP OpenView Performance Insight"/>
Show Profile	<input checked="" type="checkbox"/>

- 3 Do one or more of the following:
 - a In the Site Title box, type the title that you want to appear in your Web browser's title bar.
 - b To hide the Profile link on the link bar, click to clear the **Show Profile** box; to display the Profile link on the link bar, select the **Show Profile** box.
- 4 Click **Apply**, and then click your browser's Refresh button.

Changing Link Bar Options

To change the options for the Web Access Server's link bar:

- 1 From the Links bar, select Administration.
- 2 In the Administration tree, expand Site Appearance, and then click **Link Bar Options**. The Link Bar Options page opens.

Link Bar Options

Attribute	Value
Title	<input type="text" value="HP OpenView, Performance Insight"/>
Title Size	<input type="text" value="4"/>
Align	<input type="text" value="right"/>
Height	<input type="text" value="100"/>
Background Color	<input type="text" value="#FFFFFF"/> Select Color
Background Image	<input type="text" value="/webview/images/backgroundtile.gif"/>
Image	<input type="text" value="/webview/images/hpinvent.png"/>
Image Width	<input type="text" value="74"/>
Image Height	<input type="text" value="61"/>
Image Alt	<input type="text" value="Hewlett Packard"/>
Image X Location	<input type="text" value="25"/>
Image Y Location	<input type="text" value="25"/>
Image Background Width	<input type="text" value="200"/>
Link Background Focus On Color	<input type="text" value="#336699"/> Select Color
Link Background Focus Off Color	<input type="text" value="#6699CC"/> Select Color

[Apply](#) [Reset](#)

- 3 Do any of the following:
- a In the Title box, enter a title for the link bar. This title appears below the link bar.
 - b Click the Title Size arrow, and then select the point size you want for the title.
 - c Click the Align arrow, and then select an alignment option. This option aligns the link bar in relation to your Web browser's window.
 - d In the Height box, type a height in pixels for the link bar.
 - e In the Background Color box, do either of the following:
 - Enter an HTML standard color name (aqua, black, blue, fuchsia, gray, green, lime, maroon, navy, olive, purple, red, silver, teal, yellow, or white) or a hexadecimal color value.
 - Click **Select Color**. The Color Selector window opens from which you can select a color from a palette or enter specific color values. After you select a color, click **Apply** then **OK**.
 - f In the Background Image box, enter a location for the image that will appear behind the link bar. The image must be located on the Web server.

If you do not enter a location for the image, the Web Access Server looks for the image on the Web server in the `piweb` folder. For example, on Windows, the absolute path for this directory is as follows:

```
installation_directory\jrun\servers\piweb\tw-app\webview\images
```

In this instance, `installation_directory` is the directory into which you installed OVPI.

- g In the Image box, enter the location for the image that will appear to the left of the link bar. See [Step f](#) for information about where the image should be located.
- h In the Image Width and Image Height boxes, enter the width and height for the image in pixels.
- i In the Image Alt box, enter the text that you want to appear when the pointer remains on the image.

- j In the Image X Location and Image Y Location boxes, enter the x- and y-coordinate location (pixels) for the image in your Web browser window.

The top-left corner of your browser window is 0,0. By default, the coordinates place the image 25 pixels down and 45 pixels to the right.

- k In the Image Background Width box, enter the amount of space in pixels that the image will occupy in the link bar area.

- l In the Link Background Focus On Color and the Link Background Focus Off Color boxes, do one of the following:

- Enter an HTML standard color name (aqua, black, blue, fuchsia, gray, green, lime, maroon, navy, olive, purple, red, silver, teal, yellow, or white) or a hexadecimal color value.
- Click **Select Color**. The Color Selector window opens from which you can select a color from a palette or enter specific color values. After you select a color, click **Apply** then **OK**.

These options control the color of the tab that is located behind each link on the link bar. *Focus off* is when the pointer is not on a link; *focus on* is when the pointer is on a link.

- 4 When you have finished making changes, click **Apply**, and then click your browser's **Refresh** button.

Changing Link Bar Aliases

To change link bar aliases:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Site Appearance, and then click **Link Bar Aliases**. The Link Bar Aliases page opens.

Link Bar Aliases

Attribute	Value
1 Show Home Link	<input checked="" type="checkbox"/>
2 Show Catalog Link	<input checked="" type="checkbox"/>
3 Show Preferences Link	<input checked="" type="checkbox"/>
4 Show Administration Link	<input type="checkbox"/>
5 Show Schedule Link	<input checked="" type="checkbox"/>
6 Show Help Link	<input checked="" type="checkbox"/>
Home Link Alias	<input type="text" value="Home"/>
Catalog Link Alias	<input type="text" value="Catalog"/>
Preferences Link Alias	<input type="text" value="Preferences"/>
Administration Link Alias	<input type="text" value="Administration"/>
Schedule Link Alias	<input type="text" value="Schedule"/>
Help Link Alias	<input type="text" value="Help"/>

- 3 Do one or more of the following:
 - a To control the order in which the links will appear on the link bar, click the number arrow next to each link, and select a number.
The default order of the links is as follows: Home, Catalog, Preferences, Administration, Schedule, and Help.
 - b To hide a link on the link bar, click to clear a **Show *Link_Name* Link** box; to display a link on the link bar, select the **Show *Link_Name* Link** box.
 - c To change the name of a link, enter a new name in a ***Link_Name* Link Alias** box.
- 4 When you have finished making your changes, click **Apply**, and then click your browser's **Refresh** button.

Changing Navigation Frame Options

To change navigation frame options:

- 1 From the Links bar, select **Administration**.

- 2 In the Administration tree, expand Site Appearance, and then click **Navigation Frame Options**. The Navigation Frame Options page opens.

Attribute	Value
Navigation Width	<input type="text" value="300"/>
Navigation Margin Color	<input type="text" value="#336699"/> <input type="button" value="Select Color"/>
Navigation Background Color	<input type="text" value="#CCCCCC"/> <input type="button" value="Select Color"/>

- 3 Do one or more of the following:
 - a In the Navigation Width box, enter the width in pixels for the navigation frame.
 - b In the Navigation Margin Color and the Navigation Background Color boxes, do one of the following:
 - Enter an HTML standard color name (aqua, black, blue, fuchsia, gray, green, lime, maroon, navy, olive, purple, red, silver, teal, yellow, or white) or a hexadecimal color value.
 - Click **Select Color**. The Color Selector window appears from which you can select a color from a palette or enter specific color values. After you select a color, click **Apply** then **OK**.
- 4 When you have finished making changes, click **Apply**, and then click your browser's **Refresh** button.

Changing Results Frame Options

To change results frame options:

- 1 From the Links bar, select **Administration**.

- 2 In the Administration tree, expand Site Appearance, and then click **Results Frame Options**. The Results Frame Options page opens.

Margin Color —

Background Color —

Attribute	Value
Results Margin Color	#336699 <input type="button" value="Select Color"/>
Results Background Color	#CCCCCC <input type="button" value="Select Color"/>

This page lets you change the colors used for the results frame of the Web Access Server. The margin color affects the horizontal bar at the top of the frame; by default, this is blue. The background color affects the color of the entire frame; by default, this is gray.

- 3 In the Results Margin Color or Results Background Color boxes, do either of the following:
 - Enter an HTML standard color name (aqua, black, blue, fuchsia, gray, green, lime, maroon, navy, olive, purple, red, silver, teal, yellow, or white) or a hexadecimal color value.
 - Click **Select Color**. The Color Selector window opens from which you can select a color from a palette or enter specific color values. After you select a color, click **Apply** then **OK**.
- 4 When you have finished making changes, click **Apply** in the Results Frame Options page, and then click your browser's **Refresh** button.

Deploying Reports

To deploy a report:

- 1 From the Links bar, select **Catalog**.

- 2 In the Catalog tree, click **Deploy New Report**. The Report Deployment page opens.

Report Deployment

Name	Value		Summary
Report File	<input type="text"/> <input style="border: none; background-color: #95a5a6; padding: 2px 5px;" type="button" value="Browse..."/>		Local report file to deploy
Save To	<input type="text" value="/users"/> <input type="text"/> <input style="border: none; background-color: #3498db; padding: 2px 5px;" type="button" value="Browse"/>		Path where the report will be saved
Display Name	<input type="text"/>		Display name shown in the report hierarchy
Description	<input type="text"/>		Optional report description

- 3 Do one of the following:
- In the Report File box, enter the path and file name of the report to deploy
 - Click **Browse** to locate and select the report to deploy.

- 4 In the Save To box, do the following:

a Click the Save To arrow, and select **system** or **users**.

b Do one of the following:

- Enter the folder in which you want to deploy the report.

If you are deploying a report to the `Users` folder, ensure that you enter a folder that represents a username from an existing user account.

- Click **Browse** to locate and select a folder in which to deploy the report.

After the report is deployed, one of the following will occur:

- The report is deployed in the specified folder under the `System` folder.
- The report will appear in the user's folder under the `Users` folder.

- 5 In the Display Name field, enter a name for the report. This will be the name used for the report when you view it using the Catalog link.

If you do not supply a display name, the report's file system name is used. For example, if you deploy a report named `ncanfield.rep` and do not supply a display name, `ncanfield.rep` is used for the report's name when you go to view the report using the Catalog link.

- 6 In the Description field, enter an optional description for the report.
- 7 Click the **Deploy** button.

After the report is successfully deployed, the following message is displayed.

```
You have successfully deployed "report_name".
```

In this instance, *report_name* is the name of the report you entered in [Step 3](#) of this procedure.

Changing Settings for the Services

This section explains how to change settings for the following:

- Web Access Server ([page 329](#))
- SSL ([page 333](#))
- System Manager ([page 337](#))
- Mail Server ([page 345](#))

Changing the Web Access Server Settings

You can change the following settings for the Web Access Server: port number, realm, website document root, log file path, logging, debugging, headless operation, time-out value, and authentication method.

To change the settings for the Web Access Server:

- 1 From the Links bar, select **Administration**. The Administration page appears.

- 2 In the Administration tree, expand Service Configuration and then click **Web Access Server**. The Web Access Server Settings page opens.

Attribute	Value
Port	80
Realm	tw-app
Website Document Root	/
Log File Location	C:\OVPNlog
Enable Logging	<input checked="" type="checkbox"/>
Enable Debugging	<input type="checkbox"/>
Timeout	3 minutes
Authentication Method	HTTP Basic
Account ID	guest

Apply

- 3 Do one or more of the following:
 - a In the Port box, enter the port number of the Web Access Server.
 - b In the Realm box, enter a realm name for the server; it can contain up to 50 characters.

The realm name is used to authenticate users. A user's logon information (username and password) works only for the realm to which it is assigned. You should change the realm name if you are using a proxy server in front of the Web Access Server.

For example, if a proxy server's realm name is `realm1` and the Web Access Server's realm name is `tw-app`, a user must log on twice. If you change the Web Access Server's realm name to `realm1`, the user has to log on only once.

The realm name appears in the Enter Network Password dialog box; this dialog box appears after you click the **Log on** button on the Web Access Server.

- c In the Website Document Root box, enter a directory to relocate the home page's directory hierarchy. This changes the URL at which the Web Access Server's components are located.

By default, the hierarchy is / (root); all of the OVPI components are placed at the root level of the Web Access Server.

This option is useful in environments where there is a proxy server and OVPI is one of several components in the proxy. If you have several user applications, one of which is OVPI, you might want to change this to /ovpi. For example, this would change the default URL from

`http://hostname/reports/home` to the following URL:

`http://hostname/ovpi/reports/home`.

- d** To change the path of the server's log file, enter the path in the Log File Location box.
- e** To specify the logging status, do *one* of the following:
 - To enable logging, select the **Enable Logging** check box.
 - To disable logging, click to clear the **Enable Logging** check box.
- f** To specify the debugging status, do *one* of the following:
 - To enable debugging, select the **Enable Debugging** check box.
 - To disable debugging, click to clear the **Enable Debugging** check box.
- g** *UNIX only*: To control headless operation, do *one* of the following:
 - To enable headless operation, select the **Enable Headless Operation** check box.
 - To disable headless operation, click to clear the **Enable Headless Operation** check box.

After you install OVPI, your system is configured for headless support. *Headless support* uses software to render graphics; thus, no graphics card is required.

Alternatively, you can type either of the following commands to disable or enable headless support, respectively:

```
headlessctl enable
```

```
headlessctl disable
```

If you want to render your graphics using hardware, you can disable headless support if your system has a graphics card. Additionally, if you are using the Web Access Server on a UNIX system, you must configure your system for X. See [Configuring X for the Application Server on page 379](#) for details.

- h** In the Timeout box, enter the number of minutes that the Web Access Server should wait for an operation (for example, loading a report) to occur.

A valid value is any integer greater than zero and less than 1000.

- i** Click the **Authentication Method** arrow, and select one of the following:
 - **HTTP Basic.** This method requires that you log on to the Web Access Server by supplying the username and password for your user account in the Enter Network Password dialog box.
 - **URL.** This method requires that you supply a username when you view reports using a URL. For more information about using URLs to view reports, refer to the *Performance Insight Guide to Building and Viewing Reports*.
 - **None.** This method requires no authentication; however, users will access the system with the same privileges as the user account specified in the Account box.



If you specify a user account in the Account ID box that does not have Administrator privileges, you will no longer be able to access the Administration link. For more information about how to reset Administrator privileges, see [Resetting Administrator Privileges on page 332](#).

- **Account ID.** Enter the name of a user account. All other user accounts will acquire the same privileges as this account when accessing the Web server.
- 4** Click **Apply**.
 - 5** When you have finished your changes, stop and start the Web Access Server. See [Chapter 16, Managing Services](#) for this procedure.

Resetting Administrator Privileges

This section explains how to deal with an issue that may occur when you change the authentication settings for the Web Access Server.

If you have not specified any authentication for your Web server and have specified a user account without Administrator privileges in the Account box, you will need to reset the Administrator privileges to access the Administration link of the Web Access Server.

To reset Administrator privileges:

- 1 Open the following file with a text editor:

```
installation_directory\data\config.prp
```

In this instance, *installation_directory* is the directory into which you installed OVPI.

- 2 Find the following line in the config.prp file:

```
appserver.security=none
```

- 3 Change the line to the following:

```
appserver.security=basic
```

- 4 Navigate to the following location:

```
installation_directory\jrun\servers\piweb\tw-app\WEB-INF
```

In this instance, *installation_directory* is the directory in which you installed OVPI.

- 5 Make a copy of the web.auth file and paste the copy in the same directory.
- 6 Delete the web.xml file.
- 7 Rename the copy of the web.auth file to web.xml.
- 8 Stop and start the Web Access Server. For more information about how to stop and start the Web Access Server, see [Chapter 16, Managing Services](#).

Changing the SSL Service Settings

You can enable or disable Secure Socket Layer (SSL) communication between the Web Access Server and a supported Web browser.

To enable SSL, you create a self-signed digital certificate. This self-signed digital certificate will allow you to enable SSL; however, for increased security, you should send this certificate to a certificate authority in order to receive a signed digital certificate. Once you have a signed digital certificate, you must import it.

To change the settings for the SSL Service:

- 1 From the Links bar, select **Administration**. The Administration page appears.
- 2 In the Administration tree, expand Service Configuration and then click **SSL Service**. The SSL Service Configuration page opens.

Property	Value
Enable SSL	<input type="checkbox"/>
Secure Port	443
Keystore Password	<input type="text"/>
Certificate Information	<input type="checkbox"/> Generate a New Self-Signed Certificate (No Key Found)

Apply Reset

- 3 Do one of the following:
 - To disable SSL:
 - a Click to clear the **Enable SSL** check box.
 - b Click **Apply**.
 - c Stop and restart the Web Access Server to make the change take effect. See [Chapter 16, Managing Services](#) for this procedure.
 - To enable SSL, do the following:
 - a Select the **Enable SSL** check box.
 - b Leave the value in the Secure Port box as is (the default is 443) *or* change the port number in the Secure Port box.



If you enabled SSL communication when you installed OVPI, the Keystore Password box will not appear.

- c Enter the keystore password in the Keystore Password box.

- d** After you generate your self-signed digital certificate, it will be stored in the certificate repository. This password lets the Web Access Server access the certificate repository.
- e** Select the **Generate a New Self-Signed Certificate** check box to generate a certificate.
- f** Click **Apply**.
- g** Stop and restart the Web Access Server to make the change take effect. See [Chapter 16, Managing Services](#) for this procedure.

Importing Your Digital Certificate

After you enable SSL, you will have a self-signed digital certificate. A certificate is more likely to be trusted by others if it is signed by a Certification Authority (CA).

To do this, you must send a certificate signing request (CSR) to the CA, who will authenticate you, and then sign your certificate.

Generating a Certificate Signing Request (CSR)

To generate a CSR, do the following:

- 1** Navigate to the following directory:

```
installation_directory/jre/bin
```

- 2** Enter the following at the command line to generate your CSR:

```
keytool_path -certreq -file filename.csr -keystore  
installation_directory/jrun/lib/keyStore -alias ovpi
```

keytool_path is the location of the keytool utility in the DPIPE_HOME directory.

filename is the name you want to give to the .csr file.

installation_directory is the directory into which you installed OVPI.

For example, if OVPI is installed to /opt/OVPI, the command would be the following:

```
/opt/OVPI/jre/bin/keytool -certreq -file filename.csr  
-keystore /opt/OVPI/jrun/lib/keyStore -alias ovpi
```

The following prompt will display:

Enter keystore password:

- 3 Enter your keystore password, and press **ENTER**.

This is the same password that you provided when enabling SSL during the installation of OVPI or using the Web Access Server.

▶ When you type your password, it is echoed to the screen; thus, make sure you do not type it in front of anyone.

The `.csr` file now contains the request which should begin with the following line: -----BEGIN NEW CERTIFICATE REQUEST-----

Importing a Signed Digital Certificate

After your certificate has been signed, the CA will return it to you. You will then need to import your signed certificate into the keystore. It will override your self-signed certificate. Note that some Certification Authorities will ship a server certificate in addition to your signed certificate.

To import a signed certificate:

- 1 Navigate to the following directory:

installation_directory/jre/bin

- 2 *Server certificates only*: Enter the following at the command line to import the server certificate:

```
keytool -import -file filename -keystore installation_directory/  
jrun/lib/trustStore
```

- 3 Enter the following at the command line:


```
keytool -import -file filename -keystore installation_directory/  
jrun/lib/keyStore -alias ovpi
```

In this instance, *installation_directory* is the directory into which you installed OVPI (for example, `C:\ovpi`), and *filename* is the name of the certificate file from the CA.

The following prompt will display:

Enter keystore password:

- 4 Enter your keystore password, and press **ENTER**.

 When you type your password, it is echoed to the screen; thus, make sure you do not type it in front of anyone.

This is the same password that you provided when enabling SSL during the installation of OVPI or using the Web Access Server.

Changing System Manager Settings

The *System Manager* maintains the list of all Web Access Servers and databases. The Performance Insight Systems page displays all of your Web Access Server and database systems; it also lets you do the following:

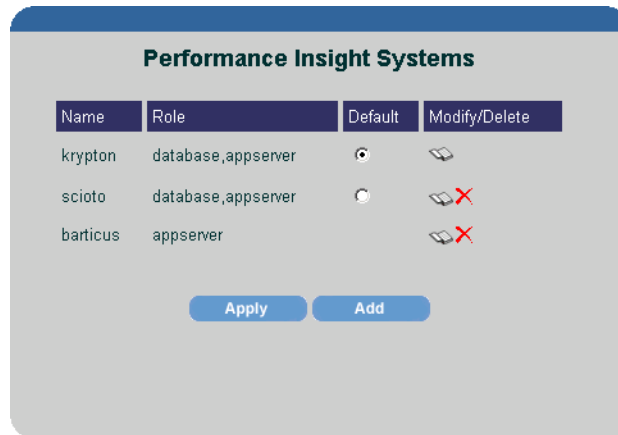
- Choose a default system ([Choosing the Default System on page 338](#)).
- Add systems ([Adding a System on page 338](#)).
- Modify the settings of current systems ([Modifying a System's Settings on page 341](#)).
- Delete systems ([Deleting a System on page 345](#)).

