

HP Performance Insight

for the Windows®, HP-UX, Solaris, and Linux operating systems

Software Version: 5.41

Administration Guide

Document Release Date: November 2012
Software Release Date: November 2010



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2009-2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Windows® is US registered trademark of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Oracle and Java are registered trademark of Oracle Corporation and/or its affiliates.

Adobe® is a trademark of Adobe Systems Incorporated.

Acknowledgements

This product includes Xerces XML Java Parser software, which is Copyright (c) 1999 The Apache Software Foundation. All rights reserved.

This product includes JDOM XML Java Parser software, which is Copyright (C) 2000-2003 Jason Hunter & Brett McLaughlin. All rights reserved.

This product includes JClass software, which is (c) Copyright 1997, KL GROUP INC. ALL RIGHTS RESERVED.

This product includes J2TablePrinter software, which is © Copyright 2001, Wildcrest Associates (<http://www.wildcrest.com>)

This product includes Xalan XSLT Processor software, which is Copyright (c) 1999 The Apache Software Foundation. All rights reserved.

This product includes EXPAT XML C Processor software, which is Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd and Clark Cooper Copyright (c) 2001, 2002 Expat maintainers.

This product includes Apache SOAP software, which is Copyright (c) 1999 The Apache Software Foundation. All rights reserved.

This product includes O'Reilley Servlet Package software, which is Copyright (C) 2001-2002 by Jason Hunter, jhunter_AT_servlets.com. All rights reserved.

This product includes HTTPClient Package software, which is Copyright (C) 1991, 1999 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

This product includes Perl software, which is Copyright 1989-2002, Larry Wall. All rights reserved.

This product includes Skin Look And Feel software, which is Copyright (c) 2000-2002 L2FProd.com. All rights reserved.

This product includes nanoXML software, which is Copyright (C) 2000 Marc De Scheemaeker, All Rights Reserved.

This product includes Sixlegs PNG software, which is Copyright (C) 1998, 1999, 2001 Chris Nokleberg

This product includes cURL & libcURL software, which is Copyright (c) 1996 - 2006, Daniel Stenberg, <daniel@haxx.se>. All rights reserved.

This product includes Quartz - Enterprise Job Scheduler software, which is Copyright 2004-2005 OpenSymphony

This product includes Free DCE software, which is (c) Copyright 1994 OPEN SOFTWARE FOUNDATION, INC., (c) Copyright 1994 HEWLETT-PACKARD COMPANY, (c) Copyright 1994 DIGITAL EQUIPMENT CORPORATION, Copyright (C) 1989, 1991 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

This product includes DCE Threads software, which is Copyright (C) 1995, 1996 Michael T. Peterson

This product includes Jboss software, which is Copyright 2006 Red Hat, Inc. All rights reserved.

This product includes org.apache.commons software developed by the Apache Software Foundation (<http://www.apache.org/>).

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport user ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

1	Introduction	23
	Users of Performance Insight	23
	Network Performance Monitoring	24
	Addressing Network Challenges Through PI	24
	Key Features of PI	25
	Functioning of PI	27
	Available Report Packs	29
	PI Documentation Set	32
2	Overview of HP Performance Insight	35
	Functions of PI	35
	Collect Data	35
	Determining the Types of Data to Collect	36
	Determining the Sources of Data to Collect	36
	Provisioning Reports	37
	Polling Parameters	38
	Process Data	38
	Archive Data	39
	Report Data	40
	PI Report Packs	41
	Creating Custom Reports	41
	PI Components	42
	PI Database Schema	42
	PI Performance Manager	42
	Web Access Server	43
	PI Clients	43
	PI Architecture	46
	Standalone Setup	46

Distributed PI Setup	46
Central Server	46
Satellite Server	47
Remote Poller	47
Web Access Server	47
Data Tables and Property Tables	48
Data Table	48
Property Table	48
Example	48
PI Administration	50
PI Management Console	50
Web Access Server	50
Command Line Utilities	50
3 Configuring HP Performance Insight	51
Additional Configuration Tasks	59
Configuring User Access	59
Creating User Accounts	59
Creating a User Group with a Filter	61
Managing Database Changes	64
Collecting Data from SNMPv3 Nodes	64
Configuring SNMPv3 Nodes	65
Using the brass_v3_seed Utility	65
Importing SNMPv3 Nodes	66
Creating Community String Profiles	67
Assigning Nodes to a Type Group	68
Configuring Remote Pollers for SNMPv3 Nodes	69
Configuring a Distributed Installation	69
Configuring Polling Policies for Remote Pollers	69
Changing the Polling System	70
Creating New Policies for Multiple Pollers	71
Distributing Polling Between Multiple Pollers	74
Option 1: Assign a Subset of the Polling Policies to Each Poller	74
Option 2: Assign a Subset of the Nodes to Each Poller	74
4 Maintaining HP Performance Insight Installation	81

Post-Installation Administration Tasks	81
Verifying that HP Performance Insight is Running	81
Scheduling Reports	83
Modifying User Accounts	83
Modifying a User Group	84
Modifying a User Group Filter	85
Deleting a User Group Filter	85
Daily Administration Tasks	86
Checking the Status of Data	86
Weekly Administration Tasks	87
Backing Up Your Oracle Database	87
Backing Up Your Sybase Database	87
Backing Up a Sybase SQL Installation Server	88
Backing Up the Sybase SQL Server	92
Backing Up the Sybase SQL Server Using a Script	98
Infrequent Administration Tasks	102
Backing Up Your System	102
Starting and Stopping HP Performance Insight Processes	103
Starting and Stopping Processes on Windows	103
Starting a Service	104
Starting and Stopping Processes on UNIX	105
Relocating the HP Performance Insight Sybase Database	106
Verifying the Relocation of the HP Performance Insight Database	107
Setting-up HP Performance Insight in an HTTP Proxy or Portal Environment	108
Example: Setting Up an Apache Proxy Server	112
5 Using the PI Management Console	115
Starting the Management Console	118
UNIX System	118
Windows System	118
Logging On	119
The Management Console Interface	120
Menu Bar	121
File Menu	121
Edit Menu	121
Tools Menu	122

Help Menu	123
Toolbar	124
The Home Snapin.	125
Viewing HP Performance Insight Status.	125
Viewing HP Performance Insight Release Notes	127
Viewing HP Performance Insight General Information	127
Changing Database Servers	128
Saving User Preferences	129
Displaying Online Documentation	131
6 Installing Report Packs	133
Report Packs.	134
Datapipes	135
Report Pack and Datapipe Dependencies.	136
Starting Package Manager	137
Windows Desktop	137
HP Performance Insight Management Console (UNIX and Windows)	137
Command Line (UNIX and Windows).	138
Installing Packages	139
Uninstalling Packages	148
7 Discovering the Network	151
The Discovery Process	152
Using the SNMP Discovery Wizard.	152
Using the Type Discovery Process	159
The Type Discovery Wizard	159
8 Managing Polling Policies	163
Polling Strategies	164
Polling Groups.	164
Polling Frequency	165
Directed Instance Polling	166
Starting Polling Policy Manager	168
Viewing Functions	169
Sorting Columns	169
Changing the Size of a Column.	170

Refreshing a Policy List.	170
Selecting Items	170
Viewing Polling Policies	171
Viewing the Current Polling Policies.	171
Creating Collections.	173
Creating a Collection from a MIB	173
Creating a Collection from the HP Performance Agent or the HP Operations Agent	197
Communicating in a Firewall Environment	197
Creating a Collection from the HP Performance Agent or the HP Operations Agent	200
Collecting Data from Agents with HTTPS Access.	222
Collecting Data from Agents through Proxy	229
Creating a Collection from a CSV File	230
Creating a User Defined Collection	241
Creating Polling Policies	252
Example: Creating a Polling Policy for an HP Performance Agent Collection	255
Verifying that the Policy is Polling	256
Editing Polling Policies	257
Deleting a Polling Policy	259
Managing Polling Groups	260
Creating a Polling Group	261
Creating a Single Node Group	262
Creating All Nodes of the Same Type Group	263
Creating All Nodes in Same View Group	264
Create a Combination of Type and View Group	265
Creating a Specific Instance Group	267
Editing a Polling Group	270
Managing Nodes.	271
Finding a Node	272
Creating Nodes	272
Editing Nodes	274
Deleting Nodes	274
Importing Nodes	275
Creating the Import File	275
Importing the File	278

Managing Community String Profiles	279
Creating a Community String Profile	280
Editing a Community String Profile	280
Deleting a Community String Profile	281
Managing SNMP Profiles	282
Creating an SNMP Profile	283
Editing an SNMP Profile	284
Deleting an SNMP Profile	284
Managing Datapipe Installations	285
Creating a Datapipe Installation	286
Editing a Datapipe Installation	287
Deleting a Datapipe Installation	287
Modifying Property Table Values	288
9 Internal Monitoring of PI using Performance Graphs	291
Starting HP Performance Insight Performance Graphs	292
Viewing HP Performance Insight Performance Graphs	293
Process Duration	294
Create Process Duration Graphs	294
Polling Health	296
Create Polling Health Graphs	296
Row Counts	298
Create Row Counts Graphs	298
10 Managing Objects	301
Understanding the Object Model	302
Locating Managed Objects	304
Changing the Objects View	306
Finding a Specific Object	308
Viewing Log Information	308
Modifying the Objects Tree Display	309
Adding New Object Models	312
Importing Nodes	316
Creating the Import File	316
Importing the File	316
Exporting Nodes	317

Using Forms	318
Modifying the Database	318
Creating New Managed Objects	319
Deleting Nodes	321
To edit the Exclusion List of nodes	328
11 Grouping Managed Objects	329
Grouping Types	330
Grouping Methods	331
Grouping Nodes	331
Grouping Other Managed Objects	332
Group Categories	332
Viewing Managed Object Groups	333
Changing the Groups View	334
Finding a Specific Group	336
Creating Groups	337
Creating an Enumerated List Group	337
Creating a Derived Group	339
Creating a Property Group	342
Creating a Folder Group	344
Importing Groups	347
Exporting Groups	348
Managing Groups	349
Viewing a Group	349
Editing a Group	350
Deleting a Group	352
Viewing Log Information	352
Setting the Display Options	353
Managing Group Polling Policies	356
Creating a Group Polling Policy	356
Modifying a Group Polling Policy	358
Deleting a Polling Policy	359
Changing the Polling Policy View	359
Using Groups in Reports	361
12 Managing Your System	363

Starting the System and Network Administration Application	364
System/Network Administration Window.....	365
Managing Your Web Access Server Systems	366
Adding a Web Access Server	366
Viewing and Modifying Web Access Server Properties.....	369
Synchronizing Web Access Servers	372
Changing Your Web Access Server	373
Deleting a Web Access Server	374
Managing Your Databases.....	375
Adding a Database	375
Viewing Database Properties.....	377
Changing Database Initialization Parameters	380
Increasing Database Size.....	381
Database Component Sizing Guidelines	381
Increasing the Size of the Database	382
Changing the Database Server	385
Add a Remote Poller.....	386
13 Managing Tables.....	391
Starting Table Manager.....	392
Viewing Table Information	394
Viewing Functions	394
Sorting Information.....	394
Changing Column Size	394
Refreshing Information.....	394
Finding Alias Names.....	395
Viewing Data Tables by Type or Category	395
Viewing Property Tables	398
Changing the Table Column Display.....	399
Viewing Data and Internal Table Details.....	400
Viewing Data Table Column Information	401
Adding a Data Table Description.....	401
Viewing Data Table Column Information	401
Viewing Data Table Index Information.....	402
Retention Profiles	403
Index Profiles	403

Managing Index Profiles	404
Viewing Property Table Details	409
Viewing Property Table Information	410
Adding a Property Table Description	410
Viewing Property Table Column Information	411
Viewing Property Table Index Information	412
Viewing or Managing an Index Profile	413
Displaying Table Data	414
Displaying Reports	414
Using the Menus	415
Managing Data	417
Printing Table Data	417
Saving Table Data	418
Exporting Table Data	419
Formatting Data	420
Formatting Statistics	420
Formatting Cells	422
Formatting Columns	423
Formatting Rows	426
Changing Table Queries	426
Selecting Statistics	427
Sorting the Data in the Report	427
Changing the Time Period	428
Adding Constraints	432
Changing Constraint Values	435
Modifying Constraints	435
Deleting Constraints and Clauses	436
Selecting Nodes and Interfaces	436
Changing the Maximum Number of Rows Returned	438
Creating Tables	440
Creating Data Tables	440
Creating Property Tables	453
Setting Table Defaults	457
Setting Data Table Defaults	457
Time-Period Table Default	458

14 Using Catalog Manager.....	461
Starting Catalog Manager.....	463
Catalog/Report Link Management Window	464
Managing the Deployed Items Catalog.....	465
Deploying Reports.....	465
Undeploying Reports	469
Undeploying a Report Using the Delete Menu Option	470
Undeploying a Report Using the Report Deployment Wizard	470
Deploying URLs	472
Creating New Directories in the Catalog	473
Moving Reports and Directories.....	473
Copying Reports and Directories	475
Viewing and Changing Directory Properties.....	476
Viewing and Changing Deployed Item Properties	477
Viewing and Changing Report Properties	477
Viewing and Changing URL Properties.....	478
Deleting Reports and URLs	479
Refreshing the Catalog.....	479
Managing the Link Catalog.....	480
Creating Link Groups.....	480
Creating Report Links	482
Linking a Report to another Report	482
Linking a Report to an Application	491
Linking a Report to a URL	496
Opening Report Links	504
Modifying Links in the Viewing Pane	505
Copying Links	506
Moving Links	507
Deleting Links.....	508
Viewing and Changing a Link's Properties.....	508
Exporting and Importing Links	509
Exporting a Link Group	510
Importing a Link	510
Refreshing the Links Catalog.....	510
15 Using the Web Access Server.....	511

Accessing the Web Access Server	512
Web Access Server Overview	514
Administration Link	515
User Accounts	515
Service Configuration	516
Diagnostics	518
Log Files	518
Java Settings	519
Deployed Items	520
Catalog Link	522
Maintaining User Accounts	525
Setting Password Validation Rules	525
Creating a User Account	527
Modifying a User Account	531
Deleting a User Account	531
Sending Email to a User Account	532
Maintaining User Groups	533
Understanding How User Groups Work	533
Creating a User Group	534
Modifying a User Group	537
Deleting a User Group	537
Creating a User Group Filter	538
Modifying a User Group Filter	539
Deleting a User Group Filter	539
Example: Data Filtering Using Groups	540
Displaying Diagnostics	543
Changing Log File Settings	544
Viewing Log Files	546
Changing Java Settings	548
Changing Deployed Item Settings	549
Creating Views	552
Creating a View	552
Modifying a View	554
Changing the Appearance of the Web Access Server	556
Changing Site Options	556

Changing Link Bar Options	556
Changing Link Bar Aliases	559
Changing Navigation Frame Options	561
Changing Results Frame Options	562
Deploying Reports	563
Changing Settings for the Services	565
Changing the Web Access Server Settings	565
Resetting Administrator Privileges	568
Changing the SSL Service Settings	570
Importing Your Digital Certificate	573
Generating a Certificate Signing Request (CSR)	573
Importing a Signed Digital Certificate	574
Changing System Manager Settings	575
Displaying the Performance Insight Systems Page	575
Choosing the Default System	576
Adding a System	576
Modifying a System's Settings	579
Deleting a System	583
Changing Mail Server Settings	583
Modifying Graph X-Axis Time Formats	585
16 Archiving Data	589
Archive Tables	590
Data Retention Profiles	591
Archiving Data Process	592
Managing Retention Profiles	593
View Retention Profiles	593
Add Retention Profiles	594
Modify Retention Profiles	595
Remove Retention Profiles	595
Viewing Data and Internal Table Details	597
Viewing Data Table Column Information	598
Setting Data Table Retention Profile	598
Assigning a Different Retention Profile to a Data Table	598
Assigning a Retention Profile to Multiple Data Tables	600
Adding a Data Table Description	600

Viewing Data Table Column Information	601
Viewing Data Table Index Information	602
Viewing Data Tables by Type or Category	603
Guidelines to set up Archiving	605
Limitations of Archiving	610
17 Location Independent Reporting and Copy Policies	611
Location Independent Reporting (LIR)	611
LIR Mapping	612
LIR Implementation Example	613
Install Wizard Updates for LIR Support	613
Manage LIR Mappings	617
Enable LIR	618
Disable LIR	618
Add a New LIR Mapping	619
View LIR Mappings	620
Delete LIR Mappings	620
Modify LIR Mapping	620
Copy Policies	622
Copy Policy Implementation Examples	622
Start Copy Policy Manager	623
View Configured Copy Policies	624
Manage Copy Policies	625
Create a Copy Policy	625
Edit Copy Policy	628
Delete Copy Policy	629
18 Managing MIBs	631
Starting the MIB Browser	633
The MIB Browser Window	635
Navigation Pane	636
Viewing Pane	636
Menu Bar and Toolbar	638
File Menu	638
View Menu	639
Actions Menu	639