Displaying the Performance Insight Systems Page

To open the Performance Insight Systems page:

- 1 From the Links bar, select **Administration**. The Administration page appears.

- 2 In the Administration tree, expand Service Configuration and then click **System Manager**. The Performance Insight Systems page opens.



This page lists all Web Access Servers and databases to which your computer is connected. It also identifies the default system and specifies the role of each system. A system's Role can be as an Web Access Server (appserver) or both a database server and an Web Access Server.

From this page, you can make a selection that enables you to add, modify, or delete a system. You can also change the default system.

Choosing the Default System

The first system to which you establish a connection becomes the default system. To change the default:

- 1 In the Performance Insight Systems page, shown in [Step 2](#) above, click the **Default** radio button for the system that you want as the default system.
- 2 Click **Apply**.

Adding a System

To add a system:

- 1 Open the Performance Insight Settings page as described in [Displaying the Performance Insight Systems Page on page 337](#).
- 2 Click **Add**.

The System Information page opens.

Property	Value
Name	<input type="text"/>
Host Name	<input type="text"/>

Next Cancel

- 3 To enter the general properties of the system, do the following:
 - a In the Name box, type a name for the system.
 - b In the Host Name box, type the host name of the system.
 - c Click **Next**.

The System Roles page opens.

Property	Value
Database	<input type="checkbox"/>
Application Server	<input type="checkbox"/>

Next Cancel

- 4 To assign a role or roles to the system, do either or both of the following:
 - a Select the Database box to assign the role of database to the system.
 - b Select the **Web Access Server** box to assign the role of Web Access Server to the system.

c Click **Next**.

Depending on the role or roles you assigned to the system, one or both of the following pages displays:

- Web Access Server Properties (go to [Step 5](#))
- Database Properties (go to [Step 8](#))

5 If you selected the Web Access Server role, the Web Access Server Properties page displays.

Property	Value
Port	80
Context	/
Connect Using SSL	<input type="checkbox"/>
Authorization Realm	tw-app

Back Finish Cancel

6 If necessary, change any of the default values on the Web Access Server Properties page.

- In the Port box, change the port number for the Web Access Server.
- In the Context box, change the path to the Web Access Server.
- Click in the **Connect Using SSL** box if the database server is using SSL and you want secure communications with it.
- In the Authorization Realm box, change the realm name that the database server uses for authentication.

7 Click **Next** (or **Finish** if you did not select the Database role).

- 8 If you selected the Database role, the Database Properties page displays.


Property	Value
Admin Server Port	80
Authorization Realm	tw-app

Back Finish Cancel

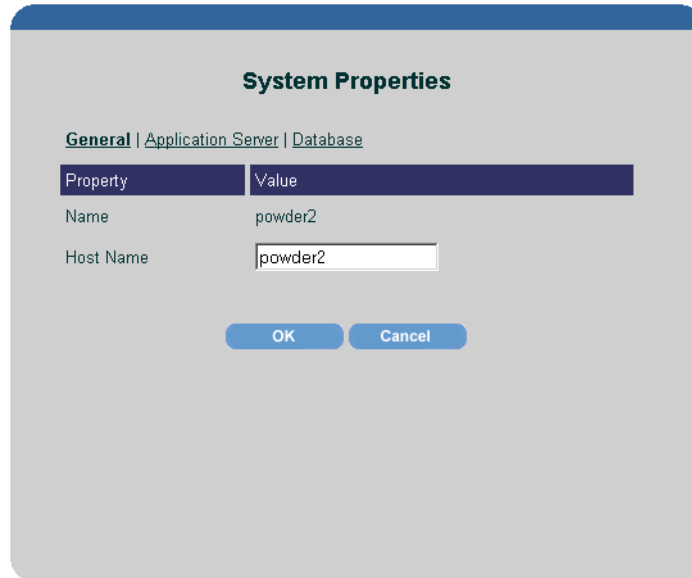
- 9 If applicable, do any of the following in the Database Properties page:
- In the Admin Server Port box, change the port number at which the Web Access Server listens for connections.
 - In the Authorization Realm box, change the realm name that the database server uses for authentication.
- 10 Click **Finish**.

Modifying a System's Settings

To modify the settings of a system:

- Open the Performance Insight Settings page as described in [Displaying the Performance Insight Systems Page on page 337](#).
- Click  (the Edit icon) next to the system that you want to modify.

The System Properties page opens.



The links that display just above the column headings represent the roles a system can assume.

- 3 To change the general properties of the system, do the following:
 - a Click the **General** link. The System Properties page opens.
 - b In the Host Name box, change the host name of the system.
 - c Click **OK**.
- 4 To change Web Access Server properties:

- a Click the **Web Access Server** link. The **System Web Access Server Properties** page opens.

The screenshot shows a dialog box titled "System Application Server Properties". At the top, there are three tabs: "General", "Application Server" (which is selected), and "Database". Below the tabs is a table with two columns: "Property" and "Value".

Property	Value
Use System Application Server	<input checked="" type="checkbox"/>
Port	<input type="text" value="80"/>
Context	<input type="text" value="/"/>
Connect Using SSL	<input type="checkbox"/>
SSL Port	<input type="text" value="443"/>
Authorization Realm	<input type="text" value="tw-app"/>

At the bottom of the dialog box, there are two buttons: "OK" and "Cancel".

- b If needed, do any of the following:
- To disable the system as the Web Access Server, click to clear the Use System Web Access Server box.
 - To enable the system as the Web Access Server, select the Use System Web Access Server box.
 - In the Port box, change port number for the Web Access Server.
 - In the Context box, change the path to the Web Access Server.
 - Click in the **Connect Using SSL** box if the database server is using SSL and you want secure communications with it.
 - In the SSL Port box, change the TCP/IP socket port that the server listens to for receiving and sending SSL requests.
 - In the Authorization Realm box, change the realm name that the database server uses for authentication.
- c Click **OK**.
- 5 To change the properties of a database:

- a Click the **Database** link. The System Database Properties page opens.

Property	Value
Use System Database	<input checked="" type="checkbox"/>
Description	<input type="text"/>
Port	<input type="text" value="5000"/>
Vendor	<input type="text" value="Sybase"/>
Database Instance	<input type="text" value="POWDER2"/>
Username	<input type="text" value="dsi_dpipe"/>
Password	<input type="text" value="dsi_dpipe"/>
Max Connections	<input type="text" value="10"/>

- b Do any of the following:
- To disable the system as the database, click to clear the **Use System Database** box.
 - To enable the system as the database, select the Use System Database box.
 - In the Description box, enter a description for the database.
 - In the Port box, enter the port number of the database.
 - In the Database Instance box, type the Database Instance. By default, they are as follows: DSQUERY (Sybase) or \$ORACLE_SID (Oracle).

This is the name of the database where the client applications will look for information by default.

For Sybase, you can find the name of this server in the Sybase interfaces file; this file is located on the system where the Sybase SQL Server is installed. For more information about the Sybase interfaces file, refer to the *Performance Insight Installation Guide*.


- In the Username box, enter the username for the database.

- In the Password box, enter the password for the database.
 - In the Max Connections box, type the maximum number of database connections that the database pool will maintain for this database.
- c** Click **OK**.
- 6** When you have finished making your changes, stop and start the Web Access Server.

For information about how to stop and start OVPI services, see [Chapter 16, Managing Services](#).


Deleting a System

To delete a system:

- 1** Open the Performance Insight Settings page as described in [Displaying the Performance Insight Systems Page on page 337](#).
- 2** Click  (the Delete icon) next to the system that you want to delete. A confirmation dialog box appears.
- 3** Click **OK**.

Changing Mail Server Settings

You can e-mail a report's URL to another user by clicking the **Email** button

 when you view a report. (Refer to the *Performance Insight Guide to Building and Viewing Reports* for details.) You must have an SMTP mail server to use this feature.

Any user can enable and disable this feature by using the Allow Report Emailing option on the Report View Settings page. To display this page, from the Links bar, select **Preferences**. In the Preferences tree, expand Reports and then click **Viewing**.

To change the settings for your SMTP mail server:

- 1** From the Links bar, select **Administration**. The Administration page appears.

- 2 In the Administration tree, expand Service Configuration and then click **Mail Server**.

The Mail Server Settings page appears.

Mail Server Settings

Attribute	Value
Host	<input style="width: 100%;" type="text" value="localhost"/>
Port	<input style="width: 100%;" type="text" value="25"/>
Header Message	<input style="width: 100%;" type="text" value="This message originated from HP OVPI User Account: <username>"/> <small>Use the <username> tag to substitute the username within the phrase.</small>
URL Message	<input style="width: 100%;" type="text" value="Report Link: <url>"/> <small>Use the <url> tag to substitute the report URL within the phrase.</small>

[Apply](#)

- 3 Do any of the following:
 - a In the Host box, enter the host name of the SMTP mail server.
 - b In the Port box, enter the port number of the SMTP mail server.
 - c In the Header Message box, enter the message that you want to appear in the Subject line of the mail message.

The <username> variable will place the username of the account from which the mail message is sent in the message.
 - d In the URL Message box, enter the way in which you want the report's URL presented in the text of the mail message.

The <url> variable places the report's URL in the message.
- 4 Click **Apply**.

Modifying Graph X-Axis Time Formats

You can set the x-axis time formats for graphs in the Web Access Server and Report Builder by editing the `graph.prp` file. For example, hourly data in the `config.prp` file is represented by `E hh:mm a` which translates into Day in Week, Hour:Minute, A.M./P.M. (for example, Wed 12:05 pm). [Table 1](#) lists the available variables.

Table 1 Graph X-Axis Time Format Variables

Symbol	Meaning	Format	Example
G	Era designator	Text	AD
y	Year	Number	1996
M	Month in year	Text & number	July & 07
d	Day in month	Number	10
h	Hour in A.M./P.M. (1~12)	Number	12
H	Hour in day (0~23)	Number	0
m	Minute in hour	Number	30
s	Second in minute	Number	55
S	Millisecond	Number	978
E	Day in week	Text	Tuesday
D	Day in year	Number	189
F	Day of week in month	Number	2 (2nd Wed in July)
w	Week in year	Number	27
W	Week in month	Number	2
a	A.M./P.M. marker	Text	P.M.
k	Hour in day (1~24)	Number	24

Table 1 Graph X-Axis Time Format Variables (cont'd)

Symbol	Meaning	Format	Example
K	Hour in A.M./P.M. (0~11)	Number	0
z	Time zone	Text	Pacific Standard Time
'	Escape for text	Delimiter	
"	Single quote	Literal	'

To set the x-axis time formats for graphs:

- 1 Edit the `graph.prp` file, which you will find in the following locations:

- *UNIX*:

`DPIPE_HOME/data/graph.prp`

- *Windows*:

`DPIPE_HOME\data\graph.prp`

In this instance, `DPIPE_HOME` is the directory into which you installed OVPI.

The time formats displayed in the `config.prp` file may appear similar to the following:

```
xaxis.hourly.time.format=E hh:mm a
xaxis.daily.time.format=E MMM d
xaxis.weekly.time.format=MMM d
xaxis.monthly.time.format=MMMMM
xaxis.summary.hourly.time.format=E hh:mm a
xaxis.summary.daily.time.format=E MMM d
xaxis.summary.weekly.time.format=MMM d
xaxis.summary.monthly.time.format=MMMMM
```

- 2 Change the formats you want using the variables listed in [Table 1](#).
- 3 After you make any changes, do the following:
 - a From the File menu, select **Save**.
 - b From the File menu, select **Exit**.

- 4 Stop and start the Web Access Server. See [Chapter 16, Managing Services](#) for details.

Tuning the System

This chapter describes things you can do to improve the performance of your HP OpenView Performance Insight (OVPI) system. Many factors can positively or negatively affect OVPI's performance. These factors can be grouped into the following areas:

- Hardware
- Operating system configuration parameters
- Sybase configuration parameters
- OVPI processing options



Not all identified areas apply to each installation. Factors such as platform and workload may also influence the overall effectiveness of the configuration area being discussed.

Hardware Considerations

OVPI is a database application that is designed to run on one to many databases. The hardware factors that affect overall performance are:

- Processor type and clock speed
- Number of processors
- Amount of memory

- Disk drive interfaces
- Number of physical disk devices

Processor Type and Clock Speed

In general, with slow processors, the processor is the major performance bottleneck. But, as the speed of the processor increases, disk and network I/O becomes the limiting factor for the overall speed of OVPI processing.

HP recommends that you use the fastest available processor for the selected platform, so that disk and network I/O become the limiting factor.

Number of Processors

OVPI and Sybase can be configured to take advantage of multiprocessor environments and run parallel (concurrent) tasks. Where possible, employ a multiprocessor configuration.

Memory

To get the best performance out of Sybase, you should obtain as much memory as possible, regardless of platform. Sybase uses memory mainly for disk cache and procedure cache buffers. The more available disk cache the less I/O a system does; thereby improving performance.

Table 1 lists the guidelines for selecting the amount of memory for a given OVPI system.

Table 1 Memory Guidelines

System Memory (per Processor)	Application
512 MB	Small to medium installations
1 GB	Medium to large installations
More than 1 GB	Very large installations

Shared Memory / Swap File Size

The amount of shared memory should be a value between the physical amount of memory and the swap file size. The swap file size should be 2 - 3 times the amount of physical memory in the machine.

Table 2 shows the maximum memory allocation for the database server on various operating systems. It specifies the maximum amount of shared addressable memory allowed in the Adaptive Server Enterprise 11.9.2.5 ESD 1 for each type of operating system.

Table 2 Maximum Allocation for Database Server

Operating System	Maximum Shared Addressable Memory
HP-UX	2.75 GB
Solaris	4 GB
Windows	2 GB

To improve performance you can prevent the system from swapping out the shared memory. The notes for the various operating systems follow:

- HP-UX
Set the following parameter: `SWAPMEM_ON = 1`. Note that this setting decreases the amount of shared memory available.
- Solaris
Use the Intimate Shared Memory (ISM) feature to lock the shared memory. Note that the Sun4c family of machines does not support this feature at this time.

Disk Drive Interfaces

If the system configuration has enough memory, the next likely bottleneck is disk I/O.

Choose the appropriate disk interface (controller) to ensure that data can move quickly from the disk to memory. In general, Fiber Channel and SCSI interfaces outperform IDE or Enhanced IDE interfaces. As with memory, the faster the interface the faster data can move on the I/O bus to memory for processing.

Faster interfaces also provide greater throughput for systems that have multiple I/O devices connected to the interface controller. Some platform vendors have optional (hardware) caching controllers that provide significant performance improvements over non-caching controllers. For larger systems, you can also use multiple interfaces.

Fiber Channel is the interface of choice for most medium and large installations. If you have a choice between using Fiber Channel or SCSI, choose Fiber Channel; otherwise, use SCSI 320 / 160 MB/sec LVD interfaces, if possible.

Hard drives should be at least 10,000 RPM.

Number of Physical Disk Devices

In general, the more disk drives the better the throughput, assuming that critical files are strategically placed on these drives. For example, if all data that is used for a normally operating system is placed on a single device and the remaining devices are used for backup, no benefit is derived from a multiple disk configuration.

When considering the use of multiple disks, make sure the database transaction log is on a separate device from the database itself. If you have a fast RAID volume, you can put the database and the transaction log on the same volume if there is not enough space available. Also, place the operating system's page/swap/virtual memory files on a separate device (or devices). If a fast I/O system is present, then increase the number of CPUs for Sybase.

Some vendors offer disk array products that not only have multiple drives but also provide sophisticated hardware that allows for hardware partitioning of database tables, imbedded cache, disk mirroring, RAID, and other options that improve throughput and data redundancy. These subsystems are the preferred choice for large installations or for small-to-medium installations where hardware budget is not a consideration.

The RAID should be set up with mirroring and striping (1+0); striping and mirroring (0+1); or mirroring, striping, and checksum (5). The 1+0 and 0+1 setups use more space than 5 setup, but they are faster. However, the 1+0

setup is easier to recover. If you use RAID, set the stripe size to 128 KB and the number of columns equal to half the number of disks in the volume. In other words, $number_of_columns = number_of_disks / 2$.

Configuring Operating System Parameters

Depending on the type of operating system that you are using, you may need to adjust various configuration parameters. The notes for the various operating systems follow:

- HP-UX

Various algorithms and formulas throughout the kernel use the **maxusers** parameter. The value of this parameter is the limit for the system resource allocation rather than the actual number of users allowed on the system. It is also part of the definition for some system table sizes and other parameters such as **nproc**, **ncallout**, **ninode**, and **nfile**. The **nproc**, **ncallout**, **ninode**, and **nfile** parameters should have fixed values. Set **maxusers** to 124.

- Solaris

The default value of the **maxusers** parameter is the number of megabytes of physical memory or 1024, whichever is lower. You cannot set it higher than 2048. If you want to adjust the number of user connections, use the **ulimit** utility to set the **nofiles (descriptors)** parameter.

Configuring Oracle Parameters

This chapter does not explain how to tune Oracle. The OVPI installation program provided defaults for several Oracle database parameters. (Refer to Chapter 3 of the *Performance Insight Installation Guide* for these parameters.) If you experience problems, consult with your Oracle database administrator.

Configuring Sybase Parameters

This section has two components:

- It describes the Database Tuner tool, which you access using the Management Console. This tool lets you view and modify many of the parameters that affect overall Sybase throughput.
- Some configuration parameters are not accessible through the Database Tuner. This chapter also explains the use of the Sybase configuration command, `sp_configure`, which you can use to display and modify these parameter settings.



When performance tuning any system or application, you change only one parameter at a time. If you change more than one at a time and the net change results in degraded performance, you will not be able to easily identify the change that caused the negative effect.

Using the Database Tuner

This tool lets you view and modify many of the parameters that affect overall Sybase throughput.

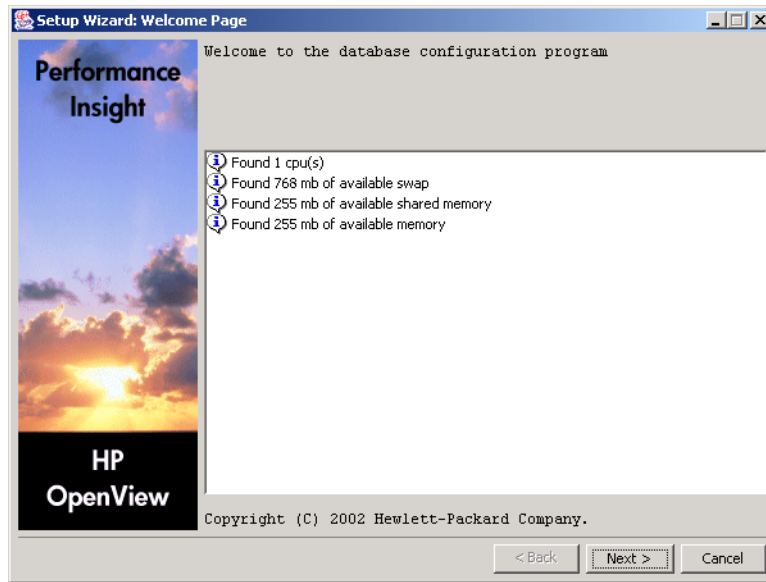


The Database Tuner is only available for Sybase databases.