Help Menu	639
Loading MIBs.....	640
Issues to Consider.....	640
Including the Top of the MIB Tree	640
Validating Object Types	641
Loading a MIB	642
Unloading a MIB	646
Querying SNMP Data	648
Setting the Node	648
Setting the Community String.....	648
Setting the SNMP Version	649
Finding a Group or Object in the MIB Map	649
Viewing the Data	650
Setting SNMP Values	653
Creating Collection Tables.....	654
Collection Table Features.....	654
64-Bit Values in SNMPv1.....	654
Collection Table Requirements	656
Creating a Collection Table	656
19 Tuning the System	673
Hardware Considerations	674
Processor Type and Clock Speed	674
Number of Processors	674
Memory	675
Shared Memory / Swap File Size	675
Disk Drive Interfaces	676
Number of Physical Disk Devices	677
Configuring Operating System Parameters	677
Configuring Oracle Parameters	678
Configuring Sybase Parameters	678
Using the Database Tuner	678
About Sybase Parameters	684
Using the sp_configure Command.....	690
Operating System - Asynchronous I/O Processing	691
Transaction Log Cache	691

Transaction Log I/O Size	693
Traceflag Parameters	694
Other Configuration Parameters	696
Binding the Property Tables	697
20 Interpreting Log Files	701
HP Performance Insight Log Files	702
Viewing Log Files	703
Log File Message Format	703
timestamp	704
time_zone	704
GMT_offset	704
appl_name	705
module_name	705
severity	705
process_ID	705
parent_proc_ID	706
message_ID	706
message	706
Example Log File Messages	706
WARNING and INFO Messages	706
ERROR Messages	706
The audit.log File	707
Log File Format	707
The metrics.log File	710
Log File Format	710
mw_collect Messages	712
bcp_gateway Messages	713
trendcopy messages	714
The builder.log and viewer.log Files	715
A NNM-PI Integration	717
Features and Benefits	717
Index	719

1 Introduction

This chapter provides an introduction to HP Performance Insight (PI) and describes the challenges you might face when monitoring the performance of your network and how PI can help you address those challenges.

Before reading this manual, you are expected to have a good working knowledge of network management protocols such as the Simple Network Management Protocol (SNMP). You must also understand the various types of network equipment and their functional differences.

For information on installing Performance Insight on Oracle or Sybase database, see the *HP Performance Insight Installation and Upgrade Guide for Oracle Databases* or *HP Performance Insight Installation and Upgrade Guide for Sybase Databases* respectively.

Users of Performance Insight

Performance Insight is designed for network managers, technicians, and service managers who are responsible for monitoring and managing IT networks or by users who want to track the performance of their networks. PI helps the following users in general:

- Business executives: Use PI for high-level overviews
- Service managers: Use PI to generate Service Level Reports
- Network managers: Use PI's proactive tools to identify potential problems before service is impacted
- Capacity planners: Use PI to identify over-utilized or under-utilized resources

Network Performance Monitoring

When managing your IT network operations, you might face the following challenges continually:

- Tracking network performance
- Monitoring the speed and inflow of network traffic
- Optimizing capacity planning
- Providing comprehensive and definitive performance reports to customers and service management teams
- Meeting service level agreements?
- Monitoring availability and performance of your services and the networks, systems, and applications that make up those services?
- Tracking the components affected by non-performing services.
- Predicting future network related problems and potential failures and outages

Performance Insight and its reporting solutions (commonly known as the Report Packs) help you address these challenges.

Addressing Network Challenges Through PI

Performance Insight is a highly scalable and customizable reporting tool. PI performs the following functions:

- Collects management data from networks, systems, applications, and services.
- Analyzes and summarizes the data and generates reports that enable you to monitor how the services are performing.
- Discovers network problems before they impact your network, applications, or systems.

In addition, you can use PI to perform the following functions:

- Diagnosing problems and capacity planning
- Identifying over-utilized and under-utilized elements in your environment
- Tracking system and network performance and their effects on the performance of your services.
- Documenting current network performance for internal use and meeting customer service level agreements

Key Features of PI

The following are the key features and resultant benefits of PI:

Key features	Benefits
Forecast/capacity planning reports using Available Report Packs	<ul style="list-style-type: none">• Identify potential problems before they impact customers.• Reduce expenses by load balancing on under-utilized resources
Customizable reports	<ul style="list-style-type: none">• Arrange or change tables and graphs to meet your staff's needs• Change labels and provision new objects to meet all audience needs• Create new reports on the data that is most relevant to you and your customers• Put your logo or your customer's logo on reports
Customer-suitable web-based reporting	<ul style="list-style-type: none">• Allow your customers to view the performance of systems, networks, and services• Generate revenue by selling access to service monitoring

Key features	Benefits
User-settable thresholds	Report on exceptions related to your operation
Service level monitoring and reporting	Track SLA parameters before they impact users or invoke penalties
Integration with HP Network Node Manager (NNM)	By integrating fault management with performance management, the NNM-PI Integration Module enhances problem diagnostic capabilities
Powerful management console	Use a single user interface to administer and control all aspects of your performance management solution
SSL-level security	Get performance management that meets the highest standards for security
PI deployment in a distributed environment	Deploy PI across geographies and optimize the management of network load

Functioning of PI

PI discovers network devices and systems in your network and polls data from them. PI can poll data from SNMP devices, other HP Software products, and non-SNMP devices.

You can specify the type of data to be polled and also the frequency for polling using specific reporting templates called report packs. The report packs contain the polling policies that define what data to collect and when to collect. PI includes many out-of-the-box Report Packs that you can use to generate reports based on your requirements.

After the data is collected, PI assigns timestamps to the records and stores the data in data tables. PI then converts this raw data into delta (or rate) data. Delta data is the difference between two values over a period of time. Depending upon the instructions provided in the report pack, the delta data is aggregated into hourly, weekly, monthly, quarterly, and yearly summaries, 30-day, 60-day, and 90-day forecasts; percentiles; and averages.


You can use the processed data to generate reports. You can generate reports automatically using out-of-the-box Report Packs included with PI. The report packs available with PI help you to create, deploy, schedule, and generate reports.

Alternatively, you can create your own custom Report Packs using Report Builder which is a client application of PI. Through Report Builder you can create report templates that define how the data will appear in the report, the data that needs to be collected from the database, and how the data will be linked to the report.

You can view the reports through a web-based interface (Web Access Server) or through the client applications Report Builder and Viewer.

The report packs also highlight the most urgent issues, provides historical trends and forecasts potential problems.

The following image illustrates the main functions of Performance Insight



Available Report Packs

The following is a list of out-of-the-box PI Report Packs, along with their features:

Report Pack	Features
Report Pack for System Resources	Reports on system performance for file systems, memory, CPU, and grade of service
Report Pack for Infrastructure Usage	<ul style="list-style-type: none">• Provides comprehensive insight into how network devices, LAN and WAN interfaces, network segments and VLANs are performing• Provides the complete view of how well the most critical components of the network fabric are functioning
Report Pack for LAN/WAN Edge	<ul style="list-style-type: none">• Enables enterprise and service provider network operators to view usage patterns, trends, error rates and capacity of network edge routers, including Frame Relay and ATM interfaces
Report Pack for Traffic Profile	<ul style="list-style-type: none">• Provides network visibility, so that you can better understand how the network is being utilized by your end users• Includes RMON2 and NetFlow data sources

Report Pack	Features
Report Pack for Quality Assurance	<ul style="list-style-type: none"> • Provides insight into the quality of the services delivered to your IT customers • Measures response times and throughput between IP devices • Includes analysis of Cisco SAA, Cisco Ping, Cisco CAR MIB, and Cisco IP STAT MIB
Report Pack for MPLS VPN	Monitors, reports and alarms on MPLS-based networks delivering virtual private network services
Report Pack for IP Telephony	<ul style="list-style-type: none"> • Enables the user to move to IP telephony • Provides insight into number of gateways required in a voice-quality IP telephony architecture • Offers analysis of Cisco Call Manager and Cisco Gateway servers

Report Pack	Features
Report Pack for WAN Core	Provides a performance reporting solution for service providers managing a Frame Relay or ATM backbone
Report Pack for Network Node Manager	<ul style="list-style-type: none"> • Highlights historical trends of the most common and critical events from HP Network Node Manager • Provides insight into availability of network devices and their Mean Time To Repair (MTTR) and Mean Time Between Failures (MTBF)
Report Pack for Internet Services	<ul style="list-style-type: none"> • Provides a reporting solution for services monitored by HP Internet Services • Details availability, service level violations, response time, and throughput

PI Documentation Set

The following is a list of manuals in the PI document set (apart from this manual). These documents are in the PDF format and you can access them by clicking **Help** in the PI user interface:

- Installation and Upgrade Guide for Oracle Databases on HP-UX, Solaris, and, Linux operating systems.
- Installation and Upgrade Guide for Oracle Databases on Windows operating system.
- Installation and Upgrade Guide for Sybase Databases on HP-UX, Solaris, and, Linux operating systems.
- Installation and Upgrade Guide for Sybase Databases on Windows operating system.
- Guide to Building and Viewing Reports.
- Reference Guide.
- Network Node Manager-Performance Insight Integration User Guide.