To use the Database Tuner:

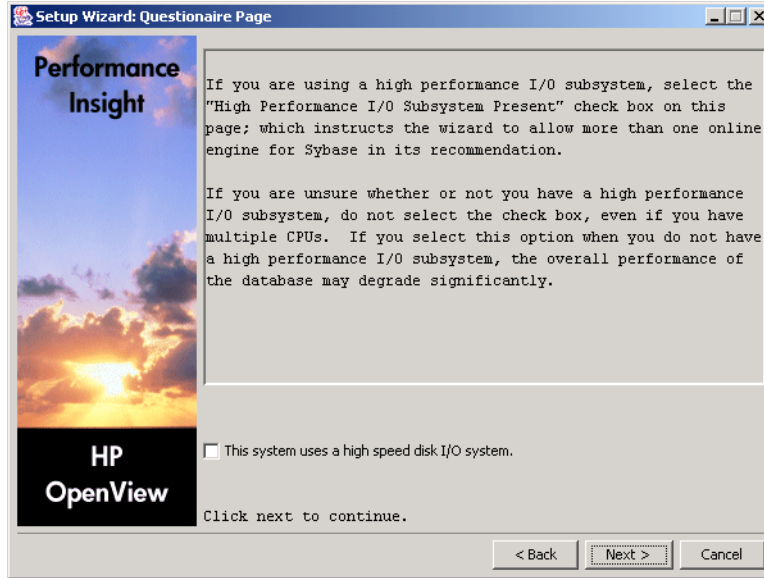
- 1 *UNIX systems only*: Log on to the system where you will run the Database Tuner as the `root` user.
- 2 Run the Database Tuner using either of the following methods:
 - Command line
Open a shell window (UNIX) **or** a Command Prompt window (Windows), and type `dbtuner`.
 - Management Console
Start the Management Console. (See [Starting the Management Console on page 76](#) for details.) Log on to the Management Console using an account with administrator privileges, and then choose **Database Tuner** from the Tools menu.

The Welcome window lists some of the physical characteristics (for example, the amount of available memory) of the system on which you are running the wizard.



- 3 Click **Next**.

The Questionnaire window displays.



4 Do one of the following:

- If you are using a high-performance I/O subsystem, such as RAID, select the **This system uses a high speed disk I/O system** check box, and click **Next**.

If you select this check box, it will affect the Max Online Engines parameter. See [Max Online Engines on page 363](#) for details.

- If you are not using a high-performance I/O subsystem, click **Next**.

The Information window displays information about your Sybase setup.

Performance Insight

HP OpenView

Enter the following information to indicate your current Sybase setup.

Version:
11.9.2.5

Port Number:
5000

SA Password:
[]

Home Directory:
D:\Sybase [Browse...]

Click Next to continue.

< Back [Next >] Cancel

- 5 Review the information, make any changes, and then click **Next**.

The Configuration window opens.

Performance Insight

HP OpenView

Enter the following information to indicate your current database setup.

Parameter	Adjusted	Recommended
SQL Perfmom Integration	0	0
additional network memory	5013504	5013504
max network packet size	16384	16384
max online engines	1	1
number of devices	20	20
number of locks	80000	80000
number of open indexes	3000	3000
number of open objects	3000	3000
number of user connecti...	100	100
procedure cache percent	20	16
recovery interval in minu...	1	1
total memory	131	128

Accept Recommendations [Reset Page Values]

Change Value

Min: no selection
Max: no selection
Units: no selection

Click Next to continue.

< Back [Next >] Cancel

This page displays the configuration parameters with their current and recommended values. Values in the **Adjusted** column are the current settings. Values in the **Recommended** column are recommended for your Sybase database.

The values in the **Adjusted** column may appear in different colors; these colors indicate the following:

- Values that are in the recommended range appear in black.
- Values that exceed the recommendation appear in blue.
- Values that fall below the recommendation appear in red.

For more information about the parameters, see [About Sybase Parameters on page 361](#). This section will also tell you whether you need to restart Sybase to make a change to a parameter take effect.

- 6 Review the parameter values, and do either of the following:
 - Click **Accept Recommendations** to change the values in the **Adjusted** column to match the values in the **Recommended** column.
 - Change the value of any parameter by doing the following:
 - Select the parameter. The current value of the parameter appears in the **Value** box, which is located to the left of the **Change Value** button.

The minimum and maximum values allowed for the parameter, and the parameter's units of measurement (for example, bytes) appear below the **Change Value** button.
 - Edit the value, and then click **Change Value**.
 - Click **Next** to continue.

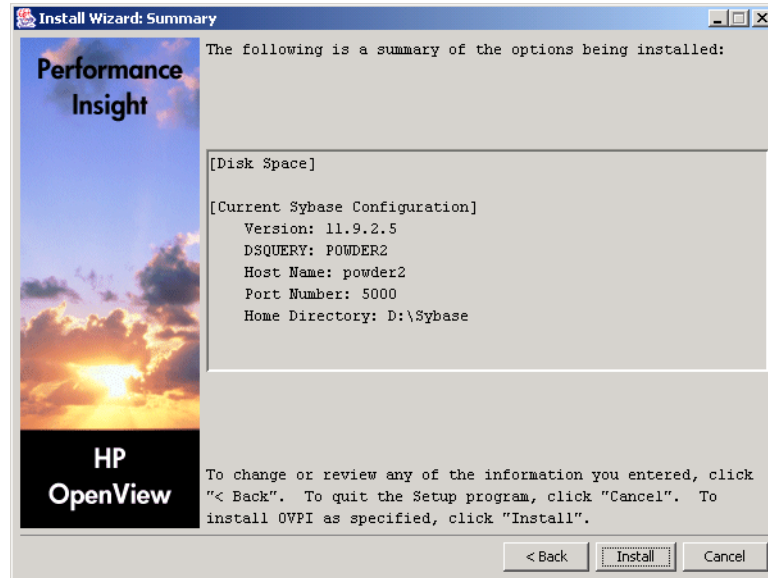
If needed, click **Reset Page Values** to reset parameters to their current values.

The new value appears in the Adjusted column, but the change will not take effect until you complete this procedure.



HP strongly recommends that you change only one parameter at a time. If you change more than one at a time and the overall change results in degraded performance, you will not be able to easily identify the change that caused the negative effect.

The Installation Wizard: Summary window displays.



- 7 Do the following:
 - a Review the configuration summary.
 - b If necessary, click **Back** to return to a previous page to make a change.
 - c Click **Install**.

About Sybase Parameters

This section describes the parameters that you can configure when performing [Step 6 on page 360](#) using the Sybase Configuration wizard. Refer to your *Sybase Administration Guide* for details.



When you make a change to a parameter that requires restarting Sybase for the change to take effect, you may restart Sybase after completing all of your changes or restart Sybase after completing each change.

SQL Perfmon Integration

The **SQL Perfmon Integration** parameter enables and disables the ability to monitor Adaptive Server statistics from the Windows NT Performance Monitor. A value of 1 enables Adaptive Server performance monitoring in the Windows NT Performance Monitor. A value of 0 disables it. Due to a limitation in Sybase, this value is set to 0 and cannot be changed.



This parameter is only applicable to Windows NT systems.

Additional Network Memory

The **additional network memory** parameter sets the maximum size of additional memory that can be used for network packets that are larger than the default packet size. Adaptive Server rounds down the value you enter to the nearest 2K value.

You may want to modify this parameter and max network packet size to increase data throughput between the server and client. Increasing these values may improve performance for applications that transfer large amounts of data.

If you increase **max network packet size** but do not increase **additional network memory**, clients cannot use packet sizes that are larger than the default size, because all allocated network memory is reserved for users at the default size. Users who request a large packet size when they log in will receive a warning message telling them that their application will use the default size.

This value is set to the recommended value and cannot be changed.

Max Network Packet Size

The **max network packet size** parameter specifies the maximum network packet size that can be requested by clients communicating with Adaptive Server. This value must be a multiple of 512.

If some of your applications send or receive large amounts of data across the network, these applications can achieve improved performance by using larger packet sizes. For best performance, choose a server packet size that works efficiently with the underlying packet size on your network. The two goals are:

- Reducing the number of server reads and writes to the network

- Reducing unused space in network packets (increasing network throughput)

If you increase **max network packet size** but do not increase **additional network memory**, clients cannot use packet sizes that are larger than the default size, because all allocated network memory is reserved for users at the default size. Users who request a large packet size when they log in will receive a warning message telling them that their application will use the default size.

If you change this value, you must restart Sybase to implement the change.

Max Online Engines

In single processor environments, Sybase is configured to use the **max online engines** default value of one (1) CPU. In multiprocessor environments, Sybase can be configured to use the additional CPUs. The use of multiple CPUs can cause a gain in throughput by allocating one or more processors for client processes.

- If your system does not use multiple CPUs with a high performance disk I/O system, do not change the `max online engines` default value of one (1) CPU; otherwise, performance may be negatively impacted due to internal contention.
- If your system uses multiple CPUs with a high performance disk I/O system; for example, RAID, you can take advantage of the multiple CPUs by increasing the value of the `max online engines`. In this case, the recommended value is derived by subtracting 1 from the number of CPUs your system uses. For example, if you have 4 CPUs, the value in the **Recommended** column will be 3.

If you change this value, you must restart Sybase to implement the change.

Number of Devices

This refers to the maximum number of devices that will be allowed to store the data in the database. The size of each device depends, to some extent, on the operating system, but it should rarely be greater than 2 GB. Because of this limitation, a large installation will require more devices than will a small installation.

If you change this value, you must restart Sybase to implement the change.

Number of Locks

The **number of locks** parameter sets the total number of available locks for all users on Adaptive Server.

The total number of locks needed by Adaptive Server depends on the number and nature of the queries that are running. If the total number of locks specified is insufficient, application problems such as the following can occur:

- Cannot obtain lock
- Failure to connect

Make sure the quantity of locks is adequate. By default, Sybase defines 5,000 locks. HP recommends a minimum of 20,000 locks. Implementations that want to run a large number of child collectors or parallel instances of **trendcopy** and **trendit** should use 40,000 to 80,000 locks.

If you change this value, you must restart Sybase to implement the change.

Number of Open Indexes

The **number of open indexes** is the maximum number of indexes than can be used simultaneously on the Adaptive Server. Because Sybase has changed the way this number is computed, OVPI sets the value to 3000 to prevent most related errors.

HP recommends keeping this value unless Sybase generates errors indicating that the pool of indexes is exhausted.

This value is set to the recommended value and cannot be changed.

Number of Open Objects

The **number of open objects** is the maximum number of objects than can be used simultaneously on the Adaptive Server. Because Sybase has changed the way this number is computed, OVPI sets the value to 3000 to prevent most related errors.

HP recommends keeping this value unless Sybase generates errors indicating that the pool of objects is exhausted.

This value is set to the recommended value and cannot be changed.

Number of User Connections

The **number of user connections** parameter sets the maximum number of user connections that can be connected to Adaptive Server at the same time. It does not refer to the maximum number of processes; that number depends not only on the value of this parameter but also on other system activity.

Choose the number of user connections carefully. Underallocation results in tasks being unable to connect to Sybase. Too many user connections waste memory that could be allocated to the data cache.

HP recommends 100 user connections. This is sufficient in most cases. In cases where memory is constrained, a more exact calculation may result in additional memory savings. To calculate the optimum number of user connections, add up the following:

- Total number of child SNMP collectors (`mw_collect -c` option). If there is no `-c`, the default is 5.
- Total number of child RMON collectors (`rmon_collect -c` option). If there is no `-c`, the default is 5.
- Total number of concurrent **trendit** processes desired.
- Total number of concurrent **trendcopy** processes desired.
- Total number of concurrent client processes desired.
- Add 5 to cover miscellaneous Sybase processes (for example, `bcp`, `isql`).

If you change this value, you must restart Sybase to implement the change.

Procedure Cache Percent

The **procedure cache percent** parameter specifies the percentage of memory allocated to the procedure cache after Adaptive Server's memory needs are met. Adaptive Server's memory needs are the sum of memory necessary for locks, user connections, the code itself, and other resources. The remaining memory is divided between the procedure cache and the data cache, with procedure cache being allocated first, based on the value of **procedure cache percent**.

Since the optimum value for procedure cache percent is different from application to application, resetting it may improve Adaptive Server's performance. For example, if you run many different procedures or ad hoc queries, your application will use the procedure cache more heavily, so you may want to increase this value.

If you change this value, you must restart Sybase to implement the change.

Recovery Interval in Minutes

The **recovery interval in minutes** parameter sets the maximum number of minutes per database that Adaptive Server uses to complete its recovery procedures in case of a system failure.

The recovery interval has no effect on long-running, minimally logged transactions (such as index creation) that are active at the time that Adaptive Server fails. It may take as much time to reverse these transactions as it took to run them.

Adaptive Server uses this setting and the amount of activity on each database to decide when to checkpoint each database. When Adaptive Server checkpoints a database, it writes to disk all data pages in cache that have been modified. This may create a brief period of high I/O. The checkpoint also performs a few other maintenance tasks, including truncating the transaction log for each database for which the **truncate log on chkpt** option has been set. About once per minute, the sleeping checkpoint process “wakes up,” checks the **truncate log on chkpt** setting, and checks the recovery interval to determine if a checkpoint is needed.

If you change this value, you do not need to restart Sybase to implement the change.

Total Memory

The total memory parameter sets the size of memory, in 2K units, that Adaptive Server allocates from the operating system. The more memory that is available, the more resources Adaptive Server has for internal buffers and caches, reducing the number of times the server has to read data from disk.

Before setting this value, be sure to assess the other memory needs on your system. When Adaptive Server starts, it must acquire the full amount of memory set by the total memory parameter. If you have set this value very high, and the full amount of memory is not available due to other memory demands, Adaptive Server will not be able to start.

If you change this value, you must restart Sybase to implement the change.



UNIX ports may require changes to shared memory configuration parameters.

Furthermore, take care to ensure that client tasks and Sybase do not cause the system to begin paging or swapping. You can check for this with Performance Monitor under Windows NT or by using `sar` or `vmstat` on the various UNIX ports.

4K Pool and 16K Pool

The **4K pool** and **16K pool** refer to data cache pools, which you can configure.

By default, Sybase ships with only one data cache called *default data cache*. This cache has only 2K buffers in it. The installation process will add more buffer pools to the default data cache. It will allocate 15 MB of the memory to a 4K buffer pool. It will then allocate 40% of the rest to a 16K buffer pool and 60% of the rest to the 2K buffer pool.

If you change either or both of these values, you do not need to restart Sybase to implement the changes.

Using the `sp_configure` Command

The **`sp_configure`** command lets you display and modify Sybase configuration parameters.



When you make a change that requires a call to **`sp_configure`** to change a Sybase parameter, the system copies the existing configuration file and creates a new configuration file with the requested change. For example, if you made 20 changes to existing values with **`sp_configure`**, you will have 20 additional configuration files; it does not matter if you change the same value 20 times or change 20 different values. Each file name will have a 3-digit extension that increments from `.001`.

To obtain a configuration listing from Sybase, type the following commands:

```
isql -Usa -P
1> sp_configure
2> go
```

The following sections list the parameters you may expect to see and recommendations for those parameters. When you make changes that require

- ▶ When you make changes that require restarting the server for the change to take effect, you may restart the server after completing all the changes or restart the server after completing each change.

Operating System - Asynchronous I/O Processing

Some operating systems require that you enable features to allow Sybase to make the most out of the available operating system resources.

For HP/UX, you should enable asynchronous I/O processing to ensure maximum I/O throughput for Sybase. The output will have a section that looks something like this

```
Group: O/S Resources
Parameter Name                Default Memory Used  Config Value  Run Value
-----
o/s asynch i/o enabled        0                    0             0           0
```

Notice the line that reads `o/s asynch i/o enabled`. If the run value is zero, then the operating system kernel does not have asynchronous I/O enabled. In this case, ask the System Administrator to enable it.

- ▶ The text of the output may differ slightly depending on the version of Sybase used.

Transaction Log Cache

All databases on a given server automatically use the default data cache for the transaction log. To improve performance, HP strongly recommends that you create a 4K buffer pool for the default data cache. In addition, create a dedicated log cache for the OVPI database. Use the information in [Modifying the Default Data Cache on page 368](#) and [Creating a Dedicated Log Cache on page 369](#) to complete these tasks.

Modifying the Default Data Cache

Use the commands listed below to create an 8MB, 4K buffer pool for the default data cache.


```
isql -Usa -P
sp_poolconfig "default data cache", "8M", "4K"
go
```

Creating a Dedicated Log Cache

Use the commands listed below to create a 12MB named-cache for OVPI database log pages only (`logonly` option).

```
isql -Usa -P
sp_cacheconfig "logcache", "12M", logonly
go
```

Once the named-cache has been created, bind it to the transaction log of the OVPI database using the procedure below.



The database must be set to single-user mode before you can bind the transaction log to the log cache.

- 1 Place the database into single-user mode as follows (be sure there are no active OVPI transactions running before proceeding):

```
isql -Usa -P
sp_dboption dpipe_db, "single user", true
go
```

The output looks something like this:

```
Database option 'single user' turned ON for database 'dpipe_db'.
Run the CHECKPOINT command in the database that was changed.
(return status = 0)
```

- 2 Switch to the `dpipe_db` database (use `dpipe_db`) and issue a checkpoint command followed by the `sp_bindcache` command as follows:

```
1> use dpipe_db
2> go

1> checkpoint
2> go

1> sp_bindcache logcache, dpipe_db, syslogs
2> go
```

- 3 Switch back to the master database to return the `dpipe_db` database to its normal (multi-user) state:

```
1> use master
2> go

1> sp_dboption dpipe_db, "single user", false
2> go
```

The output looks something like:

```
Database option 'single user' turned OFF for database
'dpipe_db'.
Run the CHECKPOINT command in the database that was changed.
(return status = 0)
```

- 4 Checkpoint the dpipe_db database by typing the following commands:

```
1> use dpipe_db
2> go

1> checkpoint
2> go
```

- 5 Allocate a buffer pool to the newly created cache (logcache):

```
isql -Usa -P
sp_poolconfig "logcache", "4M", "8K"
go
```

- 6 Restart Sybase to enable the new log cache.

Transaction Log I/O Size

If you are using a dedicated log cache, you need to modify the OVPI transaction log I/O size. If you are using a common log cache, you do not need to modify the OVPI transaction log I/O size.

Determine the OVPI transaction log I/O size by entering the following commands:

```
isql -Usa -P
1> use dpipe_db
2> go

1> sp_logiosize
2> go
```

The output of this command looks like this:

```
The transaction log for database 'dpipe_db' will use I/O size of 2
Kbytes.
(return status = 0)
```

By default, Sybase is configured to use a 4K transaction log. But, Sybase does not allocate any 4K buffers in the default cache configuration for the transaction log to use. Thus, a 2K transaction log is used for systems that have not changed the cache configuration. Sybase recommends a 4K log I/O size for user databases (particularly OLTP activity). Some OVPI installations manage better with an 8K log I/O size.

To change the log I/O size, allocate a cache pool for only the transaction log (8K for example) by entering the following command:

```
isql -Usa -P
1> use dpipe_db
2> go

1> sp_logiosize "8"
2> go
```

The output of this command looks like this:

```
The transaction log for database 'dpipe_db' will use I/O size of 8
Kbytes.
(return status = 0)
```


Traceflag Parameters

The 359 traceflag improves the performance of some queries in OVPI. It causes the server to select the correct index during procedure compilation; otherwise, the server may select an incorrect index, which causes index scans and may significantly degrade the performance. HP strongly recommends that you utilize the 359 traceflag parameter.

UNIX Systems

Add the 359 traceflag parameter, **-T359**, to the command line in the server startup run file as follows:

- 1 Log on as superuser (root).
- 2 Access Sybase by typing: **su - sybase**
- 3 Change to the installation directory by entering: **cd install**

- 4 Open the `RUN_servername` file by entering: `vi RUN_servername`
In this instance, `servername` is the name of the OVPI Sybase server.
- 5 Append, to the end of the line containing `${SYBASE}/bin/dataserver`, the traceflag parameter, `-T359`.
 -  All characters in the `-T359` string are contiguous. DO NOT use the space character when entering the string except to separate it from other parameters.
- 6 Save and close the file.
- 7 Stop and restart Sybase.
- 8 Verify that the traceflag is active by performing the following steps:
 - a Access the `install` directory under the Sybase home directory.
 - b Use your favorite editor to open the error log file and search for `Active traceflags`.
 - c Verify that `359` is in the list of active traceflags. Also, ensure that the timestamp on the `Active traceflags` entry corresponds to the actual date and time that the traceflag was added to the server startup run file.

Windows

Add the `359` traceflag parameter, `-T359`, to the registry as follows:

- 1 From the Windows task bar, click **Start**, and select **Run**.
- 2 In the Run dialog box, type `regedit` and click **OK**.
- 3 In the Registry Editor window, expand the list to the following directory:
`HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\
Server\servername\Parameters`
In this instance, `servername` is the name of the OVPI Sybase server.
- 4 Highlight the **Parameters** folder.
- 5 From the menu select: **Edit > New > String Value**.
The new value appears with a temporary name.
- 6 Type the name of the new value following typical naming conventions; for example, `Arg6` or `Arg7`, and then press **Enter**.

- 7 Double-click on the new value name.

The **Edit String** dialog appears.

- 8 In the **Value data** field, type: **-T359** and then click **OK**.



All characters in the `-T359` string are contiguous. **DO NOT** use the space character when entering the string except to separate it from other parameters.

- 9 Close the **Registry Editor** window.
- 10 Stop and restart Sybase.
- 11 Verify that the traceflag is active by performing the following steps:
 - a Access the `install` directory under the Sybase home directory.
 - b Use your favorite editor to open the error log file and search for `Active traceflags`.
 - c Verify that 359 is in the list of active traceflags. Also, ensure that the timestamp on the `Active traceflags` entry corresponds to the actual date and time that the traceflag was entered into the registry.

Other Configuration Parameters

The following configuration parameters should have their `Config Value` and `Run Value` set to the same value as their default value:

- `cpu grace time`
- `max number network listeners`
- `max async i/os per engine`
- `max async i/os per server`

To verify and set the default values for these parameters:

- 1 Type the following commands:

```
isql -Usa -P
1> sp_configure "config_param"
2> go
```

In this instance, `config_param` is one of the specified configuration parameters.

For example, if you were to list the value for cpu grace time, the output will look similar to the following:

Parameter Name	Default	Memory Used	Config Value	Run Value
cpu grace time	500	0	500	500

2 To set the default values, do the following:

```
isql -Usa -P
1> sp_configure "config_param", 0, "default"
2> go
```

In this instance, *config_param* is one of the specified configuration parameters.

3 Restart Sybase to make your changes take effect.

Binding the Property Tables

You can enhance performance by placing all property tables into their own cache. This way, when the property tables are referenced, no page I/O is required. To implement this, you must do the following:

- Create a named cache.
- Bind the individual property tables to the named cache.
- Verify that the bindings were successful.

Creating a Named Cache

You must create a named cache (for example, keycache) in which to put the property tables. It is advisable to do this after the OVPI database has reached its normal size to avoid under- or over-allocating the amount of cache necessary. After you create the named cache (keycache), you must restart Sybase for the named cache to be visible.

For example, to create a named cache called keycache and make the size of cache equal to 1 MB, type the following commands:

```
isql -Usa -P
1> sp_cacheconfig keycache, "1M"
2> go
```

Make sure that you restart Sybase. See [4K Pool and 16K Pool on page 367](#) to configure the cache you just created.

Binding the Individual Property Tables to the Named Cache

After you have created the named cache, you must bind the property tables to the named cache. You can select the property tables of your most frequently used data tables, property tables of your largest data tables, or any property table that would help the performance by reducing disk access.

To select property tables, type the following commands:

```
isql -Usa -P
1> use dpipe_db
2> go
```

To return the names of all data tables with their corresponding property table names, type the following commands:

```
1> select dsi_data_table, dsi_key_table from dsi_key_tables
2> go
```

After you have identified the names of all the property tables that you want to cache, repeat the following command for the name of each property table:

```
1> sp_bindcache keycache, dpipe_db, "key_tabl_name"
2> go
```

In this instance, *key_tabl_name* is the name of the property table to cache; you must enter the quotes.

For example, to bind the property tables **K_IFEntry_Disc**, **KIRDevPorts**, and **K_frswitch_port** to the cache named keycache, enter the following commands:

```
1> sp_bindcache keycache, dpipe_db, "K_IFEntry_Disc"
2> go

1> sp_bindcache keycache, dpipe_db, "KIRDevPorts"
2> go

1> sp_bindcache keycache, dpipe_db, "K_frswitch_port"
2> go
```

Verifying that the Bindings Were Successful

Type the following command to verify that the binding of the property tables was successful:

```
isql -Usa -P
1> sp_helpcache
2> go
```

The status should be v (valid) for each binding. The output will look similar to the following:

Cache Name	Config Size	Run Size	Overhead
default data cache	0.00 Mb	35.87 Mb	1.93 Mb
keycache	1.95 Mb	1.95 Mb	0.15 Mb

Memory Available For Named Caches	Memory Configured To Named Caches
37.84 Mb	1.95 Mb

There is 35.89 Mb of memory left over that will be allocated to the default cache

```
----- Cache Binding Information: -----
Cache Name  Entity                                     Type      Status
-----
keycache    dpipe_db.dsi_dpipe.Kmon_ethstats_         table     V
keycache    dpipe_db.dsi_dpipe.KDrouter_traffic       table     V
keycache    dpipe_db.dsi_dpipe.Kmon2_alhost_          table     V
keycache    dpipe_db.dsi_dpipe.KDhp_filesystem        table     V
keycache    dpipe_db.dsi_dpipe.Kmon2_almatrix_        table     V
keycache    dpipe_db.dsi_dpipe.Kdbaselineday         table     V
keycache    dpipe_db.dsi_dpipe.KHhp_compsystem        table     V
.
.
.
```


You can perform these functions as needed:

- Add or modify the named cache when you suspect that the cache allocation is out of date. For example, you may need to make the cache smaller or larger.
- Bind new property tables to the named cache.

Managing MIBs

The SNMP protocol supports the management of network resources through a data structure called a Management Information Base (MIB). MIBs have groups of managed objects, which are tables of variables that can contain a node's statistical values and configuration parameters. In this context, a *node* means any network communications accessible device that can be polled using SNMP.

Each node supports the MIBs that the node manufacturer feels are appropriate to the function of the node. Definitions for standard MIBs appear in Request for Comment (RFC) documents from the Internet Engineering Task Force. Definitions for enterprise MIBs come from the manufacturers of the nodes.

If you want to verify the information from a node, you can use a MIB Browser to query or modify the status and configuration information on it by making requests to the agent running on it. Furthermore, if you want to select a different set of objects to collect from the node, you can create a custom collection table.

You can use any MIB Browser that is available to you to query or modify the status and configuration information on a node; however, you need to use the HP OpenView Performance Insight (OVPI) MIB Browser to create a custom collection table.

This chapter describes how to do the following using the OVPI MIB Browser:

- Load a standard, vendor, or device-specific MIB. If the MIBs provided with OVPI are not sufficient for your data collection needs, you can load a custom MIB definition.

See [Loading MIBs on page 386](#).

- Display the contents of the currently loaded MIB file as a graphic MIB map tree similar to a UNIX or DOS directory structure. This tree structure lets you view and analyze the contents of the MIB.
- Conduct SNMP Get requests to query the status and configuration information on a node. The MIB Browser provides Get and Set functions. The MIB Browser can poll a node for a MIB group or object and display the response on the screen.

See [Querying SNMP Data on page 392](#).

- Set SNMP values.

See [Setting SNMP Values on page 397](#).

- Create a custom collection table by selecting objects from the loaded MIBs.

See [Creating Collection Tables on page 397](#).



In this chapter, the terms *objects* and *variables* are synonymous in the context of MIBs and refer to *MIB managed objects*.

Starting the MIB Browser

You can start the MIB Browser from the Management Console or the command line.

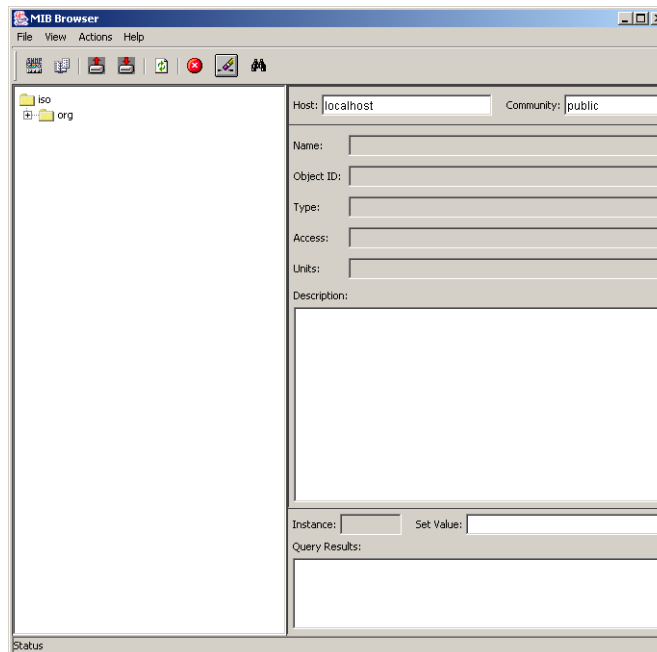
To start the MIB Browser, use one of the following methods:

- Management Console
 - a Start the Management Console. See [Starting the Management Console on page 76](#) for details.
 - b From the Tools menu, select **MIB Browser**. The MIB Browser window appears.
- Command line

Type the following command:

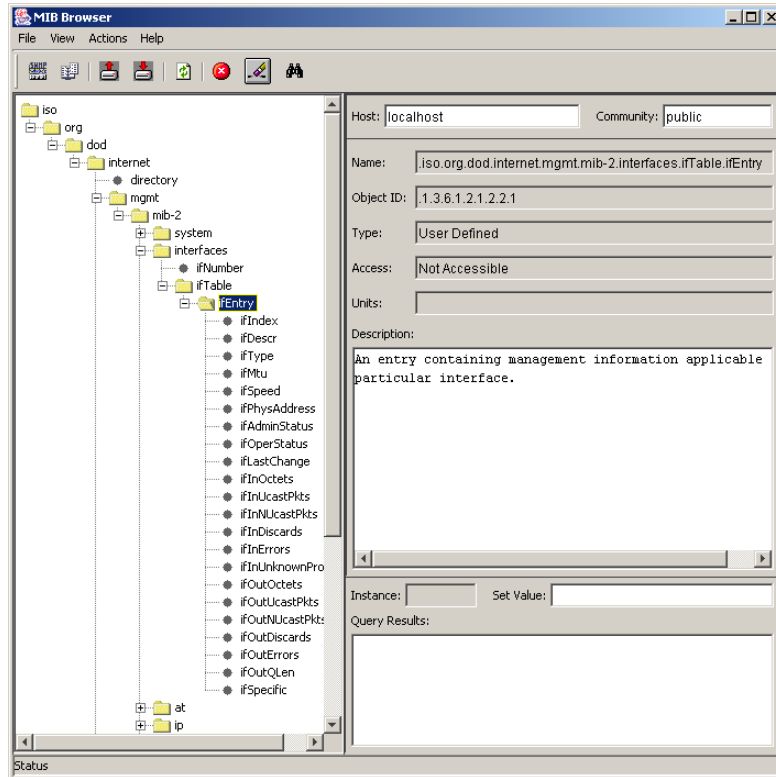
```
mibbrowser
```

The following figure shows an example of the initial MIB Browser window.




The MIB Browser Window


The MIB Browser window displays information about a selected MIB object. The following figure shows an example of the MIB Browser window with an expanded MIB tree.



The MIB Browser window has the following panes: The navigation and viewing panes.

You can adjust the size of the panes by moving the vertical line that separates the panes in either direction to the desired size. Do one of the following depending on your operating system:

- Windows. Place the cursor over the vertical line that separates the panes until you see a double arrow () and then move it in either direction to the desired size.

- UNIX. Point to the box on the vertical line that separates the panes (), left-click and hold, and then move in either direction to the desired size.

Navigation Pane

This pane shows the MIB tree, which is a map of all the currently loaded MIBs. When you load a MIB, it appears in the tree after you refresh the tree. When you select any object from the tree, the information about that object appears in the right pane of the MIB Browser.

You expand the hierarchical levels of the MIB tree just as you do when you view the contents of a Windows system using Windows Explorer. You can expand the folders to see the additional pages or folders it may contain, or collapse those levels.

You open the folders in the navigation pane as follows:

- To view the contents of a folder, click the plus sign (+).
- The folder's contents is displayed in the viewing frame. If you cannot see all of the objects in the MIB, use the scroll bar in the navigation pane to see the object.
- To close an open folder, click the minus sign (-) or double-click the folder.

Viewing Pane

This pane shows the details for the selected object in the tree. It shows the following information:

- Target host information. This displays the name of the node and its corresponding community string. This information is important when you use the SNMP Get and Set functions.

The information that is displayed is described as follows:

Host	This box identifies the target node. Type the name or IP address of the node that contains the MIB objects to select. The default value is localhost.
Community	This box identifies the current community string to access the host. The default value is public. Type the appropriate value in this box depending on the type of operation such as read or write.

- General information about the selected MIB object.

The information that is displayed is described as follows:

Name	This identifies the textual object identifier (OID) of the selected MIB object. It has each node of the path separated by a period and ending with the name of the selected managed object.
Object ID	This identifies the numeric OID, which is the path through the MIB tree to the selected object represented as a string of numbers separated by periods. The last value in the object identifier is the selected object position in its table.
Type	This identifies the SNMP data type assigned to the selected object. It specifies the storage type that OVPI will use in the reporting database and the algorithms used in processing the objects. It corresponds to the SYNTAX entry for each object in the MIB file.
Access	Identifies the type of permission to obtain the information from the selected object. The valid access modes are as follows: read-only, read-write, not-accessible.
Units	This specifies the type of quantity for the selected object.
Description	This provides the description for the selected object.

- SNMP Get request results. This contains an area to set a value with SNMP Set.

The information that is displayed is described as follows:

Instance	This identifies the selected instance of the managed object. You can select the instance from the Query Results box.
Set Value	Type a new value for the selected instance of the object. Use the SNMP Set action to actually change the value, if allowed. Remember to set the appropriate community string at the top in the Community field.
Query Results	This shows the result of the SNMP Get request for the selected managed object at that point in time.



Menu Bar and Toolbar

The MIB Browser has the following menus: File, View, Actions, and Help. There are also toolbar buttons for the most frequently used tasks.

File Menu

Table 1 list the options on the File menu and their corresponding toolbar buttons.

Table 1 MIB Browser File Menu Options

Option	Description	Toolbar Button
New	From the New submenu, you can select Collection . It lets you create your own SNMP collection table from a loaded MIB. See Creating Collection Tables on page 397 for details. This option is not required in standard OVPI installations. You do not need it if you are using OVPI solutions reporting.	
Open	This option lets you load MIB definition files into the MIB Browser. See Loading MIBs on page 386 for details.	
Exit	Exits from the MIB Browser.	Not applicable

View Menu

The View menu contains the following option






Refresh This option lets you reset the tree in the navigation pane and any information that appears in the SMNP Get request results section of the viewing pane.



Actions Menu

Table 2 lists the options on the Actions menu and their corresponding toolbar buttons.

Table 2 MIB Browser Actions Menu Options

Option	Description	Button
SNMP Get	This begins a GetNextRequest tree walk of the MIB tree for the group or object selected. It operates on scalar or columnar variables. See Querying SNMP Data on page 392 for details.	
SNMP Set	This lets you set a value for the selected instance from the Query Results box. It applies to scalar or columnar variables. See Setting SNMP Values on page 397 for details.	
Stop	Stops the output for the current SNMP Get request.	
Clear	Clears the output in the Query Results box.	
Find	This option lets you locate a variable that is in the MIB map in the left pane.	

Help Menu

The Help menu contains the About MIB Browser option, which displays version and copyright information for the MIB Browser.

Loading MIBs

OVPI stores MIBs in the default location: `DPIPE_HOME/mibs`; `DPIPE_HOME` is the environment variable that points to the directory where OVPI resides. When you load a MIB, you select a MIB from that directory or any directory that has MIB files.

For information about unloading MIBs, see [Loading a MIB on page 389](#).


Issues to Consider

You need to consider the following issues when you load a MIB:

- Including the top of the MIB tree
- Validating object types for collecting and aggregating data

Including the Top of the MIB Tree

The MIB file must include the fully qualified ISO tree down to the start of the MIB or the top of the MIB tree must be in a previously loaded file. If the top of the MIB tree is missing, you will see a message similar to the following:



In this example, the top of the MIB tree is in another file.

Validating Object Types

Sometimes MIBs use inappropriate object types. The most common examples of this are when fields that function as a Counter or a Gauge are defined as INTEGER.



Field labels are case sensitive.

If you want to collect the data from those fields and then correctly aggregate it, you must correct the field labels.

The DESCRIPTION field can indicate how the data in a field is used as shown in the following examples, where the first example shows the definition for Counter and the second shows the definition for gauge:

Example: DESCRIPTION Field for Counter

Fields that should be defined as Counters monitor how often something occurs over time. Their DESCRIPTION field often contains the word *count* or *total number*, as in Total number of octets since the device booted. For example, the definition for the following field should be a Counter:

```
sysLinkCellsTx
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION
"Count of ATM cells transmitted on the link"
::= { sysLinkEntry 10 }
```

The correction follows:

```
sysLinkCellsTx
SYNTAX Counter
ACCESS read-only
STATUS mandatory
DESCRIPTION
"Count of ATM cells transmitted on the link"
::= { sysLinkEntry 10 }
```

Example: DESCRIPTION Field for Gauge

Fields that function as a Gauge usually indicate how much or how many of something there is at any given point in time. Their DESCRIPTION field often contains the words *in use*, *used*, or *available*. For example, this field should be defined as a Gauge:

```
sysCurrentBuffers
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION
"The number of system buffers currently in use."
::= { sysLinkEntry 22 }
```

The correction follows:

```
sysCurrentBuffers
SYNTAX Gauge
ACCESS read-only
```

```

STATUS mandatory
DESCRIPTION
"The number of system buffers currently in use."
 ::= { sysLinkEntry 22 }


```




To create the database collection tables with the appropriate field types, correct these data types in the MIB file before you load the MIB. If the changed object appears in a SEQUENCE OF list, change the type in the sequence list as well.

Loading a MIB

To load a different MIB definition file:

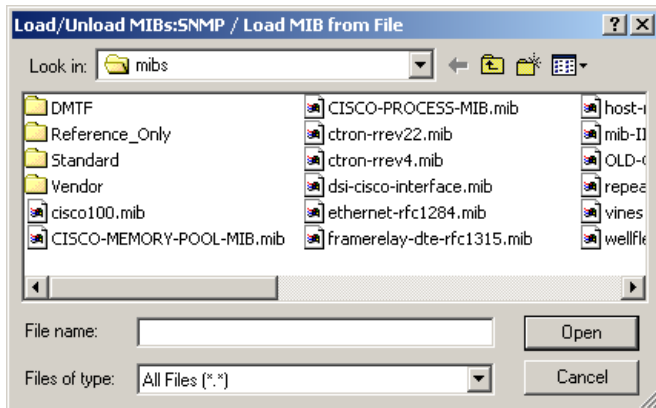
- 1 Do one of the following:
 - From the File menu, select **Open**.
 - Click .

The Load/Unload MIBs:SNMP dialog box opens, which shows the currently loaded MIBs.



- 2 Click **Load** to view the mibs directory, which is the default.


The Load/Unload MIBs:SNMP / Load MIB from File dialog box opens, which shows the mibs directory.



At this point, if the MIB is in a different directory, select the directory that contains the MIB.

- 3 Select a MIB, and click **Open**.


The MIB is loaded and appears in the Load/Unload MIBs:SNMP window. If the MIB does not load, you will see an error message similar to the following:



In this case, the error is a syntax error, and the message identifies the location of the error. If you get an error that you cannot resolve, contact the vendor to resolve the problem.

If the top of the MIB tree is in a separate file, select the file with the top of the MIB tree in it first, and then load the file that is missing the top of the MIB tree; otherwise, you will get an error. See [Including the Top of the MIB Tree on page 387](#) for details.