You need Adobe Acrobat Reader to read these guides. If you do not have Acrobat Reader, you can install it from the `docs` folder, located in the PI install directory.

To view the online documentation:

- 1 From the Help menu, click **Help Topics**.
 - If you have not selected 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 selected a browser for the client application you are using, a browser window opens, displaying the list of online documents. Go to [step 4](#).
- 2 In the Browser tab, do one of the following:
 - Type the Web browser's complete directory path and application file name in the Browser Application box.


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

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


- Click . The Open dialog box opens. Locate and double-click the Web browser's application file name.

3 Click **OK**.

A browser window opens, displaying 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 guide opens in Acrobat Reader.

2 Overview of HP Performance Insight

This chapter describes the different components of HP Performance Insight (PI) and how they work. It also gives you an understanding about the environments in which you can use PI.

Functions of PI

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

- Collect Data
- Process Data
- Archive Data
- Report Data

Collect Data

This function of PI 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 PI database.

PI determines the type of data to collect, where to collect it from, and when to collect it. PI can collect data from the following sources:

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

In this context, SNMP-type devices include Remote Monitoring (RMON) devices.

Determining the Types of Data to Collect

PI uses a list of object identifiers (OIDs) from SNMP-type devices to determine what data to collect. The OIDs are in the Management Information Base (MIB) tables. PI automatically provides these tables when you install a report pack.

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

When PI collects data from a non-SNMP device, it uses TEEL files to import and store the data from the device. PI automatically provides the definition file for the corresponding tables and data file when you install a report pack. Alternatively, if you want to create your own definition, 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:** Enables you to use a wizard to ping a range of network addresses to locate the nodes that respond. See [The Discovery Process](#) on page 152 for details.
- **Polling Policy Manager:** Enables 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 8, Managing Polling Policies](#) for details.
- **node_manager:** Enables you to import an existing list of nodes into the system. You can set up an ASCII file with the required information and then import the nodes using this utility. For more information about the `node_manager` command, refer to the *HP Performance Insight Reference Guide*.

After specifying the nodes use the Type Discovery process to identify and group the nodes. Typically node groups are based on the specifications in the Report Packs. However, if you want to poll only a specific set of nodes, you can create your own groups, which are referred to as view lists. See [Creating a Polling Group](#) on page 261 for more information about creating a node group. When a poller polls the nodes under these lists, it places the data that it gathers into data tables.

To collect data through SNMP enabled HP performance agents, use SNMP Discovery to discover the nodes. To collect data through non-SNMP enabled HP Operations performance agents, use the `node_manager` command to import the list of nodes.

Data for non-SNMP devices is collected from a flat file. To obtain the data from the flat files set up a TEEL file that provides instructions for importing the data and define the tables that store data from the 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 to collect the data properly.

SNMP data, on the other hand, is collected directly from the source node. After the data is collected and stored in the database, both SNMP and non-SNMP data are processed in the same way.

Provisioning Reports

Report packs might 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.

Modify an existing file after the type discovery process runs for the first time and before any other process starts. After creating or modifying the file, you can import it to the correct location. For instructions about modifying and importing these files, see the respective report pack documentation.

Polling Parameters

After the system locates the nodes to poll, it needs to know what data to collect and how frequently. The system obtains these instructions from a polling policy. Report packs contain the polling policy definitions. If you set up a distributed environment you may need to modify an existing polling policy or create a new one. See [Chapter 8, Managing Polling Policies](#) for more information about polling policies.

Process Data

PI converts the raw data into rate data as necessary, performs summarization routines on data, and ages the data out of database. The summarization routines convert rate data into hourly, weekly, monthly, quarterly, and yearly data. The summarization routines also produce forecasting data based on past performance. This data enables you to forecast future problems.

PI 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 of managed object are routers, interfaces, customer, and location.

After the data is collected, the system assigns timestamps to the records and stores the data in data tables. The system then converts the raw data into delta (or rate) data, as necessary. Delta data is the difference between two values over a period of time.

The aggregation process continues on the delta data. The system reads data from a source (input) table, aggregates it, and stores it in the destination (output) table. The aggregation process creates different destination tables for different types of statistics. Statistics include hourly, weekly, monthly, quarterly, and yearly summaries; 30-day, 60-day, 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, PI either ages or archives the data. Aging involves the concept of retention time, which is the number of days the data is stored 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 457 for more information about changing the retention time for a table.

Archive Data

You can save your data for longer duration without affecting the summarization and reporting performance of PI. PI stores the older data in a different tablespace in the PI database schema. The PI reporting solution can access the data in these tablespaces in the manner as the data in standard PI database schema.

Data tables store the data collected from the network devices. Storing data in the data tables for a long time slows down the performance of the summarization and reporting solution. Therefore, the data stored in a data table is deleted once the specified retention time for that table is reached. PI uses archive tables to store the data. This helps you avoid summarization and reporting solution performance problems and still retain the data for a long time for reporting purpose.

Archiving involves:

- **Archive flag**– indicates whether an archive table is required for a data table or not. By default, the value of the archive flag is set to OFF. If you want to archive the data of a data table, set the value of the archive flag to ON.
- **Archive retention time**: indicates the number of days the data is stored in the archive table.
- **Active table**: the data table with which an archive table is associated.

While installing a report pack, the `datapipe_manager` creates a retention profile for each data table based on the retention time specified in the TEEL file. The system uses default values for the profile name and archive retention time, based on the data table's time type, to create a retention profile.



If a retention profile of the same name and retention time exists in the profile dictionary, the system uses that profile and does not create a new profile.

If the archive file is set to ON for a data table, the following happens:

- 1 The `datapipe_manager` creates the archive table for a data table and assigns the retention profile of the data table (which is also called the active table) to the archive table.
- 2 The aging process moves the data from the data table to the archive table when the specified retention time for that data table expires. For example, if the retention time for a data table is set to 10 days and the

archive retention time is set to 60 days, the aging process moves the data from the data table to the archive table after 10 days. In the archive table, data is available for the next 60 days.

- 3 The aging process deletes the data from the archive table after 60 days. This means that data is stored in the active table and the archive table combined for 70 days. See [Archiving Data](#) on page 589 for more information about changing the archive retention time and archive flag for a data table.

Report Data

PI automatically generates reports from summarized data through a reporting solution known as *report pack*. A report pack provides report templates and all the supporting files and scripts for a particular type of solution. You can view these reports through a web-based interface (Web Access Server) or through the client applications— Report Builder and Viewer. See [PI Report Packs](#) on page 41.

You can also create your own customized reports with Report Builder (a separate product). See [Creating Custom Reports](#) on page 41

Both the automatic and custom methods create the report template, which is a set of instructions that describe the report.

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 but the 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 enables you to change the data contained in individual database tables. All elements available for reports are available for forms. In addition, you can add form elements that enable 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 the information about an entity for a specific report pack.

PI Report Packs

After you install 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 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. In the report template, you can define how to display the data, what information to collect from the database, and how to link the data to the report.

PI Components

PI is made up of the following components:

- [PI Database Schema](#)
- [PI Performance Manager](#)
- [Web Access Server](#)
- [PI Clients](#)

PI Database Schema

The PI database is the primary repository for the data that is collected, processed, and reported on by other PI applications. The [org.apache.commons](#) PI database configures the database server for PI and creates the PI schema. For Oracle, PI database creates new tablespaces or gets installed in existing tablespaces set up for PI. For Sybase, it creates the database segments for PI. The Management Server is installed with this component.