- 4 At this point, you can load another MIB by repeating [Step 2](#) and [Step 3](#) above until you have loaded all the desired MIBs.
- 5 Click **Close**.

- 6 Do one of the following:
 - Click .
 - From the View menu, select **Refresh**.
- 7 Locate the newly loaded MIB from the tree on the left pane if you want to view information about its elements.

After you load a MIB, it is available for use in OVPI until you use the unload procedure to remove it.


Unloading a MIB

To unload a MIB definition file:

- 1 Do one of the following:
 - From the File menu, select **Open**.
 - Click .


The Load/Unload MIBs:SNMP dialog box opens, which shows the currently loaded MIBs.

- 2 Select one or more MIB files from the Loaded SNMP MIBs box to unload.




You can select multiple MIBs from the list by clicking on the first one and then using the **Shift** key to select contiguous MIBs or the **Ctrl** key to select non-contiguous MIBs.

If you select a MIB file that has the top of the MIB tree in it and there is another MIB file that uses the selected file for the top of the MIB tree, you will see an unresolved object id error message similar to the following:



In this case, either select both MIB files to unload at the same time, or unload the dependent MIB file first.

- 3 Click **Unload**. A confirmation message opens, which is similar to the following:



When there are multiple MIBs selected, the number of selected MIBs appears in the message.

- 4 Click **OK**. The MIB is removed from the list. Any unloaded MIBs no longer appear in the map only after you exit the MIB Browser, and then restart it.
- 5 Click **Close**.

Querying SNMP Data

You can get and display data for the loaded MIB on the current node by choosing either the **SNMP Get** and **SNMP Set** options from the Actions menu. To access the data, you need to set the node and the community string.


Setting the Node

The current node, which appears in the Host box in the viewing pane, is the target for the SNMP functions. Enter the name or the IP address of the target node in the Host box. The default is localhost; you can change it by deleting the current name and typing a new name.

Setting the Community String

The MIB Browser uses defined community strings for a node to access its data. If OVPI knows about the current node, the MIB Browser uses the community strings for that node. New nodes use the default. The default community strings are public for read and private for write.


If you want to change the value, delete the current value and type a new value. Remember to set the value to the appropriate string for the operation you want to perform. For example, if you want to perform a read operation for the selected object, verify that the value in the Community box is public (if that is the appropriate value); otherwise, you will get an error message similar to the following message:




Finding a Group or Object in the MIB Map

You can locate a group or object in the MIB map if you know all or part of the name of the group or object.

To find the name of a group or object that you want to locate in the MIB map:

- 1 Do one of the following:
 - From the Actions menu, select **Find**.
 - Click .

The Find dialog box opens.



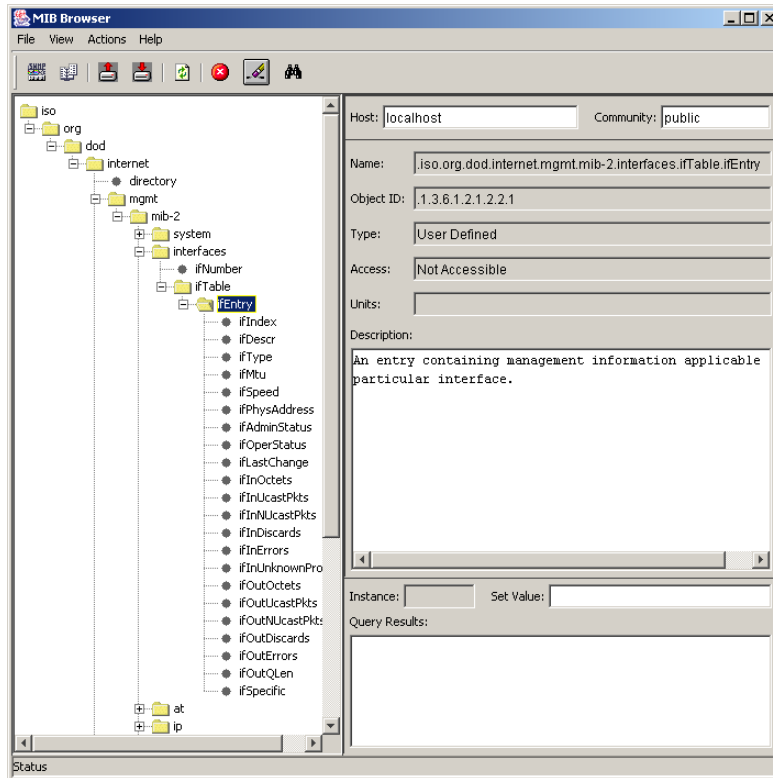
- 2 Type as much of the item name as you want or know in the *Find* box.
The default search ignores case differences between your entry and the actual MIB names because SNMP group and object names often have embedded capital letters.
- 3 If necessary, do the following:
 - If you want to locate the name exactly as typed with respect to uppercase and lowercase letters, select **Search is case sensitive**.
 - If you want to search for only the name that is in the box, click on **Match whole word**; otherwise, the default search compares your entry to any portion of the MIB names.
- 4 Click **Find**.
- 5 If you want to locate another instance of your entry, click **Find Next**.
- 6 When you are finished with the search, click **Close**.


Viewing the Data

If you want to confirm or verify the data that is on the node, you can use the **SNMP Get** option to view the current data on that node. The **SNMP Get** option provides data for scalar or columnar variables. To view the current data for the node specified in the **Host** field, do the following steps.

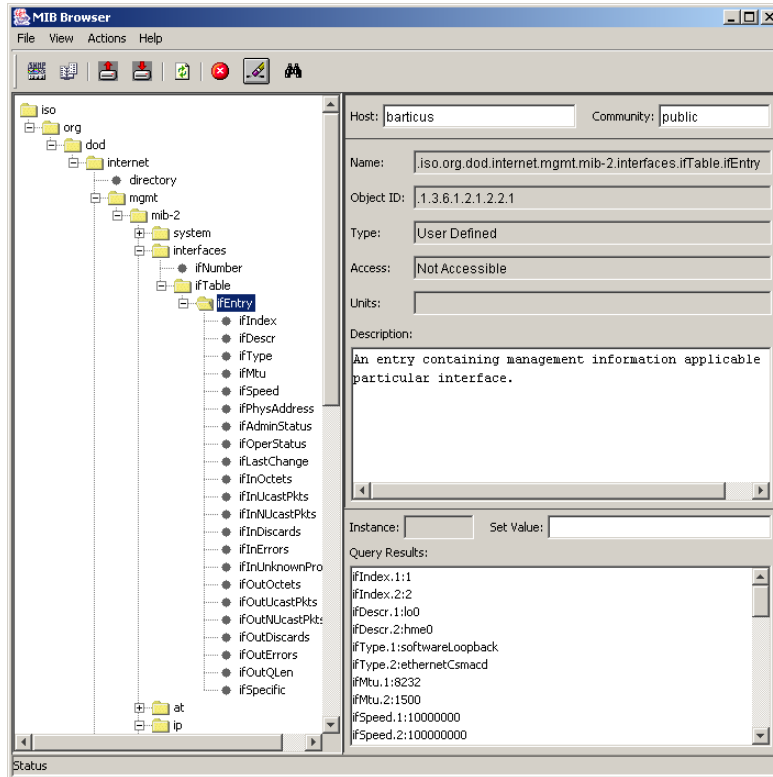
- 1 Select a MIB group or object by clicking the name in the MIB map.

The following figure shows the `mib-II.mib` file loaded with the **ifEntry** group folder selected.



- 2 To begin a GetNextRequest tree walk of the MIB tree for the group or object selected, do one of the following.
 - From the **Actions** menu, select **SNMP Get**.
 - Click .

The values appear in the bottom right corner of the MIB Browser in the **Query Results** field. The following example shows a partial listing of entries received for a GetNextRequest on the **ifEntry** group:



If you select a specific object, only that object is polled, but you get the value of that object for all instances of the table it is in.


If you select a group (folder), the MIB Browser polls for every object in every instance of that group, including the contents of any groups under the selected group. The data appears in the **Query Results** field grouped by object. That is, all the values for the first field are displayed, then all the values of the next field, and so on.

- 3 If you do not want to wait for the polling to finish, you can stop the GetNextRequest by clicking the **Stop** button or selecting **Stop** from the **Actions** menu. The polling stops and any values already received appear in the **Query Results** field.


Setting SNMP Values

You can set the values of writable MIB objects using the **Set Value** field on the MIB Browser. The current node is the target of the SNMP function.

To set a value:

- 1 Select the object name in the MIB map that you want to set.
- 2 Perform the **SNMP Get** function by doing one of the following.
 - From the **Actions** menu, select **SNMP Get**.
 - Click .

The result appears in the **Query Results** field.

- 3 Select the instance or variable that you want to change from the **Query Results** field. The instance or variable name appears in the **Instance** field and the value appears in the **Set Value** field.
- 4 Enter the new value for the object in the **Set Value** field.
- 5 Click .

If the set is successful, the changed value appears in the **Query Results** field.

Creating Collection Tables

The OVPI MIB Browser allows you to select a specific set of MIB managed objects from one or more groups so that the selected objects can be polled together in a single table; this is known as a *collection table*.

Collection Table Features

Collection tables have the following features:

- They are a composite set of MIB groups and/or objects selected for data collection.
- They provide a data source for custom designed reports.

- They allow you to focus your collection on data actually needed for reports, which increases polling efficiency, speeds up data collection, and reduces the amount of database space used.
- They allow you to generate a single report that includes objects from different groups in a MIB.
- They allow you to overlay a pair of 32-bit counters over a 64-bit field. See [64-Bit Values in SNMPv1 on page 398](#).

To create the collection table, you use the **Collection** option to select the objects to collect. When you create a collection table from the OVPI MIB Browser, the system creates the corresponding TEEL file for the table.



You do not need to create a collection table to run any of the existing solutions provided by OVPI. The solutions include any required collection tables.

64-Bit Values in SNMPv1

SNMPv2 introduced 64-bit counters; however, SNMPv1 agents only support 32-bit counters. One common way that SNMPv1 agents return the higher precision values stored in 64-bit counters is to overlay a pair of 32-bit counters over the 64-bit field and provide the full value in two pieces.

In the following example, a 64-bit counter called `sitCounter` is broken into a high order 32-bit field called `sitCtrSplit1High` and a low order 32-bit field called `sitCtrSplit1Low`. By collecting the two 32-bit fields and reassembling them, OVPI is able to store and process the full 64-bit value.

```
sitCtrSplit1High OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The high 32-bits of the 64-bit counter, sitCounter."
    ::= { sitEntry 7 }
```

```
sitCtrSplit1Low OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The low 32-bits of the 64-bit counter, sitCounter."
```

```

 ::= { sitEntry 8 }

sitCounter OBJECT-TYPE
    SYNTAX      Counter64
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The SectionInTransit units received since the device
        booted."
 ::= { sitEntry 9 }

```

It is critical that the two halves be reassembled in the correct order. *High Part* refers to the high order, most significant 32-bit portion of the 64-bit counter. *Low Part* refers to the low order, least significant part of the 64-bit counter.

64-bit counter	
High Order 32-bit counter	Low Order 32-bit counter

Collection Table Requirements

Collection tables have the following requirements:


- The maximum width is 1,962 characters, which is the total character count for all of its columns.
- The maximum number of columns is 250. Note that there are 13 default columns that you must allow for when you select the objects for the data table. If you choose to define a property table, there are seven (7) default columns that you must allow for when you select the objects for it.
- The objects must have a common index. For example, if one object counts errors per interface, then all objects must provide data per interface. Note that MIB objects in the same table have a common index by definition. A collection table that contains objects with different indexes, results in invalid data in your database.
- It must have a corresponding property table. You can select an existing property table from a list or create a new property table. If you create a new property table, you can use the default by-variables or select different by-variables.

This section contains the basic information for creating a collection table. If you want additional information about creating the tables, contact your HP OpenView sales representative to obtain the *Performance Insight TEEL Reference Guide*. This document provides information about the attributes of tables and creating TEEL files.

Creating a Collection Table

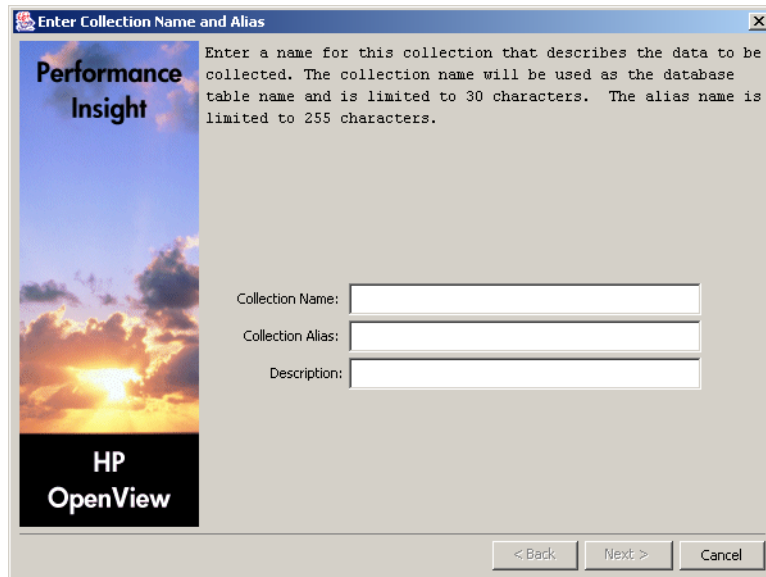
To create a custom collection table:

- 1 Open a collection dialog. Do one of the following steps.

- Click .
- From the **File** menu, select **New** ⇒ **Collection**.

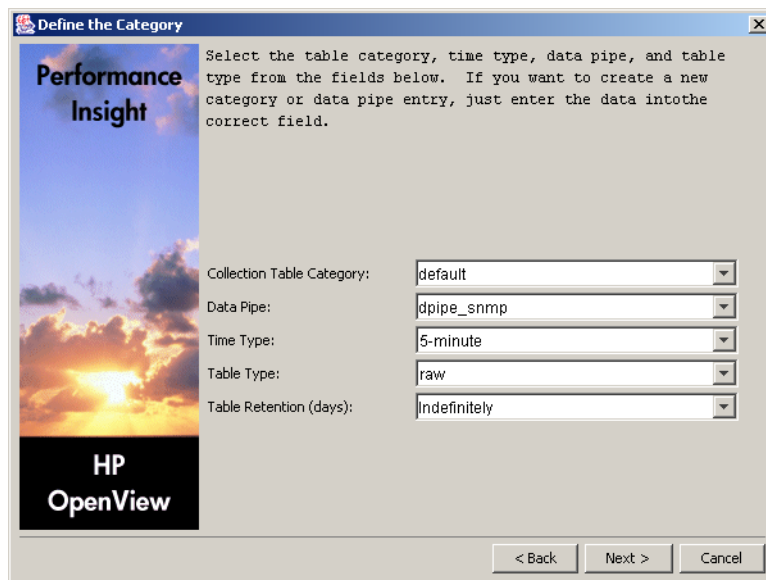
(A login dialog may appear at this point. Enter the login information to continue.)

The **Enter Collection Name and Alias** dialog box opens.



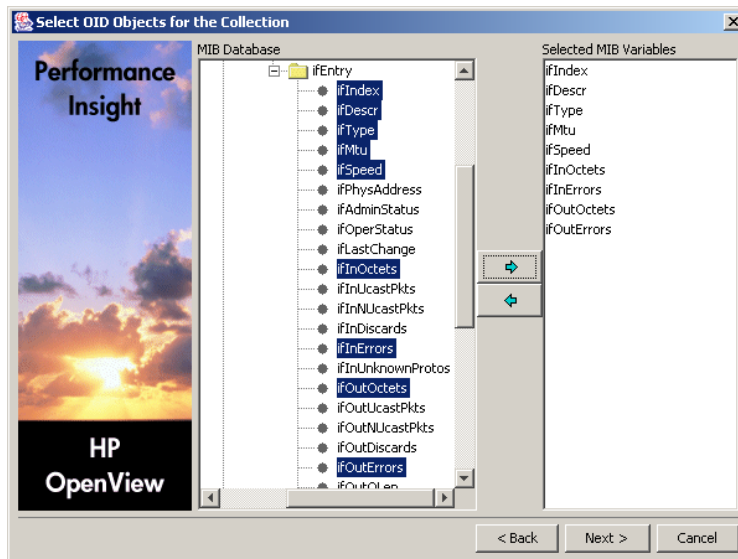
- 2 Enter the SQL name, Alias name, and description for the collection table.

- a In the **Collection Name** field, type the SQL name of the table. The name must be unique and can be up to 30 characters. Valid entries include uppercase and/or lower-case alphabetic characters, all numeric digits, and the special character underscore (_). This is a required field.
 - b In the **Collection Alias** field, type the alias name of the table. The name must be unique and can be up to 255 characters. Valid entries include uppercase and/or lower-case alphabetic characters, all numeric digits, and the special characters hyphen (-) and underscore (_). This is a required field.
 - c In the **Description** field, type a description for the table. It can be up to 255 characters, which can be blank. This is not a required field.
- 3 Click **Next**. The **Define the Category** dialog box opens.



- 4 Select the table category, data pipe, time type, table type, and table retention attributes of the table.
- a In the **Collection Table Category** field, select an existing category from the pull down list or type a new category name. The category name groups sets of tables together. If you do not know which name to choose, select `default`. This name can be up to 30 characters.

- b In the **Data Pipe** field, select an existing data pipe name from the pull down list. It identifies the collectable property of the table. Currently, this name can only be `dpipe_snmp`.
 - c In the **Time Type** field, select an existing time characteristic for the data in the table from the pull-down list. The default is `5-minute`.
 - d In the **Table Type** box, select a table type. For the available table type options, see [Step 4 on page 263](#).
 - e In the **Table Retention** field, select the number of days from the pull-down list or type the number of days to keep the data in the table. If you never want to age the data out from the table, select `Indefinitely` from the pull-down list.
- 5 Click **Next**. The **Select OID Objects for the Collection** dialog box opens.




- 6 Select the MIB objects.

If you previously selected the objects from the MIB Browser, those objects will appear in both panes. The highlighted objects will be in the tree in the left pane, which is the **MIB Database** field, and the names of the selected objects will be in the right pane, which is the **Selected MIB Variables** field.


If you did not previously select the objects or if you want to add more objects, you can select one or more objects from the tree in the **MIB Database** field at this time. You can select multiple objects consecutively by

holding the **Shift** key and clicking the left mouse button or individually by holding the **Ctrl** key and clicking the left mouse button. Each object becomes a column in the table.


Remember that you can only have a total of 250 columns, and there are 13 default columns; so, you can add up to 237 additional columns; otherwise, you will see an error message.

After you select the objects, click the right arrow button, , to add the selected objects to the **Selected MIB Variables** field.

If you select objects that exceed the total character count limit, which is 1,962 characters, the system will not add the objects that exceed the limit to the **Selected MIB Variables** field; you will see the following message.

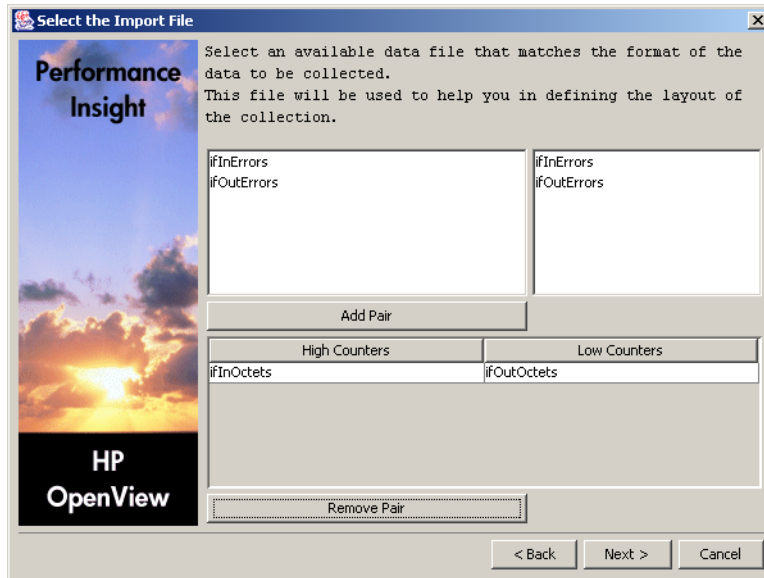


If you select objects that are not valid as columns, such as a folder, you will get an error message with the title, **Unable to add new column**, and it will list the objects; these objects will not appear in the **Selected MIB variables** list.



If you do not want a selected object, highlight that object in the **Selected MIB Variables** field and then click the left arrow button to remove the selected object from the list.

- 7 Click **Next**. The **Define the Hi-Lo Counters Pairs** dialog box opens, which shows all the selected counters in each column.



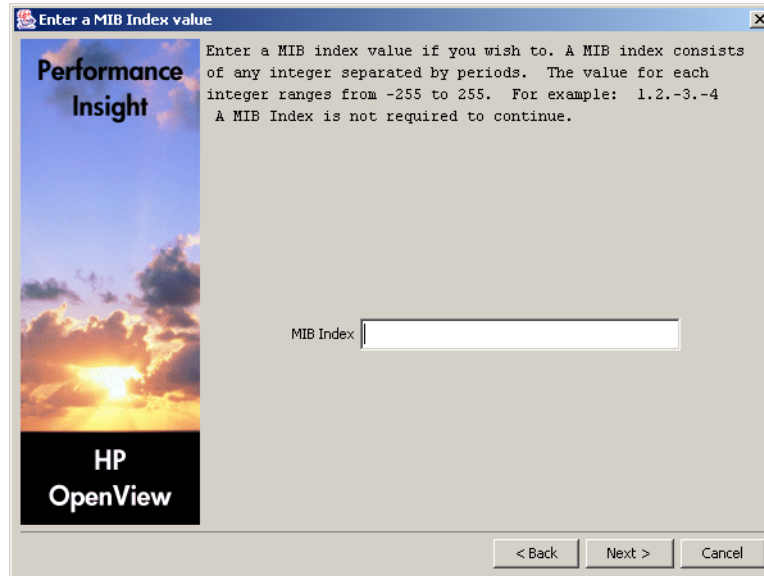
At this point, you can select pairs of counters to define 64-bit counters. See [64-Bit Values in SNMPv1 on page 398](#) for more information about 64-bit counters. This step is optional.

- a To define a 64-bit counter, select a counter from each column and click the **Add Pair** button.

Continue in this manner to add pairs as is necessary to define all the desired counters for the new collection table.

- b If you mistakenly add a counter pair, highlight it in the field on the bottom and click the **Remove Pair** button to remove it from the list.

- 8 Click **Next**. The **Enter a MIB Index value** dialog box opens.



- 9 Enter the MIB index value. (This step is optional.)

The MIB index consists of integers separated by periods. Each integer is the length of the index OID. If the length is variable, then the integer is 0 (zero). If the variable is not constant, such as RMON2 time filters or RMON history bucket numbers, the length is -1.

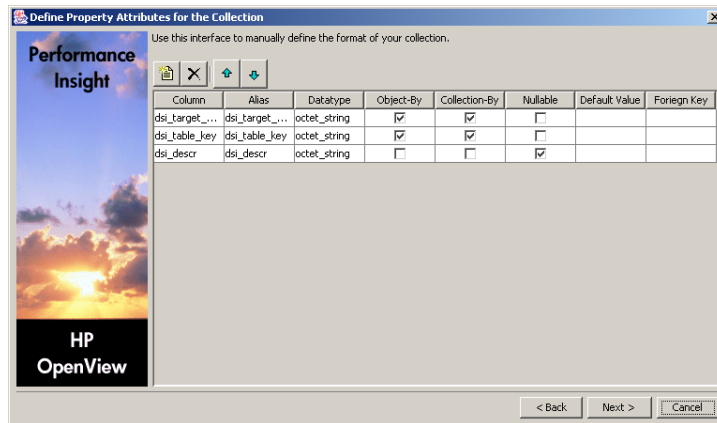
- 10 Click **Next**. The **Associate with a Property Table** dialog box opens.



- 11 Either select an existing property table from the pull down list or create a new property table.

If you select the **Create a new property table** option, type the SQL name for the new property table. The name must be unique and can be up to 30 characters. Valid entries include alphabetic characters, numeric digits, and underscore (_).


- a Click **Next**. The **Define Property Attributes** dialog box opens.



- b** You can create a property table with only the default columns or you can add more columns. In either case, you can modify the attributes of the columns. Note that if you modify any attribute, remember to leave the field with the change by clicking out of it before clicking the **Next** button.

If you want to use the default columns in the property table only, you can use them with the current settings or modify some of the attributes, as necessary. See the following steps to modify the various fields.

If you want to add additional columns to the property table, do the following steps:

- Click the **New** button, . A column line appears at the end of the existing list. A default name appears in the **Column** and **Alias** fields, which is `Column_n` where *n* is a number and it increments every time you add a new column.
- To specify the SQL name of the column, highlight the existing name in the **Column** field and type the new name. The name can be up to 30 characters and contain the following characters: alphabetic, numeric, or an underscore.
- To specify the alias name of the column, highlight the existing name in the **Alias** field and type the new name. The name can be up to 255 characters and contain the following characters: alphabetic, numeric, an underscore, or a hyphen.
- Select the data type for the column from the pull down list that appears when you click in the **Datatype** field.

For the available datatype options and their corresponding database storage types, see [Step 4 on page 263](#).

- In the **Data Length** field, type a numeric value if appropriate. This field accepts a value only if the data type is `char_string`, `hex_string`, `numeric`, `octet_string`, `snmp_char_string`, `snmp_hex_string`, or `snmp_octet_string`.
- If you want to indicate that the column is part of a unique combination of properties that define elements in the property table, click the box in the **Object-By** field so that the check mark (4) appears; in this case, the **Nullable** and **Default Value** fields should be blank. If you do not want this column to be an object by-variable,


click the box so that the check mark disappears. If you specify any other column other than `dsi_target_name` as an object by-variable, then you may not specify `dsi_table_key` as an object by-variable. However, you may specify `dsi_table_key` as an object by-variable only if you specify `dsi_target_name` as the first object by-variable.

- If you want to indicate that the column is part of a unique combination of columns that define elements to collect, click the box in the **Collection-By** field so that the check mark (4) appears; in this case, the **Nullable** and **Default Value** fields should be blank. If you do not want this column to be a collection by-variable, click the box so that the check mark disappears. If you specify any other column other than `dsi_target_name` as a collection by-variable, then you may not specify `dsi_table_key` as a collection by-variable. However, you may specify `dsi_table_key` as a collection by-variable only if you specify `dsi_target_name` as the first collection by-variable.
- If this column can contain NULL values, verify that the box in the **Nullable** field has a check mark (4) in it, which is the default. If this column should not contain NULL values, click the box so that the check mark disappears. If the **Object-By** or **Collection-By** fields are not blank, then this column should be blank.
- If the column should have a default value, type the value in the **Default Value** field that can be up to 255 characters. This field should be blank if this column is an object by-variable or a collection by-variable.
- If this column contains a foreign-key reference, select the property table that has the key from the pull-down list in the **Foreign Key** field. If this column should not contain a foreign-key reference, you can select the `no fk` value from the pull-down list that appears when you click in this field; the default is no foreign key.
- If you designate a column as a by-variable, you must put the columns in order of the by-variables such that the first column in the list is the first by-variable and the second column is the second by-variable, and so on.

If you want to change the order of the columns in the property table, select the column that you want to change and do the following:

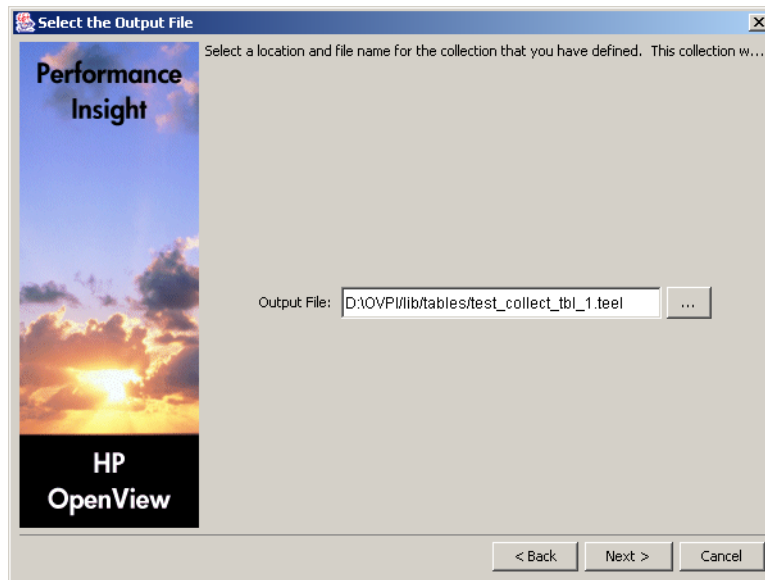
To move the column up, click the **Move Up** button, .

To move the column down, click the **Move Down** button, .

To delete a column from the property table, click the **Delete** button, .


Note that the property table has only six (6) default columns if the object by-variables and collection by-variables are the same.

- 12 Click **Next**. The **Select the Output File** dialog box opens.



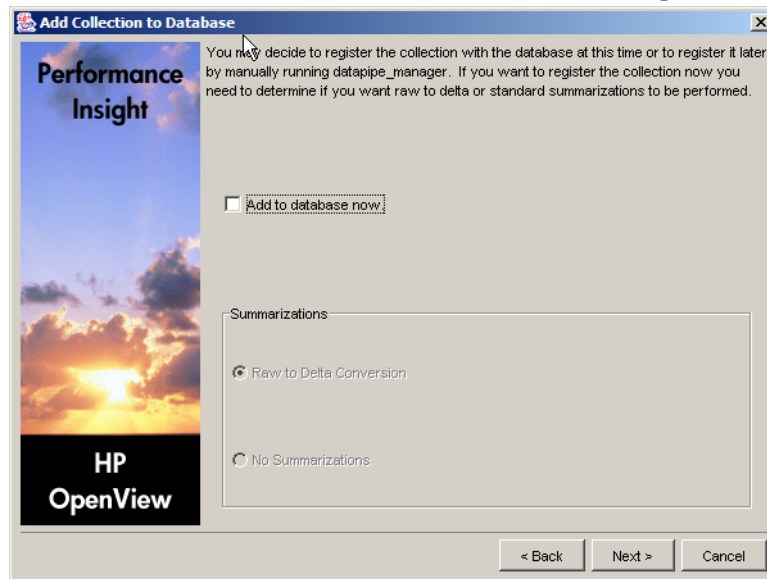
- 13 Verify that the name and location of the output file are correct.

If you need to change the name or location of the file, type the new name in the **Output File** field or do the following:

- a Click . The **Open** dialog box opens.
- b Select the directory from the **Look in** pull-down list.
- c Locate the appropriate directory.
- d Type the name in the **File name** field.

- e Click **Open**. The new name appears in the **Output File** field on the **Select the Output File** dialog box.

- 14 Click **Next**. The Add Collection to Database window opens.



- 15 If you want to generate the table now and add it to the database, verify that the **Add to database now** option is selected.

If you deselect this option, you will need to enter the following command from the command line:

```
datapipe_manager -p create -a teel_file
```

In this instance, *teel_file* is the name of the output file from [Step 13 on page 409](#).

If you need additional information about the options for this command refer to the *Performance Insight TEEL Reference Guide*.

- 16 Select the type of summarization that the system will perform on the data in the table from one of the following options:

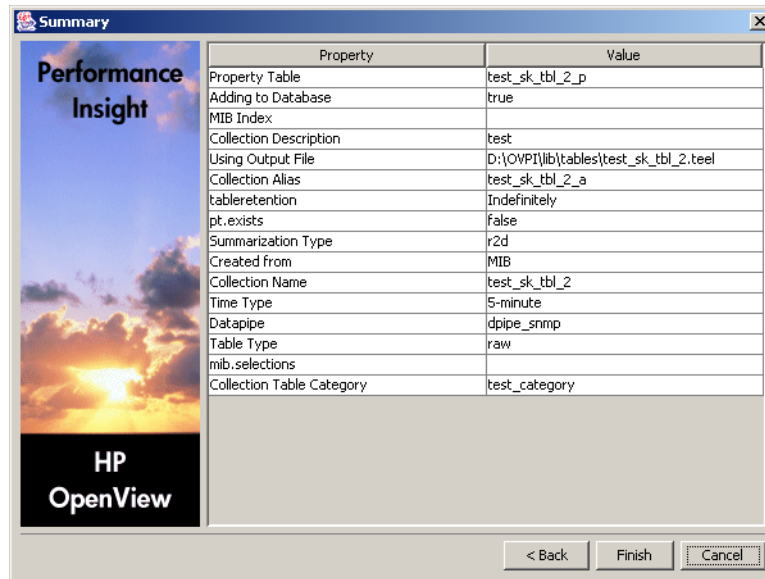
- **Raw to Delta Conversion**

The system will generate the required database objects for raw-to-delta processing. For manual processing, add `-r r2d` to the command in [Step 15](#).

— **No Summarizations**

The system will not generate any additional database objects.

- 17 Click **Next**. The **Summary** dialog box opens.



- 18 Do the following:

- a Review the summary.
- b If necessary, click **Back** to return to a previous page to make a change.
- c Click **Finish**. The TEEL file will be in the designated location.

Interpreting Log Files

HP OpenView Performance Insight's (OVPI) log files keep track of messages regarding the operation of the system. Messages appear in the log when processes begin and complete. Messages also appear when an error condition occurs.

This chapter lists the log files and describes the basic format for the messages that appear in these files.

The OVPI Log Files

OVPI generates the following log files:

- `audit.log`. Each message in this log lists an entered command and its start and end times.
- `builder.log`. Contains messages pertaining to the operation of Report Builder.
- `console.log`. Contains messages pertaining to the operation of the Management Console.
- `dbmaint.log`. Contains messages generated by the `dbmaint` utility, which synchronizes the internal tables to the new database schema.
- `install.log`. Contains messages for the OVPI installation or for an upgrade.

- `metrics.log`. Each message in this log lists an entered command and its results.
- `ovpi_uninstall.log`. Contains messages regarding an uninstall of OVPI.
- `report_YYMMDDhhmmss.log`. Contains messages for the installation of report packs using Package Manager. The file name specifies the year, month, day, hour, and time when you performed an installation, for example: `report_20030110131106.log`.
- `trend.log`. Contains messages for all of OVPI's back-end processes (for example, `trendtimer`). Each message specifies the start and end times for the logged process.
- `viewer.log`. Contains messages pertaining to the operation of Report Viewer.
- `website.log`. Contains messages pertaining to the operation of the Web Access Server. See [Viewing Log Files on page 312](#) for details about viewing the `website.log` file.

You can find all but one of OVPI's log files in the `log` folder of the OVPI installation directory. (For example, `C:\OVPI\log` for Windows systems.) The `ovpi_uninstall.log` file can be found in the following location depending on your operating system:

- UNIX
`/var/tmp/ovpi_uninstall.log`
- Windows
`%TEMP%\ovpi_uninstall.log`



All log files have the same message format except for the `audit.log` and `metrics.log` files. For more information about the audit and metric log files, see [The audit.log File on page 419](#) and [The metrics.log File on page 421](#), respectively.

Viewing Log Files

To view log files:

- 1 Navigate to the `log` folder in the OVPI installation directory using Windows Explorer (Windows) or a file manager (UNIX).

- 2 To open the log file you want to view, double-click it.

To open the log file this way, an appropriate program (such as WordPad on Windows systems) should be associated with the `.log` file type.

- ▶ On Windows, some log files cannot be open because they are exclusively held by the application using them. To open these log files, create a copy of the log file, and open the copy.

Log File Message Format

The following log files use the message format described in this section:

`builder.log`, `console.log`, `install.log`, `ovpi_uninstall.log`, `report_YYMMDDhhmmss.log`, `trend.log`, `viewer.log`, and `website.log`.

For information about the message formats for the audit and metric log files, see [The audit.log File on page 419](#) and [The metrics.log File on page 421](#), respectively.

All entries in the log files have the following comma-separated-value (CSV) format:

timestamp, time_zone, GMT_offset, appl_name, module_name, severity, process_ID, parent_proc_ID, message_ID, message

Each message contains fields separated by commas. If a field is missing, a comma should still appear as a placeholder.

timestamp

The *timestamp* field represents the time the log entry occurred. It uses a 24-hour clock with the following format:

YYYY-MM-DD hh:mm:ss:nnn

The format represents the following:

- *YYYY* is the 4-digit year.
- *MM* is the 2-digit month.
- *DD* is the 2-digit day.
- *hh* is the 2-digit hour.
- *mm* is the 2-digit minute.

- *ss* is the 2-digit second.
- *mmn* is the 3-digit millisecond. If this value is not known, it is 000.

For example, the following timestamp represents April 19, 2001 at 4:47 and 12 seconds and 123 milliseconds in the afternoon:

```
2001-04-19 16:47:12:123
```

time_zone

The *time_zone* field is the time zone where the log entry occurred.

For example, if the error occurred in Torrance during the winter, the time zone is Pacific Standard Time (PST). If the error occurred in Reston during the summer, the time zone is Eastern Daylight Time (EDT).

GMT_offset

The *GMT_offset* field is the offset of the timestamp's time zone from Greenwich Mean Time (GMT). It has the following format:

```
{±} hh:mm
```

The format represents the following:

- *hh* is the 2-digit hour.
- *mm* is the 2-digit minute.

For example, -05:00 is five hours before Greenwich Mean Time, which is either Eastern Standard Time (EST) or Central Daylight Time (CDT).

appl_name

The *appl_name* field is the unique name of the application in OVPI that caused the log entry.

module_name

The *module_name* field is a descriptive name that the programmer supplies.

severity

The *severity* field is the severity level for the log entry. The severity levels are the following:

- **INFO.** The message provides information regarding the application, and it does not require any action.
- **WARNING.** The message provides information that may require an action to avoid a more serious situation.
- **ERROR.** The message provides information about a problem that occurred while running the application. This situation requires an action to resolve the problem.
- **FATAL.** The message provides information about a problem that caused the application to fail and stop processing. This situation requires an action to resolve the problem.

If the message does not have a severity level identified, it uses the following value for the severity level:

- **DEF_ERROR.** The message provides information about a problem that occurred while running the application. This situation might require an action to resolve the problem.

process_ID

The *process_ID* field is the process identification number (Pid) for the process that caused the log entry. The system assigns a Pid to each process when it starts.

parent_proc_ID

The *parent_proc_ID* field is the process identification number (Pid) for the parent process of the process that caused the log entry.

message_ID

The *message_ID* field is a unique identification number for the message. If the message does not have an identification number assigned to it, this field is 0.

message

The *message* field is the description about the cause of the log entry.

Example Log File Messages

This section shows some examples of how errors may display in log files.

WARNING and INFO Messages

The following example shows WARNING and INFO messages in a log file:

```
2001-12-06 01:00:17:000,Pacific Standard Time,-8:00,
bcp_gateway,,WARNING,701,718,0,"no data collected for table
RIRRouterSystem, Db POWDER2"

2001-12-06 01:00:17:000,Pacific Standard Time,-8:00,trendtimer
,,INFO,187,41,0,"[Pid=544] C:\OVPI/bin/rmon_collect -n -i 60"

2001-12-06 01:00:18:000,Pacific Standard Time,-8:00,trendpm,,
INFO,475,718,0,"No procedures found to execute - trying to
generate one from command line params"
```

ERROR Messages

The following example shows ERROR messages in a log file:

```
2001-12-06 01:00:16:000,Pacific Standard Time,-8:00,dpipe_snmp
,,DEF_ERROR,726,718,0,"Error: failed to setup SNMP session with
host rough.hp.com,port 161,comm public (table RIRRouterSystem)"

2001-12-06 01:00:16:000,Pacific Standard Time,-8:00,dpipe_snmp
,,ERROR,660,718,0,"SNMP request Error[1]. host: 134.70.65.10,
comm: public, timeout: 1000, retries: 5, port 161"

.

.

.

2001-12-06 01:00:18:000,Pacific Standard Time,-8:00,trendpm,,
DEF_ERROR,475,718,0,"No output table specified"
```

The audit.log File

The `audit.log` file records the start times, end times, and durations of back-end processes. When a process begins, the file assigns a Pid (process identification number), and then it records when the process ends, showing that the Pid for the process was terminated.