PI Performance Manager

The Performance Manager provides a mechanism to store, process, and access data collected from applications and network devices throughout your network. The Performance Manager automatically aggregates, summarizes, and ages out the data.

The Performance Manager consists of a relational database and the PI server. The PI server collects and processes the data and stores it in the database. The database size depends on various factors, including the amount of data you want to store and the number of nodes that you poll. The server contains the programs that collect, process, and generate report on the data.

The Performance Manager can be installed on the Central Server which is the main server for the PI system. Typically, the data processing and reporting functions are performed by the Central server.

The Central Server usually has a large amount of disk space, memory, and processing power. Typically, there is only one Central Server.

The Performance Manager can also be installed on the satellite servers. These servers contain data and the PI system. Typically, the satellite servers perform only the first level of processing and transfer the processed data to the Central Server.

You can configure the satellite servers to transfer the data to the Central Server without processing. This, however, may negatively impact the performance if the amount of data to transfer is large.

You can also configure the satellite servers to transfer the data to the Central Server after complete processing.

A large deployment of PI can contain multiple satellite servers.

Web Access Server

The Web Access Server is both a service and an application. It is installed as a service and provides authentication and configuration services to the PI client. It also provides web-based reporting and user administration utilities. The following PI client applications are automatically installed with the Web Access Server:

- Report Builder
- Report Viewer
- Web Access Server (application server)
- Management Console

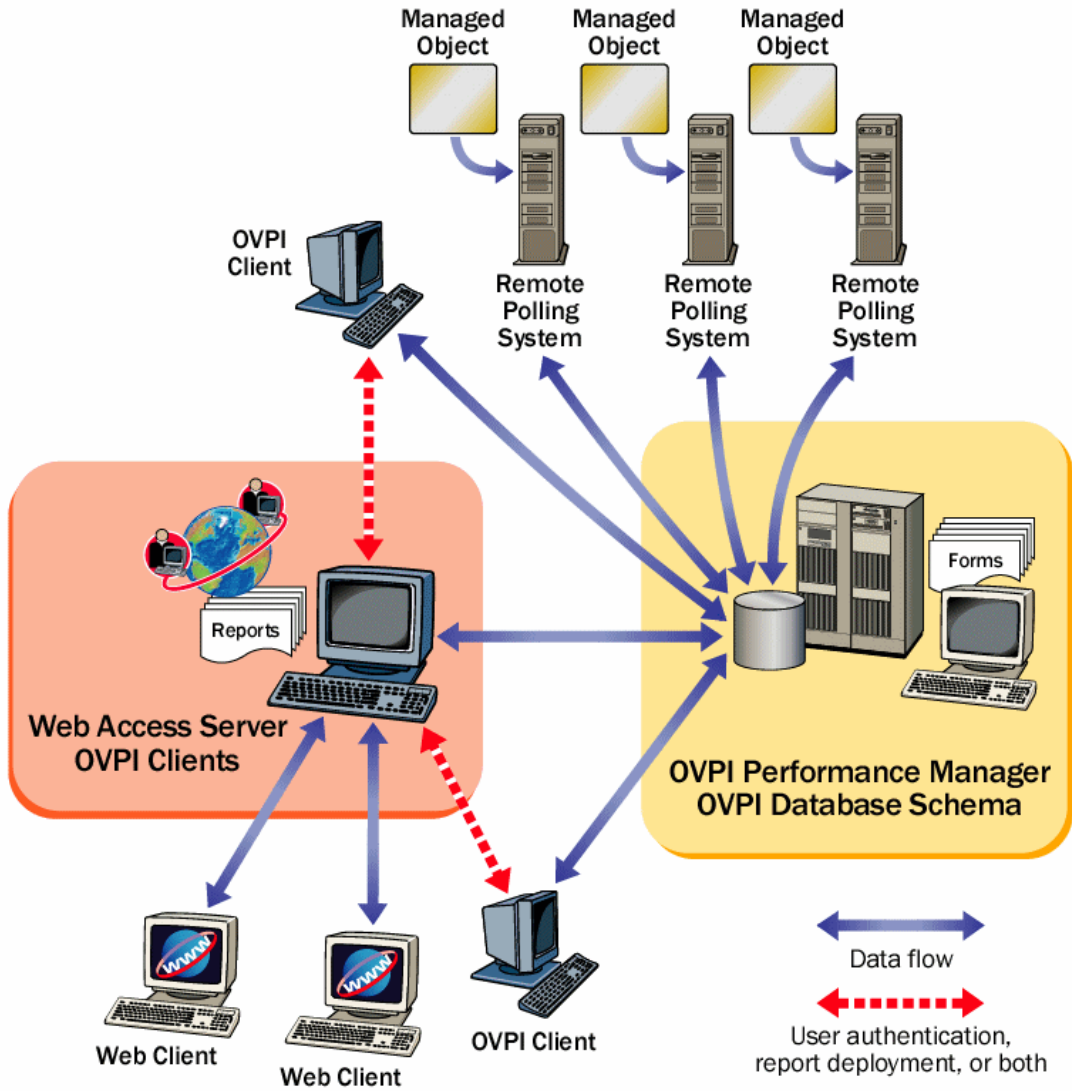
PI Clients

The clients provide the interface to PI. You can install a client on any computer connected to the network. The PI client applications are:

- Report Builder: Enables you to build and customize reports with multiple tables and graphs based on data from the PI database. See the *HP Performance Insight Guide to Building and Viewing Reports* for details.
- Report Viewer: Enables you to interactively view, modify, save, and print reports. See the *Performance Insight Guide to Building and Viewing Reports* for details.

- Management Console: Enables you to perform administrative tasks using the following applications:
 - Database Tuner: Enables you to set many database configuration parameters to improve overall data throughput.
 - MIB Browser: Enables you to analyze, query, and troubleshoot Structure of Management Information (SMI)-compliant SNMP MIBs.
 - SNMP Discovery: Enables you to find the nodes in your network and determine whether or not a node is SNMP manageable.
 - Type Discovery: Enables you to determine the device type of a node.
 - Package Manager: Enables you to install and uninstall packages and deploy and undeploy them from the Web Access Server component.
 - Object Manager: Enables you to view, edit, and delete the managed objects on your system.
 - System Manager: Enables you to add, modify, and delete database and Web Access Server systems.
 - Group Manager: Enables you to create groups from the managed objects referenced by the polling policies for your enterprise.
 - Polling Policy Manager allows you to manage polling policies, import and manage nodes, and manage polling groups.
 - Table Manager: Enables you to monitor your database, view data and property tables, and set defaults for both individual and global tables.
 - Location Independent Reporting (LIR) in a distributed environment enables PI to report on data, irrespective of the location of the node in the network.

The following figure shows an example of the PI components.



PI Architecture

The process of data collection, processing, and reporting can be divided into the following tasks:

- Data collection
- Raw data to rate data processing
- Rate data to hourly data summarization
- Hourly data to daily, monthly, forecasting summarization, or all
- Report request servicing
- Report viewing and building

The role of any PI system—Standalone System, Remote Poller, Satellite Server, Central Server, or Web Access Server—is defined by the tasks it must perform, with each task requiring a particular set of PI components.

Standalone Setup

In a standalone setup all the PI components are installed on a single server. The standalone server is capable of gathering and processing data, producing reports that can be accessed from the web server, and running the Report Viewer and Report Builder applications.

Distributed PI Setup

A distributed setup includes a Central Server and one or more satellite servers. It might or might not have remote pollers.

Central Server

All the PI components, except the Web Access Server, are installed on the Central Server. The Central Server receives partially processed data from one or more satellite servers. The Central Server, usually, processes data collected on an hourly basis and also the complete data collected. The central server performs the daily, monthly, and forecasting summarization of data.

Satellite Server

All the PI components, except the Web Access Server, are installed on the satellite server. The satellite server, usually, only processes hourly data. The partially aggregated data is transferred to the Central Server for complete aggregation. One or more remote pollers can connect to a satellite server.

The satellite server typically performs the following tasks:

- Processes the raw data into rate
- Summarizes the rate data on an hourly basis.

In some cases, the satellite server runs as a complete standalone system performing full data aggregation and reporting. However, the satellite server still transfers raw data, rate data, or partially aggregated data to a central server. This enables you to view regional reports at each satellite server and network-wide reports on the Central Server.

Remote Poller

A remote poller only collects the data from the polled devices and transfers the polled data to the satellite server. It does not have a database and it does not perform any processing.