Log File Format

All entries in the log file have the following comma-separated-value (CSV) format:

message type, program name, pid, ppid, start date-stamp, start time-stamp, end date-stamp, end time-stamp, duration in seconds, command line arguments

Message Type

The *message type* field is one of the following:

- `dtsys001`

This indicates the start of the message.

- `dtsys002`

This indicates the end of the message.

Thus, when the application starts, it logs a `dtsys001` message, and when it ends, it logs a `dtsys002` message.

Program Name

The *program name* field is the name of the executable.

PID

The *PID* field is the id of the process.

PPID

The *PPID* field is the id of the parent process.

Start Date Stamp

The *start date-stamp* field is the date when the application was started. It has the following format:

yyyymmdd

The format represents the following:

- *YYYY* is the 4-digit year.
- *MM* is the 2-digit month.
- *DD* is the 2-digit day.

Start Time Stamp

The *start time-stamp* field is the time when the application was started. It has the following format:

hh:mm:ss

The format represents the following:

- *hh* is the 2-digit hour.
- *mm* is the 2-digit minute.
- *ss* is the 2-digit second.

End Date Stamp

The *end date-stamp* field is the date when the application ended. This field is empty when the *message type* field is *dtsys001*.

It has the following format:

yyyymmdd

The format represents the following:

- *YYYY* is the 4-digit year.
- *MM* is the 2-digit month.
- *DD* is the 2-digit day.

End Time Stamp

The *end time-stamp* field is the time when the application ended. This field is empty when the *message type* field is *dtsys001*.

It has the following format:

hh:mm:ss

The format represents the following:

- *hh* is the 2-digit hour.
- *mm* is the 2-digit minute.
- *ss* is the 2-digit second.

Duration in Seconds

The *duration in seconds* field gives the duration of the execution of the process in seconds. This field is empty when the *message type* field is *dtsys001*.

Command Line Arguments

These are the command-line arguments used by the process when it runs. For example, the following line from the `audit.log` file shows that the `mw_collect` process is using the `-n` (distributed polling), `-i` (collection ID), and `-K` (reject samples) arguments when it runs:

```
dtsys001,mw_collect,888,1188,20031210,14:30:07,,,,"-n -i 5 -K 1"
```



In the example above, the following fields are blank, because the *message type* field is set to *dtsys001*: *end date-stamp*, *end time-stamp* field, *duration in seconds*.

Refer to the *Performance Insight Reference Guide* for information about the command-line arguments for any processes.

The metrics.log File

The `metrics.log` file records accounting information for the following OVPI programs: `mw_collect`, `bcp_gateway`, and `trendcopy`.

You can correlate the messages in the this log file to those in the `audit.log` file. In general, you can use the header portion of the message to find the correlating message. The *program name*, *PID*, and *PPID* fields will match exactly to the message in `audit.log` file; the *time-stamp* field in the `metrics.log` file message will be between the *start time-stamp* and *end time-stamp* fields from the message in the `audit.log` file.

Log File Format

All entries in the log file have the following comma-separated-value (CSV) format:

message type, program name, pid, ppid, date-stamp, time-stamp, number of objects to process, number of objects processed successfully, min, max, avg, object description



The meaning of *object* will be different for each program.

Message Type

The *message type* field is one of the following:

- dtsys003
This is the message summary.
- dtsys004
This is the message detail.

Program Name

The *program name* field is the name of the executable.

PID

The *PID* field is the id of the process.

PPID

The *PPID* field is id of the parent process.

Date-Stamp

The *date-stamp* field is the date when the message was generated. It has the following format:

yyyymmdd

The format represents the following:

- *YYYY* is the 4-digit year.
- *MM* is the 2-digit month.

- *DD* is the 2-digit day.

Time-Stamp

The time-stamp field is the time when the message was generated. It has the following format:

hh:mm:ss

The format represents the following:

- *hh* is the 2-digit hour.
- *mm* is the 2-digit minute.
- *ss* is the 2-digit second.

Message Data

The following fields are considered the message data fields: *number of objects to process*, *number of objects processed successfully*, *min*, *max*, *avg*, *object description*. These messages will differ depending on the program being used (*mw_collect*, *bcp_gateway*, or *trendcopy*). For a description of these messages, see [mw_collect Messages on page 423](#), [bcp_gateway Messages on page 424](#), or [trendcopy messages on page 425](#), respectively.

Object Description

The *object description* will differ depending on the message type as follows:

- *dtsys003*
The value for *object description* will be `all` if the executable was called without specifying a unique object to operate on. If unique object is specified, then the name of the object will be the object name itself.
- *dtsys004*
The value for *object description* will be the object itself.

mw_collect Messages

This section describes the message data that is displayed when *mw_collect* is the program being used. In this case, the *object* is a collection node. The *mw_collect* program will only produce *dtsys003* messages. Its message data is defined as follows:

- *number of objects to process*

This is the number of nodes `mw_collect` attempted to poll.

▶ This is not a count of unique nodes; for example, if the same node is collected for more than one table, this node will be counted multiple times.

- *number of objects processed successfully*

This is the number of successfully polled nodes.

A successful poll is defined as one where `mw_collect` was able to collect all the data without any error conditions. For example, if Node A has 5 interfaces and an SNMP timeout occurs after collecting only 3 rows, the poll is defined as not being successful.

- *min*

This is the minimum collection time, which is the minimum time the collector child ran (in seconds).

- *max*

This is the maximum collect time, which is the maximum time the collector child ran (in seconds).

- *avg*

This is the average collect time, which is the average time the collector child ran (in seconds).

- *object description*

This should be `all`, unless the `mw_collect` program ran with the `-h` option. For more information about the `mw_collect` command-line options, refer to the *Performance Insight Reference Guide*.

bcp_gateway Messages

This section describes the message data that is displayed when `bcp_gateway` is the program being used. In this case, the *object* is the rows in the data table. The `bcp_gateway` program will only produce `dtsys003` messages.

▶ Multiple records may be generated for the same table/database combination if multiple instances of the `bcp_gateway` are executed to load a table.

Its message data is defined as follows:

- *number of objects to process*

This is the number of rows that the program is trying to load. If the `bcp_gateway` command used the `-M` option, then this value will always be 0.

- *number of objects processed successfully*

This is the number of rows that loaded successfully.

- *min*

This is the minimum time to load a fixed length transaction (in seconds).

- *max*

This is the maximum time to load a fixed length transaction (in seconds).

- *avg*

This is the average time to load a fixed length transaction (in seconds).

- *object description*

This is the name of the object table name/database name.



If the total number of rows loaded by the `bcp_gateway` is less than transaction size (1000 rows) then the *min*, *max*, and *avg* times will reflect loading this number of rows. If the total number of rows loaded by the `bcp_gateway` is more than the transaction size, then the last set or rows may not be counted during the *min*, *max*, and *avg* calculations if the last set is less than transaction size.

trendcopy messages

This section describes the message data that is displayed when `trendcopy` is the program being used. In the case of `dtsys003` messages, the *object* is rows copied in every data and property table. The `trendcopy` program produces `dtsys003` and `dtsys004` messages.

The `dtsys004` message will be recorded only when `trendcopy` is started without specifying particular table name. In this case, the `dtsys004` messages will be generated for each property table/data table copy.

Its message data is defined as follows:

- *number of objects to process*

A “,” will appear in this field as a placeholder.

- *number of objects processed successfully*

This is the total number of rows copied; it includes data and property tables.

- *min*

This is the minimum time to load a fixed length transaction (in seconds); this includes key-table and data-table processing times.

- *max*

This is the maximum time to load a fixed length transaction (in seconds).

- *avg*

This is the average time to load a fixed length transaction (in seconds).

- *object description*

This is `all`, unless the table name is specified; in this case, the table name will be recorded as follows:

- If it is a data-only copy, then data table name is recorded.
- If it is a property-only copy, then property table name is recorded.
- If it is both, then data table and property table names are recorded.



If the total number of rows copied by the `trendcopy` is less than transaction size then the values of the minimum, maximum, and average times will reflect loading this number of rows. If the total number of rows loaded and copied by the `trendcopy` is more than the transaction size, then the last set or rows may not be counted during minimum, maximum, and average calculations if the last set is less than transaction size.

The `builder.log` and `viewer.log` Files

To locate the source of any report problems when using Report Builder or Report Viewer, you can view their respective log files: `builder.log` and `viewer.log`. When the Web Access Server is set to run in debug mode, the log files generate SQL statements when reports are loaded.

To run the Web Access Server in debug mode:

- 1 Choose **Options** from Report Builder or Report Viewer's Edit menu. The Options dialog box opens.
- 2 Select the **Logging** tab.
- 3 Click to select the **Enable debug tracing** option.
- 4 Click **OK**.

You can do the following to view log file messages from within the applications:

- Report Builder: Choose **Log File** from Report Builder's View menu.
The Log Viewer window opens, which displays the contents of the log file. The log file messages are categorized by type (all, info, warning, error, and debug).
To clear the messages in the window, right-click the Log Viewer window and choose **Clear** from the shortcut menu.
- Report Builder and Report Viewer: Choose **Log Output** from Report Builder or Report Viewer's View menu.
The Output Viewer window opens, which displays the log file messages as they are written to the log file.



Troubleshooting

This appendix describes the problems that you may encounter while using HP OpenView Performance Insight (OVPI) on Windows and UNIX systems.

OVPI Client Applications

This section describes problems that may occur when using OVPI's client applications.

Oracle Unable to Resolve a Service Name

Sometimes an OVPI application may display the following error:

```
ORA-12154: TNS: could not resolve service name
```

If you receive this error message, you must edit the `tnsnames.ora` file and add an entry that lets OVPI connect to Oracle.

To edit the `tnsnames.ora` file:

- 1 Find and open the `tnsnames.ora` file with a text editor; it is located in your `$ORACLE_HOME/network/admin` directory.
- 2 Add an entry at the end of the file that has the following format:

```
sid=  
(DESCRIPTION=  
(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)  
(HOST=host)(PORT=port)))  
(CONNECT_DATA=(SID=sid)))
```

where:

sid is the value of the Oracle SID.

host is the local system hostname.

port is the port number of the listener process, which is typically 1521.

For example, the following entry creates an SID called *ovpi* which specifies *squeaky* as the system where Oracle is installed and port 1521 as the port for the Oracle listener.

```
ovpi=  
(DESCRIPTION=  
(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)  
(HOST=squeaky)(PORT=1521)))  
(CONNECT_DATA=(SID=ovpi)))
```

- 3 Save your changes.
- 4 Stop and start the Oracle Listener program by doing one the following depending on your operating system:

Windows

- a From the Windows Desktop, click the Start button, point to Settings, and select **Control Panel**. The Control Panel window opens.
- b Double-click **Administrative Tools**, and then double-click **Services**. The Services window opens.
- c Right-click **Oracle TNSListener**, and select **Stop** from the shortcut menu. The Status column will indicate when the process is successfully stopped.
- d Right-click the **Oracle TNSListener**, and select **Start** from the shortcut menu.

UNIX

- a As the Oracle user, type the following at the command line:

```
Isnrctl stop
```

Wait for the message that indicates the command was successful.

- b** Type the following at the command line:

```
Isnrctl start
```

- c** Verify that there is a service started for the SID name of the OVPI database.

Displaying the Navigation Frame for the Web Access Server

If your Web browser has Java scripting turned off, you will be unable to use the navigation frame of the Web Access Server.

To use the navigation frame, do one of the following depending on your browser.

- Internet Explorer

To turn on Java scripting:

- a** From the browser's Tools menu, select **Internet Options**. The Internet Options dialog box opens.
- b** Select the **Security** tab.
- c** Click the **Internet** icon, and then click the **Custom Level** button. The Security Settings window opens.
- d** Scroll down to the bottom of the box to the Scripting options, and select **Enable** for Active scripting.
- e** Click **OK** in the Security Settings window.
- f** Click **OK** in the Internet Options dialog box.

- Netscape Navigator

To turn on Java scripting:

- a** From the browser's Edit menu, select **Preferences**. The Preferences dialog box opens.
- b** Scroll down to the Advanced options, and click **Scripts & Plugins**.
- c** Select the **Navigator** check box in the Enable JavaScript for box.
- d** Click **OK** in the Preferences dialog box.

Java Stack Trace Entries in the website.log File

If the `website.log` file contains java stack trace entries that have been caused by an `SSLPeerUnverifiedException`, it is likely that you have used the Web Access Server in a network which contains a proxy.

If the Web browser used to display the Web Access Server is set to use the HTTP/HTTPS proxy server rather than a direct internet connection, the proxy server may cause a connection reset during the initial communication between the Web Access Server and the Web browser, which causes this error.

You can ignore this error, because the Web browser and the Web Access Server will renegotiate and the SSL connection will be maintained.

Web Access Server Help Does Not Display Properly

Sometimes Internet Explorer does not display all the images that are on an HTML page. This usually occurs when you connect to the Web Access Server using SSL. Sometimes Internet Explorer will not display the + and - icons and the folder and file icons in the Web Access Server's navigation panel and instead shows missing image placeholders.

If this occurs, right-click the missing icon and select **Show Picture**. This will load the icon and make it visible.

Troubleshooting Sybase Space Allocation Problems

When using the Management Console's System/Network Administration application, you may see the following error (#1105) when Sybase is unable to allocate space for a database:

```
00: 94/07/13 04:07:08.18 server: Error: 1105, Severity: 17,
State: 3
00: 94/07/13 04:07:08.47 server: Can't allocate space for object
'syslogs' in database 'dpipe_db' because the 'logsegment' segment
is full. If you ran out of space in Syslogs, dump the transaction
log. Otherwise, use ALTER DATABASE or sp_extendsegment to
increase the size of the segment.
```


This message may be accompanied by the following message:

```
00: 94/07/13 04:07:08.75 server: Automatic checkpointing is
disabled in database 'dpipe_db' because the log is out of space.
It will continue when the DBO successfully checkpoints the
database. Please free up some space or extend the database and
then run CHECKPOINT.
```

If the State value of error message 1105 is 1, your database is out of space and needs to be enlarged or cleaned out.

To free up space in the log:

- 1 Determine approximately how much space the transaction log occupies by entering the following command:


```
sp_helpdb dpipe_db
```
- 2 Dump the inactive portion of the transaction log using the **dump transaction** command with the **truncate_only** option. If this command also fails with the error message, then use the **dump transaction** command with the **no_log** option.
- 3 Repeat [Step 1](#); however, if there are still a large number of pages in the **syslogs** table, an outstanding transaction is probably preventing the log from being cleared.

If so, check **syslogshold** for old transactions for the database specified in the error message (**dpipe_db** in the example error message) by entering the following commands:

```
use master
go

select * from syslogshold
where dbid = db_id('dpipe_db')
go
```

- 4 Determine whether the oldest active transaction can be terminated; it may have been left inactive intentionally. Continue this procedure until there are no other old transactions that can be terminated.

You can clear a long-running transaction using either of the following methods:

- If the long-running transaction is due to a runaway query and the process with the open transaction has been identified, use the **kill** command to stop the process. This clears the transaction and allows the log to be truncated. If the **kill** command cannot stop the process, restart the Sybase SQL Server.
- When you restart the Sybase SQL Server, it makes the database go through normal recover; thus, any outstanding transactions are either committed or rolled back. To restart the Sybase SQL Server, refer to the *Performance Insight Installation Guide*.

Rendering Graphics Using Hardware

After you install HP OpenView Performance Insight (OVPI), your system is configured for headless support. *Headless support* uses software to render graphics; thus, no graphics card is required. If you want to render your graphics using hardware, you can disable headless support after you install OVPI.

In order to disable headless support, do the following:

- Make sure your system has a graphics card.
- Disable headless support using *one* of the following methods:
 - **Web Access Server.** See [Changing the Web Access Server Settings on page 329](#).
 - **Command-line utility.** Type either of the following commands to disable or enable headless support, respectively:

```
headlessctl enable
```

```
headlessctl disable
```
- (*UNIX systems only*) If you are using the Web Access Server, you must configure your system for X *after* installing OVPI.

Configuring X

This section describes the procedure to configure X on a standard UNIX installation for Solaris and HP-UX platforms with a graphics card on the system. You must perform this procedure *after* OVPI is installed, and then start the Web Access Server.

All of supported UNIX operating systems provide `dtlogin`, and it is assumed that `dtlogin` is running. Non-standard systems might be running something other than `dtlogin`, such as `xdm`.

If you are going to install the Web Access Server on a UNIX system, you must ensure that your system meets the following requirements:

- X must always be running on your system.
- `root` must be able to access X always.

To configure X for a standard installation *after* OVPI is installed:

- 1 Login into the system as **root**.
- 2 Verify `dtlogin` is running by entering the following command:

```
ps -aef | grep dtlogin
```

If `dtlogin` is not running, you have one of the following:

- A non-standard installation of X on your system
- A headless system

- 3 If the `/etc/dt/config/Xconfig` file does not exist, copy `/usr/dt/config/Xconfig` to `/etc/dt/config/Xconfig`.
- 4 If the **`/etc/dt/config/Xsetup`** file does not exist, copy `/usr/dt/config/Xsetup` to `/etc/dt/config/Xsetup`.
- 5 Give execute permissions to `/etc/dt/config/Xsetup` by typing the following command at the UNIX prompt:


```
chmod 755 /etc/dt/config/Xsetup
```
- 6 Edit `/etc/dt/config/Xconfig` and uncomment the following line:


```
Dtlogin*grabServer: False
```
- 7 Append the following lines to `/etc/dt/config/Xsetup`:

```
# Give X access to root
XAUTH=path_to_xauth
LOCALHOST=`hostname`
KEY=`$XAUTH list | awk '{print $2 " " $3}'`
$XAUTH -f /.Xauthority add $LOCALHOST:0 $KEY
$XAUTH -f /.Xauthority add $LOCALHOST/unix:0 $KEY
```

In this instance, enter one of the following for *path_to_xauth*:

— *Solaris*:

```
/usr/openwin/bin/xauth
```

— *HP-UX*:

```
/usr/bin/X11/xauth
```

8 Open one of the following files depending on your operating system:

— *Solaris*:

```
/etc/init.d/ovpi_httpd
```

— *HP-UX*:

```
/sbin/init.d/ovpi_httpd
```

9 If the following line is not already present in the file, insert it directly after the `PATH=/bin:$PATH; export PATH` line:

```
XAUTHORITY=/.Xauthority; export XAUTHORITY
```

X is properly configured now; however, you *must* start the Web Access Server.

10 Start the Web Access Server by doing *one* of the following:

— Shutdown and restart the system.

— Do the following:

a Stop the OVPI HTTP Server.

b If a user other than `root` is logged on to the desktop, run the following command as `root`:

```
xauth merge ~username/.Xauthority
```

In this instance, *username* is the name of the user logged on to the desktop.

c Get the PID of the `dtlogin` process. If there is more than one `dtlogin` process, get the PID of the one with the parent PID of 1.

- d Type the following command:

```
kill -HUP PID
```

In this instance, *PID* is the PID of the `dtlogin` process in [Step c](#).

- e Start the OVPI HTTP Server.