The remote poller is located close (in terms of network path) to the target/pollled devices. This reduces the distance the SNMP packets must travel.

- ▶ Ensure that duplicate polling is avoided when remote pollers are configured. Duplicate polling occurs when the same device is polled by two or more pollers.

Web Access Server

The Web Access Server system hosts the PI web site. It is also used by PI clients to authenticate their users.

- ▶ The Web Access Server often requires significant memory and CPU. Therefore, to improve performance, you might install the Web Access Server on a system on which other PI components are not installed.

Data Tables and Property Tables

Data Table

The data collected from the network devices are stored in data tables. Before storing the data, PI assigns time stamps to the records. The data tables contain time-series data about managed objects.

All PI data tables must have at least two basic data columns, the `ta_period` and `dsi_key_id_`. The `ta_period` column contains the date/time stamp for the collected data. The `dsi_key_id_` column contains the identification for the managed object that belongs to the data, which is the identifier that connects the data table to the property table.

Property Table

Each data table requires a property table to further define the data. Property tables contain lists of unique managed objects (nodes) with their attributes (properties). Each row in the property table is a unique entry for the managed object. Each row has a set of columns that are the attributes or properties of the object.

Each property table can have many data tables reference the property table. The data tables contain time-series data about the objects in the property table.

Example

If you are managing an object, such as books in a publishing company, you would have a property table that contains all the static information about the books. The static information includes the attributes (properties) about the books that typically do not change over time. Each book would be one row in the table and the attributes would be the columns in the table. On the other hand, you would also have data tables that track the dynamic information about the books, such as the number of books sold to customers.

In this example, the property table that contains the list of books is called `k_books`. Books are the managed object. The attributes of the books are title, subject, author, publisher, cover design, number of pages, and ID number.

Each of these attributes appears as a column in the table and the combination of some subset of them uniquely identifies the book. [Table 1](#) shows the layout of the columns for this example.

Table 1 Layout of Sample Property Table - k_books

dsi_key_id	Title	Subject	Author	Publisher	Cover_Design	Num_Pages	ID_Num	target_name	table_key	Description	status	status_time	dsi_bv_state
book1													
book2													
⋮													
bookN													

PI Administration

You can use the following utilities to configure, administer, and monitor PI:

[PI Management Console](#)

[Web Access Server](#)

[Command Line Utilities](#)

PI Management Console

The Management Console is an administrative interface to a range of PI operations, especially related to the functions of PI Performance Manager.

The tools available through the Management Console enable you to populate the database with device and element information and control the PI environment. For more information on Management Console, see [Chapter 5, Using the PI Management Console](#).

Web Access Server

The Web Access Server (WAS) home page enables you to manage the application server. Tools available through the PI home page include report scheduling, user administration, report library management, PI for networks process configuration, and log file viewing. WAS can be used to deploy and undeploy reports and also to delete reports as required. Through the Web Access Server you can run reports on demand and look at the results in HTML format. You can also personalize the look and feel of the home page.

Command Line Utilities

Command line utilities include bulk processing versions of interactive functions (that are done one-by-one). See the *HP Performance Insight Reference Guide* for details on each command line utility.

3 Configuring HP Performance Insight

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

The following tasks are required to configure PI:

- Tune the database.
- 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: 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 678 for details.



The Database Tuner is only available for Sybase databases.

Oracle

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

During the PI installation, there was opportunity to change the default values of these parameters. For more information, see “Changing Specific Oracle Configuration Parameters (Parameters Tab)” in Chapter 5 of the *HP Performance Insight Installation and Upgrade Guide for Oracle Databases*.

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. See your Oracle documentation for additional tuning recommendations.

Task 2: Discover Your Network

Before PI can collect data, PI 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.

PI 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 CiscoWorks). 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. See your report pack documentation before you begin the discovery process. To locate reporting solutions documentation, see *HP Performance Insight Release Notes*.

Discovering Nodes

- 1 From the Management Console's Tools menu, select **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 Type 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 [The Discovery Process](#) on page 152 for complete details. Alternatively, you can use the `trend_discover` command. See the *HP Performance Insight Reference Guide* for details.

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 275 for details about how to format the ASCII file.

To import a list of nodes:

- 1 From the File menu, select **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, see the *HP Performance Insight Reference Guide*.

Task 3: Install Report Packs

This procedure explains how to install report packs using the RNS CD.



Before you continue with this procedure, make sure that you have read the user guides for the report packs you want to install. To obtain these guides, go to:

http://ovweb.external.hp.com/lpe/doc_serv/

- 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.

Operating System

Command

Sun Solaris 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 **root**, type the following:

```
mkdir /cdrom
```

```
mount -r /dev/dsk/cd_devicename /cdrom
```

In this instance, *cd_devicename* is the name of the CD-ROM device.

HP-UX You can use the SAM utility to mount the CD-ROM, **or** you can type the following commands:

```
mkdir /cdrom
```

```
mount /dev/dsk/cd_devicename /cdrom
```

In this instance, *cd_devicename* is the name of the CD-ROM device.

- 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 55; 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 **PI 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.

Type the PI administrator username and password, and click **Next**. You can use the default PI 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 the report packs and datapipe 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.

Task 4: Discover Device Types

This discovery determines the device type of a discovered node and then records this information in the database. This information lets PI 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 56), 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 5](#).

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. See the *HP Performance Insight Reference Guide* for details.

Task 5: Change Polling Policies

A polling policy defines the frequency with 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 PI and your network performance.

For information about configuring polling policies for remote pollers, see [Configuring Polling Policies for Remote Pollers](#) on page 69.

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 8, Managing Polling Policies](#) for more details about polling policies.

Task 6: Back up your system

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



It is recommended that you stop all PI services and the database before you back up your system. For information about stopping the PI services, see [Starting and Stopping HP Performance Insight Processes](#) on page 103. For information about stopping Sybase, see the *HP Performance Insight Installation and Configuration Guide*. For information about stopping Oracle, see your Oracle documentation.

Task 7: Back up the database

It is recommended that you back up your database after you install and configure PI and, thereafter, back up the database weekly. See [Backing Up Your Oracle Database](#) on page 87 or [Backing Up Your Sybase Database](#) on page 87 for details.

Additional Configuration Tasks

Configuring User Access

All PI 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, PI 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 PI 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, 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, 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 552 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 512 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 Type 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 character, can be a maximum of 255 characters, and can only contain alphanumeric characters, spaces, underscores (_), at sign (@), period (.), and hyphens (-).
 - The user’s real name and company/department have no minimum length (they can be blank) and can be a maximum of 255 characters.
 - The password can be a minimum of four 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 575.
 - 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 552.
 - 9 To select 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 type the 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 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 525 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 565 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 to manage users (`userctl`). See the *HP 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 563 for details.

To create a group with a filter:

- 1 Start the Web Access Server. See [Accessing the Web Access Server](#) on page 512 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, type 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 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. Only table columns from property tables are valid values in this field. See [Viewing Data Table Column Information](#) on page 401 for details.

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

For example, suppose you are a service provider who has the ABC Company as a client. You want to 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 report pack. 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 to manage groups (`groupctl`). See the *HP Performance Insight Reference Guide* for more information about these commands.

Managing Database Changes

You can maintain your databases using the System and Network Administration application, which is available from the Management Console.

To use this application:

- 1 Start the Management Console. See [Starting the Management Console](#) on page 118.
- 2 Click the **Systems** icon.

For more information about using the System and Network Administration application to manage your database, see the following:

- [Adding a Database](#) on page 375
- [Viewing Database Properties](#) on page 377
- [Changing Database Initialization Parameters](#) on page 380
- [Increasing Database Size](#) on page 381
- [Changing the Database Server](#) on page 385

Collecting Data from SNMPv3 Nodes

PI can collect data from SNMPv3 nodes. PI cannot discover these nodes. You must configure PI to support data collection from these nodes.

To configure PI to collect data from SNMPv3 nodes, do the following:

- 1 Install SNMP Security Pack version 16.1.0.35 or later on the system where you installed PI.

SNMP Security Pack is a product from SNMP Research. You must contact SNMP Research to purchase this product. See <http://www.snmp.com> for more information.

The pack's BRASS product lets the PI poller identify SNMPv3 nodes with v3 community string profiles. PI has a utility (`brass_v3_seed`), which lets you supply authentication and privacy passwords.



In a distributed PI configuration, you need to install the SNMP Security Pack on any system that will be collecting data; thus, you must install the pack on all remote pollers and all systems on which PI's Performance Manager component is installed.