Numerics

- 14K Pool parameter, **367**
- 16K Pool parameter, **367**
- 32-bit counter, **398**
- 64-bit counter, **398, 399, 404**
- 64-bit counter, higher precision values, **398**

Symbols

- .ini files, saving, **84**

A

- Actions menu
 - MIB Browser, **386**
- adding
 - nodes, **256**
 - property table descriptions, **232**
- additional network memory parameter (Sybase), **362**
- Administration link, **286, 287**
 - Java Settings folder, **290**
 - Log Files folder, **290**
 - Report Settings folder, **291**
 - Services Configuration folder, **288**
 - User Accounts folder, **287**
- administrator, resetting privileges, **333**

- aliases, changing, **324**
- audit.log file, **413**
- authentication
 - changing settings, **332**
 - setting method, **332**

B

- backup script
 - creating, **66**
 - running, **67**
- backup server, **53**
 - dumps and loads, **54**
 - shutting down, **64**
- backup server, creating
 - on UNIX systems, **54**
 - on Windows NT systems, **57**
- binding property tables, **374, 375**
 - verifying, **376**
- builder.log file, **413**

C

- cache
 - data pools, **367**
 - default data, **367**
 - named, creating, **369, 374**
- cascading windows, **82**

- Catalog link, **286, 293**
 - System folder, **293**
 - Users folder, **294**
 - categories
 - grouping by, **170**
 - change server
 - menu option, **80**
 - changing
 - database servers, **83**
 - table column size, **222**
 - checklist, post-relocation, **73**
 - collectable data table
 - definition, **259**
 - collection table
 - common index, **399**
 - definitions of, **397**
 - maximum columns, **260, 399**
 - maximum width, **260, 399**
 - required tables, **398**
 - column count, **260, 399, 409**
 - column size, changing, **118**
 - common index, **399**
 - community string profiles
 - creating, **145**
 - deleting, **146**
 - editing, **145**
 - managing, **144**
 - community strings
 - changing values for, **393**
 - connections, maximum number of, **365**
 - console.log file, **413**
 - constraints
 - creating, **252 to 254**
 - definition of, **252**
 - modifying, **255**
 - counter
 - 64-bit, **398**
 - CPU grace time parameter, **373**
 - create
 - polling groups
 - rule-based, **136**
 - creating
 - community string profiles, **145**
 - constraints, **252 to 254**
 - custom constraints, **253**
 - derived groups, **176**
 - enumerated list groups, **174**
 - folder groups, **180**
 - managed objects, **165**
 - nodes, **141**
 - polling groups, **131**
 - nodes in same view, **133**
 - nodes of same type, **132**
 - single node, **131**
 - specific instances, **136**
 - type and view, **134**
 - property groups, **178**
 - SNMP profiles, **147**
 - views, **318**
 - creating data tables, **260**
- ## D
- database
 - allocation errors, **432**
 - backing up with script, **64**
 - determining
 - available space, **220**
 - dump database command, **61**
 - dumping, **61 to 63**
 - examples of, **61 to 63**
 - forms, used to modify, **164**
 - high water mark, overriding, **220**
 - loading, **63**
 - restoring, **63**
 - retention days, **224**
 - sizing guidelines, **211**
 - stop collecting data, **220**

- tuning
 - See* performance
 - See also* Table Manager
- database information
 - viewing, **219**
- database servers
 - changing, **83**
- database statistics
 - viewing, **220**
- Database Tuner
 - configuring Sybase, **356 to 361**
 - See* performance
- Database Tuner (Sybase), **364**
 - additional network memory parameter, **362**
 - max network packet Size parameter, **362**
 - max online engines parameter, **363**
 - number of devices parameter, **363**
 - number of locks parameter, **364**
 - number of open indexes parameter, **364**
 - number of user connections parameter, **365**
 - procedure cache percent parameter, **365**
 - recovery interval in minutes parameter, **366**
 - SQL perfmon integration parameter, **362**
 - total memory parameter, **366**
- database usage
 - monitoring, **220**
- data cache pools, configuring, **367**
- data collection
 - defined, **18**
 - described, **18**
- datapipe installations
 - creating, **150**
 - deleting, **151**
 - editing, **150**
 - managing, **149**
- Datapipes
 - Property Import, **21**
- datapipes, **100**
- data processing
 - defined, **18**
 - described, **21**
- data reporting
 - defined, **18**
- dataset report, definition of, **23**
- data table
 - collectable, definition, **259**
- data tables
 - adding a description, **229, 232**
 - creating, **260**
 - defined, **22**
 - displaying data in, **235**
 - retention days, **224**
 - setting defaults, **278**
 - See also* property tables
- data tables, viewing
 - by category, **223**
 - by type, **223**
- dbmaint.log file, **413**
- debugging
 - enabling and disabling
 - Web Access Server, **331**
- deleting
 - community string profiles, **146**
 - datapipe installations, **151**
 - nodes, **143**
 - polling policies, **125**
 - SNMP profiles, **149**
 - views, **320**
- delta data, defined, **22**
- derived groups
 - creating, **176**
- devices, dumping

- to disk, **60**
- to tape, **60**
- digital certificates
 - CSR, generating, **335**
 - importing, **335**
- discovering, nodes
 - using SNMP Discovery wizard, **88**
- disk drives
 - effect on performance, **353, 354**
- displaying
 - online documentation, **85**
- display options
 - Group Manager, **187**
 - Table Manager, **280**
- documentation, locating on Web, **3**
- dpipes_db, **54**
- dumping
 - database, **61 to 63**
- dump transaction command, **433**

E

- editing
 - community string profiles, **145**
 - datapipe installations, **150**
 - groups, **184**
 - polling policies, **123**
 - SNMP profiles, **148**
- editing polling groups, **139**
- editing views, **320**
- Edit menu
 - Management Console, **80**
- ee_collect, **115**
- enumerated list groups
 - creating, **174**
- execution frequency, defined, **279**
- exporting

- groups, **182**
- nodes, **164**
- table data, **241**

F

- File menu
 - Management Console, **80**
- finding
 - groups, **173**
 - nodes, **141**
 - objects, **159**
 - table alias names, **222**
- flat files
 - defined, **18**
 - obtaining data from, **20**
- folder groups
 - creating, **180**
- Folder Uniqueness value, **160**
- foreign key, **233**
- forms
 - associated with managed objects, **157**
 - defined, **156**
 - using in Object Manager, **164**

G

- groupimport utility, **295**
- Group Manager, **167**
 - deleting groups, **185**
 - display options, **187**
 - grouping methods, **168**
 - group types, **167**
 - managing polling policies, **188**
 - starting, **170**
 - viewing logs, **186**
- groups
 - deleting, **185**

derived
 creating, **176**
 defined, **168**
 enumerated list
 creating, **174**
 defined, **168**
 exporting, **182**
 finding, **173**
 folder
 creating, **180**
 defined, **168**
 importing, **182**
 maintaining, **303**
 managed objects
 defined, **167**
 editing, **184**
 types, **167**
 using in Report Builder, **193**
 viewing contents of, **184**
 property
 creating, **178**
 defined, **168**
 user accounts
 creating, **304**
 defined, **38, 304**
 deleting, **308**
 filters, **40, 51, 307, 309, 310**
 modifying, **50, 308**

Groups folder, **288**

H

hardware effects on performance, **351**

headless support
 disabling, **331, 435**
 configuring X, **435 to 438**
 enabling, **331**

Help menu
 Management Console, **82**
 MIB Browser, **386**

High Part, defined, **399**

high water mark, overriding, **220**

hysteresis units
 defined, **279**
 setting, **279**

I

importing
 groups, **182**
 nodes, **126, 163**
 users and groups, **295**

install.log file, **413**

installing packages, **102**

interfaces file, **53**

internal tables
 See data tables

intervals, polling, **122**

Intimate Shared Memory (ISM), **353**

J

Java settings
 changing, **314**
 folder, **290**

K

keymap table, **232**

L

lag time
 See time lag

link bar
 changing aliases, **324**
 changing options, **322**

locking
 user accounts, **300**

log
 See transaction log

log files, **413 to 427**

audit.log, **413**

builder.log, **413**

changing settings, **311**

console.log, **413**

dbmaint.log, **413**

examples of, **418**

format of, **415 to 418**

install.log, **413**

location of, **413**

metrics.log, **414**

names of, **413**

ovpi_uninstall.log, **414**

report.log, **414**

trend.log, **48, 414**

viewer.log, **414**

viewing, **312, 414**

website.log, **414, 432**

Log Files folder, **290**

logging, enabling and disabling

Web Access Server, **331**

logging on

Management Console, **77**

log information

viewing, **186**

Low Part, defined, **399**

M

Mail Server page, described, **289**

managed objects

creating, **165**

defined, **167**

finding, **159**

forms associated with, **157**

modifying data in database, **164**

modifying Objects tree, **160**

properties, **155**

reports associated with, **157**

viewing, **156**

views, **158**

See also Group Manager

See also Object Manager

See also objects

Management Console

exiting, **80**

logging on, **77**

menu bar, **79 to 82**

Edit menu, **80**

File menu, **80**

Help menu, **82**

Tools menu, **81**

Window menu, **82**

starting, **76**

toolbar, **83**

managing

community string profiles, **144**

datapipe installations, **149**

nodes, **140**

SNMP profiles, **147**

master, **54**

max async I/Os per engine, **373**

max async I/Os per server, **373**

maximum columns, **260, 399**

max network packet Size parameter
(Sybase), **362**

max number network listeners, **373**

Max Online Engines parameter, **358**

max online engines parameter (Sybase), **363**

maxusers parameter, **355**

memory

allocation for database servers, **353**

guidelines, **352**

shared, **353**

menu bar

for Management Console, **79 to 82**

menu options, MIB Browser, **385**

- metrics.log file, **414**
 - MIB Browser, **379 to 411**
 - Actions menu, **386**
 - changing community strings, **393**
 - Help menu, **386**
 - interface, described, **382**
 - menu options, **385**
 - starting, **381**
 - MIBs
 - collection tables, **398**
 - described, **379**
 - loading, **386**
 - managed objects, **397**
 - object, inappropriate type, **387**
 - setting object values, **397**
 - unloading, **391**
 - modifying
 - constraints, **255**
 - property tables, **152**
 - monitoring
 - database usage, **220**
 - mw_collect, **115**
- N**
- named cache, creating, **369, 374**
 - navigation frame options, **325**
 - node
 - defined, **379**
 - nodes
 - adding, **256**
 - characteristics of, **140**
 - creating, **141**
 - custom groups, **122**
 - deleting, **143**
 - editing, **143**
 - exporting, **164**
 - finding, **141**
 - grouping, **169**
 - groupings, **122**
 - importing, **126, 163**
 - managing, **140**
 - number of devices parameter (Sybase), **363**
 - number of locks parameter (Sybase), **364**
 - number of open indexes parameter (Sybase), **364**
 - number of open objects parameter, **364**
 - number of open objects parameter (Sybase), **364**
 - number of user connections parameter (Sybase), **365**
- O**
- Object Manager, **155**
 - changing the view, **158**
 - starting, **156**
 - viewing logs, **160**
 - See also* managed objects
 - objects
 - See also* managed objects
 - online documentation
 - displaying, **85**
 - operating systems
 - effect on performance, **355**
 - ORA-12154
 - TNS
 - could not resolve service name, **429**
 - output
 - viewing
 - Object Manager, **160**
 - OVPI
 - functions of, **17 to 25**
 - ovpi_uninstall.log file, **414**
 - OVPI database
 - post-relocation checklist, **73**

relocating, **71**

P

Package Manager, **99**

starting

from command line, **102**

from Performance Insight, **101**

from Windows desktop, **101**

uninstalling packages, **111**

packages

dependencies for, **100**

installing, **102**

uninstalling, **111**

paging, caution about, **367**

password

setting expiration, **37, 300**

setting validation rules, **295 to 297**

Password Rules page, described, **288**

performance

changing multiple tuning parameters,
356

disk drive I/O, **353**

disk drive quantity, **354**

effect of retention time, **280**

factors affecting, **351**

hardware considerations, **351**

memory guidelines, **352**

for database servers, **353**

operating systems, **355**

parameter definitions, **361**

shared memory, **353**

swap file size, **353**

performance tuning, **351**

See performance

polling

frequency, **115**

intervals, **122**

purpose of, **116**

polling groups

creating, **131**

nodes in same view node, **133**

nodes of same type, **132**

rule-based, **136**

single node, **131**

specific instances, **136**

type and view, **134**

editing, **139**

managing, **129**

polling policies

creating, **120**

deleting, **125**

editing, **123**

for groups, **188**

creating, **188**

deleting, **192**

modifying, **191**

viewing, **119**

Polling Policy Manager

changing column size, **118**

importing nodes, **126**

intervals, **122**

refreshing policy list, **118**

sorting information, **118**

starting, **116**

viewing functions, **118**

port number

Web Access Server, **330**

procedure cache percent parameter (Sybase),
365

property groups

creating, **178**

Property Import Datapipe, **21**

Property Table

column count, **270, 409**

default columns, **270, 409**

maximum columns, **270, 409**

property table

- maximum columns, **260, 399**
- property tables
 - adding descriptions, **232**
 - binding to named cache, **374, 375**
 - creating, **275**
 - displaying details, **231**
 - keymap table, **232**
 - modifying values, **152**
 - verify binding, **376**
 - viewing, **225**
 - See also* data tables
- proxy servers, setting up, **74**

R

- RAID
 - effect on database tuning, **358**
- realm name, Web Access Server, **330**
- recovery interval in minutes parameter (Sybase), **366**
- refreshing
 - polling policy list, **118**
- report.log file, **414**
- Report Builder
 - creating constraints, **252 to 254**
- Report file, definition of, **23**
- report pack
 - provisioning, **21**
- report packs, **100**
- reports
 - activating, **327**
 - demo packages, **100**
 - deploying, **327**
 - managed objects, **157**
 - types of files, **23**
- report settings
 - changing, **315**
- Report Settings folder, described, **291**

- restoring
 - database, **63**
- results frame options, **326**
- retention days
 - data tables, **224**
- retention time
 - defined, **278**
 - for time-period tables, **279**
 - setting, **278**
- retention time, defined, **22**
- rmon_collect, **115**

S

- sar, **367**
- saving
 - .ini files, **84**
 - user preferences, **84**
- server
 - backup
 - dumps and loads, **54**
 - shutting down, **64**
 - starting, **58**
 - changing, **80**
 - disconnecting from, **80**
- services
 - changing settings, **329**
 - for UNIX
 - starting, **70**
 - for Windows
 - changing status of, **70**
 - checking status of, **69**
 - starting or stopping, **70**
- Services Configuration folder, **288**
- setting defaults for, **279**
- Settings page, described, **290**
- setting up, proxy servers, **74**
- shared memory, **353**

- Intimate Shared Memory feature, **353**
 - locking, **353**
 - Site Appearance folder, **321**
 - site options
 - changing, **321**
 - SNMP Discovery
 - function of, **87**
 - starting, **88**
 - SNMP profiles
 - creating, **147**
 - deleting, **149**
 - editing, **148**
 - managing, **147**
 - Solaris, **353**
 - sort-by settings for tables, **281**
 - sorting
 - database information, **221**
 - sorting information, **118**
 - sp_configure, **367**
 - SQL perfmon integration parameter (Sybase), **362**
 - SQL Server, backing up, **58, 64**
 - srvbuild utility, **54**
 - SSL
 - digital certificates, **333**
 - disabling, **334**
 - enabling, **334**
 - importing, digital certificates, **335**
 - keystore password, **334**
 - SSL Service page, described, **289**
 - starting
 - Management Console, **76**
 - Package Manager
 - command line for, **102**
 - from Performance Insight, **101**
 - from Windows desktop, **101**
 - storage time
 - setting, **228**
 - summarization, table, **273**
 - swap file size, **353**
 - swapping, caution about, **367**
 - Sybase
 - add buffer pools to default data cache, **367**
 - backing up SQL Server, **58**
 - backup server, **54**
 - script, **64**
 - shutting down, **64**
 - starting, **58**
 - default data cache, **367**
 - default transaction log, **371**
 - devices, dumping, **59**
 - error #1105, **432**
 - parameter definitions, **361**
 - RAM, **352**
 - space allocation problems, **432**
 - tuning
 - See performance
 - sybssystemprocs, **54**
 - System folder, **293**
 - System Manager, **337 to 345**
 - adding systems, **338**
 - changing settings, **333**
 - deleting systems, **345**
 - modifying system settings, **341**
 - System Manager page, described, **289**
- T**
- table, **232**
 - data, collectable, **259**
 - table alias names
 - finding, **222**
 - table data
 - exporting, **241**

- Table Manager, **217 to 281**
 - changing
 - table column size, **222**
 - creating tables, **259**
 - data tables
 - viewing by category, **223**
 - viewing by type, **223**
 - data types, **224**
 - determining
 - available space, **220**
 - display options, **280**
 - finding
 - table alias names, **222**
 - monitoring
 - database usage, **220**
 - navigation, **219**
 - panes, **218**
 - refreshing the display, **222**
 - setting
 - table viewing style, **281**
 - setting sort rules, **281**
 - sorting
 - table information, **221**
 - starting, **218**
 - storage time
 - setting, **228**
 - summarization types, **273**
 - table type setting, **281**
 - time types, **224**
 - viewing
 - database information, **219**
 - database statistics, **220**
 - property tables, **225**
 - table information, **221**
 - See also* data tables
 - See also* property tables
- tables
 - column size, changing, **222**
 - group member, **232**
 - keymap, **232**
 - property
 - binding, **374, 375**
 - verify binding, **376**
 - time-period, **279**
 - See also* data tables
 - See also* property tables
- tasks (forms), **157**
- tempdb database
 - sizing guidelines, **211**
- tiling windows, **82**
 - horizontally, **82**
 - vertically, **82**
- time lag
 - defined, **279**
 - setting, **278**
- time-out value, setting, **332**
- time-period tables, **279**
 - setting retention time, **279**
- title, changing, **321**
- tl_groups table, **73**
- toolbar
 - Management Console, **83**
- Tools menu
 - Management Console, **81**
- total memory parameter (Sybase), **366**
- traceflag parameters, **371**
- transaction log
 - determining size of, **433**
 - dumping inactive portion, **433**
 - I/O size, **370**
 - sizing guidelines, **211**
 - Sybase default, **371**
- Tree Nodes Per Level value, **160**
- trend.log file, **48, 414**
- trendtimer

- stopping
 - on Windows, **70**
- troubleshooting, **429 to 434**
 - could not resolve service name error, **429**
 - for the Web Access Server
 - cannot display navigation frame, **431**
 - Help does not display properly, **432**
 - java stack trace entry in website.log, **432**
 - space allocation problems, **432**
 - Sybase error #1105, **432**
- tuning
 - See performance
- tuning OVPI, **351**
- Type Discovery, **96**
 - function of, **87**

U

- uninstalling packages, **111**
- user accounts, **295**
 - creating, **297**
 - deleting, **302**
 - e-mailing, **302**
 - locking, **300**
 - modifying, **301**
 - password validation rules, **295 to 297**
 - sending e-mail to, **302**
 - setting password expiration, **37, 300**
- User Accounts folder, **287**
- user connections, maximum number of, **365**
- userimport utility, **295**
- user preferences
 - saving, **84**
- Users folder, **294**
- Users folder, described, **288**
- utilities
 - groupimport, **295**

- importing users and groups, **295**
- userimport, **295**

V

- viewer.log file, **414**
- viewing
 - database information, **219**
 - database statistics, **220**
 - data tables, **223**
 - log information, **186**
 - property tables, **225**
- viewing style, setting for tables, **281**
- views
 - creating, **318**
 - defined, **292, 318**
 - deleting, **320**
 - editing, **320**
 - node groups, **122**
- vmstat, **367**

W

- Web Access Server
 - Administration link, **286, 287**
 - Catalog link, **286, 293**
 - changing link bar aliases, **324**
 - changing link bar options, **322**
 - changing navigation frame options, **325**
 - changing results frame options, **326**
 - changing settings, **329**
 - changing title, **321**
 - configuring X, ?? to **438**
 - deploying reports, **327**
 - user accounts, **295**
- Web Access Server page, described, **289, 290**
- website.log file, **414, 432**
- Window menu
 - Management Console, **82**
- Windows

starting or stopping services, **70**

windows

 cascading, **82**

 horizontally tiling, **82**

 tiling, **82**

 vertical tiling, **82**