- 2 Configure agents for your SNMPv3 nodes using the SNMP Security Pack's SNMPv3 Configuration Wizard. See the SNMP Security Pack documentation for more information.
- 3 Use the `brass_v3_seed` utility to supply authentication and privacy passwords for the nodes.
- 4 Import the nodes.
- 5 Create community strings for the nodes.
- 6 Assign the nodes to a type group.
- 7 Configure remote pollers.

The procedures described in this section to collect SNMPv3 data assume that you installed the SNMP Security Pack along with PI.

Configuring SNMPv3 Nodes

Configure the agents for your SNMPv3 nodes using the SNMPv3 Configuration Wizard. See the SNMP Security Pack documentation for more information.

Using the `brass_v3_seed` Utility

You must use the **`brass_v3_seed`** utility before you can start the actual data collection from an SMMPv3 node. This utility will store the passwords for the nodes in an encrypted format. This utility is located in the `bin` folder of the PI installation directory.

The `brass_v3_seed` command has the following options:

```
brass_v3_seed  -f filename
                  [ -d debug_level ]
                  [ -E ]
```

The -f Option

The `-f` option is required. It must specify the name of an input file. This input file must be formatted as follows:

```
<host>, <port>, <username>, <authorization protocol>, <authorization
password>, <privacy protocol>, <privacy password>
```

The host, port, and username fields are required. The authorization protocol is MD5 or SHA. The privacy protocol can be any one of the following: DES, AES, or 3DES.

The -d Option

The `-d` option is used to display debug messages and is optional. The various debug levels supported are 1, 2, 3. Level 1 provides the least amount of debugging information, and Level 3 provides the maximum amount of detail.

The -E Option

The `-E` option tests the availability of the nodes specified in the input file.

For example, if you were adding a host named `neast1` to an input file called `snmpv3_node.txt`, this file might look like the following:

```
neast1,161,ewing,MD5,cat,DES,dog
```

The command line for the utility might be the following, depending on whether you have chosen the debug and test options:

```
brass_v3_seed -f snmpv3_node.txt -d 1 -E
```

Importing SNMPv3 Nodes

After running the `brass_v3_seed` utility, you need to import the SNMPv3 nodes to the `ksi_managed_node` table by creating an ASCII file, and then import that file using Polling Policy Manager or the `node_manager` command.

Suppose you want to import a node called `neast1`. To import this node:

- 1 Create an ASCII file called `snmpv3_node.txt`. (For details, see [Importing Nodes](#) on page 275.)

The contents of this file will look like the following:

```
neast1, public, private,,,,1,,smp_1, 161,10,500,20,50
```

- 2 Import the file by doing one of the following:
 - Select **Import Nodes** from the Polling Policy Manager's File menu. Navigate to and select `snmpv3_node.txt` in the Import From dialog box, and click **OK**.
 - Use the `node_manager` command to import the file. See the *HP Performance Insight Reference Guide* for details. This command will look like the following:

```
node_manager -import -file snmpv3_node.txt
```

Creating Community String Profiles

After you store passwords for the nodes, you will need to create community string profiles for them.

- 1 Start Management Console. See [Starting the Management Console](#) on page 118.
- 2 Click the **Administration** icon in the navigation pane. The Administration Snapin opens.
- 3 In the navigation tree provided in the Administration Snapin, click **Node Management**.
- 4 Create community string profiles for the nodes by doing the following:
 - a Select **Manage Community String** from **Node Management**.
The current Community String Profiles are shown in the viewing pane.
 - b Click **Create**. The Create Community String Profile dialog box opens.
 - c Type the name for the profile in the Profile Name box. This name can be a maximum of 64 characters.

To enable SNMP Get Requests for v3, the string must begin with one of the following:

- 3N (No Authentication, No Privacy)
- 3A (Authentication, No Privacy)
- 3P (Authentication, Privacy)

The community string should use one of the above with the SNMPv3 username you assigned when configured for the agent using the SNMPv3 Configuration Wizard.

For example, if the SNMPv3 agent has `admin` as its username and Authentication and No Privacy permissions, the community string for this agent will be `3A/admin`. You can have other combinations of usernames and authentication and privacy permissions.

- d Type the value for the read permission in the Read box.
 - e Type the value for the write permission in the Write box.
 - f Click **OK**.
- 5 Repeat this procedure for each SNMPv3 node.



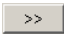
The read community string (not the community string profile name) must contain the `"3?/<user_name>"` string. The community string profile name can have any value.

Assigning Nodes to a Type Group

SNMPv3 nodes will not be automatically discovered by type and placed into specific groups. You have to manually assign the nodes to a group using the Polling Policy Manager application.

To assign the nodes to different groups:

- 1 Open the Polling Policy Manager application.
See [Chapter 8, Managing Polling Policies](#) for more information.
- 2 From the Edit menu, select **Polling Groups**.
The Edit Polling Groups dialog box opens.
- 3 Select **All Nodes of the Same Type** from the Select Kind of Group box.
- 4 Select a group to which to add the node or nodes from the Select Group to Poll From box.
- 5 Click **Edit**.

- 6 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.
- 7 When you are finished, click **OK**.
- 8 Click **Close**.

Configuring Remote Pollers for SNMPv3 Nodes

Remote pollers will not poll SNMPv3 nodes automatically. You will need to install the SNMP Security Pack product on any remote pollers. See the SNMP Security Pack documentation for information.

Configuring a Distributed Installation

In a distributed installation, the PI components are installed on multiple physically separate systems and operate as a single, distributed system. See the *HP Performance Insight Installation and Upgrade Guides* 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 a PI database. This database is not present on the remote poller system. The remote poller collects data from the nodes in the network and stores the data on a server running the PI 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 that 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 8, Managing Polling Policies](#) for more details about polling policies.

Creating New Policies for Multiple Pollers

When you have multiple pollers polling different devices, you must 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 118.
- 2 Click the **Groups** icon.

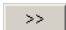
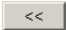
The Group Management window appears in the Management Console. The left side of the window displays the group management tree, which lists all the groups for the type of view that is selected.

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 11, Grouping Managed Objects](#).

Creating the First Polling Policy

To modify a polling policy:

- 1 Start the Management Console. See [Starting the Management Console](#) on page 118.
- 2 Click the **Polling Policies** icon.
- 3 From the File menu, select **Create Polling Policy**.
The Create Polling Policy dialog box opens.
- 4 Change the values for the following options:
 - a Type a distinct and unique name (Policy 1 in this example scenario) for the polling policy in the Policy Name box.
 - b Select the data you want to poll for.
 - c Select **Custom Groups** from the Collect Data From list.
 - d Select the name of the first group (**DevicesA** in the example scenario) from the Select Group to Poll From list.
 - e Select a polling interval.

- f Select a datapipe.
- g 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.
- h Type a description for this polling policy in the Description box.
- i Click **OK**.

You have created the first polling policy called Policy 1.

Creating the Second Polling Policy

To create the second polling policy:

- 1 Start the Management Console. See [Starting the Management Console](#) on page 118.
- 2 Click the **Polling Policies** icon.
- 3 From the File menu, select **Create Polling Policy**.
The Create Polling Policy dialog box opens.
- 4 Do the following:
 - a Type a distinct and unique name (Policy2 in this example scenario) for the polling policy in the Policy Name box.
 - b Select the data you want to poll for.
 - c Select **Custom Groups** from the Collect Data From list.
 - d Select the name of the second group (**DevicesB** in the example scenario) from the Select Group to Poll From list.
 - e Select a polling interval.
 - f Select a datapipe.
 - g 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.
 - h Type a description for this polling policy in the Description box.
 - i Click **OK**.

You have created the second polling policy called Policy 2.

Distributing Polling Between Multiple Pollers

This section describes the steps to divide polling between different pollers using View Groups.

When you install a Datapipe package on Performance Insight (PI), the package creates a set of polling policies on the server where the package is installed. Each default polling policy collects data from a “Type” group which is a set of nodes that support a particular SNMP MIB or that support a particular HP agent.

If you are using multiple pollers, you can divide the polling between these pollers using one of the following options:

Option 1: Assign a Subset of the Polling Policies to Each Poller

You can use this option if there is too much data for a single poller to collect. To configure this option, follow these step:

- 1 Open the HP Performance Insight Management Console.
- 2 Click on **Administration** in the left pane and then select **Polling Policies**.
- 3 Double-click a policy that you want to assign to a different poller.
- 4 Enter the poller in the Poll From field and click **OK**.

Repeat this step for all the policies that you want to move.

Example

You can assign the Interface collections to one poller and the Device Resource and System Resource collections to another poller.

Option 2: Assign a Subset of the Nodes to Each Poller

You can use this option to divide polling by region or by client network. In this case, there is a separate set of polling policies for each poller, and each poller collects those policies from a sub-set of the nodes.

PI supports a second category of enumerated groups called View groups. You can use View groups to divide the managed nodes between different pollers.

You can further, create a “Combination of Type and View” group for each Type group and each poller View group and create polling policies based on these “Combination of Type and View” groups.

When creating these groups, ensure that each node is polled by only one poller. If you add a new node to the list of managed nodes, you must also add the node to the appropriate View group.



There are special processes for Interface Discovery and the ifEntry data collections associated with the ATM, Frame Relay and Interface Reporting report packs that you need to follow if you use this option. See the respective user guides for more information.

If you are using System Resource datapipes, copy the *.bcpg files from the server to the \$COLLECT_HOME/SR/ directory on each remote poller.

Use Case

This use cases discusses a scenario where you need to configure the Device Resource Cisco Router Datapipe with three pollers for three different regions. The default installation of the datapipe has four Type groups and there is one collection per type.

The standard polling Type groups for the Device Resource Cisco Router Datapipe are:

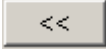
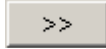
- DevResCiscoRouter
- DevResCiscoRouter_Memory
- DevResCiscoRouter_Process
- DevResCiscoRouter_System

To configure polling, create new View groups for the three pollers. Complete these tasks:

Task 1: Create View Groups

Follow these steps:

- 1 Open the HP Performance Insight Management Console.
- 2 Click on **Administration** in the left pane and then select **Polling Policies**. The Polling Policy pane opens.
- 3 From the Edit menu select **Polling Groups....** The Edit Polling Groups window opens.

- 4 Select **All Nodes in Same View** and click **Create....** The Create View Group window opens.
- 5 Enter a group name in the Group Name box (for example, Poller1) and description for the group in the Description box.
- 6 Select the nodes, you want to add to the group, from the Other Nodes: list and move them to Nodes in Group: list. To move the nodes between the lists use the  or  buttons.
- 7 Click **OK**.
- 8 Repeat steps 4 through 7 to create the other two groups as well.

Task 2: Create Combination Groups

Follow these steps:

- 1 In the Edit Polling Groups window, select **A Combination of Type and View** and click **Create** The Create a Combined Type and View Group window opens.
- 2 Enter a group name in the Group Name box (for example, P1_DevResCiscoRouter) and description for the group in the Description box.
- 3 Select **DevResCiscoRouter** from the Type Group drop-down menu.
- 4 Select and from the drop-down menu.
- 5 Select the first view (Poller 1) created in [Task 1](#) on page 75. You can choose to view the list of nodes by clicking **View List of Resulting Nodes**.
- 6 Click **OK**.
- 7 Repeat steps 9 through 14 for create all the combinations groups.

You must create the following groups:

- P1_DevResCiscoRouter
- P1_DevResCiscoRouter_Memory
- P1_DevResCiscoRouter_Process
- P1_DevResCiscoRouter_System
- P2_DevResCiscoRouter
- P2_DevResCiscoRouter_Memory

- P2_DevResCiscoRouter_Process
- P2_DevResCiscoRouter_System
- P3_DevResCiscoRouter
- P3_DevResCiscoRouter_Memory
- P3_DevResCiscoRouter_Process
- P3_DevResCiscoRouter_System

8 Click **Close** to close the Edit Polling Groups window.

Task 3: [Modify the Existing Polling Policies for the First Poller.](#)

Follow these steps:

- 1 In the Polling Policy pane, select the DevResCiscoRouter Datapipe from the Datapipe drop down menu. The four default polling policies appear in the table.
- 2 Double-click a policy (xDRCR_DeviceSystem). The Edit Polling Policies window opens.
- 3 Select **A Combination of Type and View** from the Collect Data From dropdown list.
- 4 Select the first combination group you created in [Task 2](#) on page 76 (P1_DevResCiscoRouter) from the Select Group to Poll From dropdown list.
- 5 Select the server name of Poller1 from the Polling Assigned to dropdown list.
- 6 Click **OK**.
- 7 Follow steps 2 through 6 to edit the other three default polling policies (in [step 4](#) on page 77 select the 2nd, 3rd, and 4th combination groups respectively).

After editing the four policies, you must have results similar to the following:

Polling Policy Properties	Original Values	New Values
xDRCR_DeviceSystem		
Collect Data From	All Nodes of the Same Type	A Combination of Type and View
Select Group to Poll From	DevResCiscoRouter	P1_DevResCiscoRouter
Polling Assigned to	<SERVER>	<Name of Poller1 server>
xDRCR_MemoryPool		
Collect Data From	All Nodes of the Same Type	A Combination of Type and View
Select Group to Poll From	DevResCiscoRouter_Memory	P1_DevResCiscoRouter_Memory
Polling Assigned to	<SERVER>	<Name of Poller1 server>
xDRCR_CPUTotal		
Collect Data From	All Nodes of the Same Type	A Combination of Type and View
Select Group to Poll From	DevResCiscoRouter_Process	P1_DevResCiscoRouter_Process
Polling Assigned to	<SERVER>	<Name of Poller1 server>
xDRCR_Buffer		
Collect Data From	All Nodes of the Same Type	A Combination of Type and View
Select Group to Poll From	DevResCiscoRouter_System	P1_DevResCiscoRouter_System
Polling Assigned to	<SERVER>	<Name of Poller1 server>

Task 4: Create New Polling Policies for the Other Two Pollers

Follow these steps:

- 1 From the **File** menu, select **Create Polling Policy...** The Create Polling Policy window opens.
- 2 Create four new polling policies for each poller. In each case:
 - a Enter a unique policy name in the Policy Name box.

- a Select **A Combination of Type and View** from the Collect Data From dropdown list.
- b Select the appropriate P2 and P3 combination groups from the Select Group to Poll From list.
- c Select the server name of Poller2/Poller3 from the Polling Assigned to dropdown list.
- d Select a polling interval from the Poll Interval dropdown list.

After creating the new policies, you must have results similar to the following:

Policy Properties	Poller2	Poller3
xDRCR_DeviceSystem		
Policy Name	P2_DevResCiscoRouter_System	P3_DevResCiscoRouter_System
Collect Data From	A Combination of Type and View	A Combination of Type and View
Select Group to Poll From	P2_DevResCiscoRouter	P3_DevResCiscoRouter
Poll Interval	60	60
Polling Assigned to	<Name of Poller2 server>	<Name of Poller3 server>
xDRCR_MemoryPool		
Policy Name	P2_DevResCiscoRouter_Memory	P3_DevResCiscoRouter_Memory
Collect Data From	A Combination of Type and View	A Combination of Type and View
Select Group to Poll From	P2_DevResCiscoRouter_Memory	P3_DevResCiscoRouter_Memory
Poll Interval	15	15
Polling Assigned to	<Name of Poller2 server>	<Name of Poller3 server>
xDRCR_CPUTotal		
Policy Name	P2_DevResCiscoRouter_CPU	P3_DevResCiscoRouter_CPU
Collect Data From	A Combination of Type and View	A Combination of Type and View

Policy Properties	Poller2	Poller3
Select Group to Poll From	P2_DevResCiscoRouter_Process	P3_DevResCiscoRouter_Process
Poll Interval	15	15
Polling Assigned to	<Name of Poller2 server>	<Name of Poller3 server>
xDRCR_Buffer		
Policy Name	P2_DevResCiscoRouter_Buffer	P3_DevResCiscoRouter_Buffer
Collect Data From	A Combination of Type and View	A Combination of Type and View
Select Group to Poll From	P2_DevResCiscoRouter_System	P3_DevResCiscoRouter_System
Poll Interval	15	15
Polling Assigned to	<Name of Poller2 server>	<Name of Poller3 server>



If you plan to use only two pollers, you can simplify the View group administration so that only one View group needs to be maintained. To do this, configure the “Combination of Type and View” group for one poller as a Type *and* View group and for the other poller configure the group as “type not in view”.

You might also need to configure SNMP Type Discover to run on each poller for the group of nodes in that poller's View group. For example, you can add a similar line in the `trendtimer.sched` file on the Poller1 server.

```
24:00+2:00 - - {DPIPE_HOME}/bin/trend_discover -t -v poller1
```

To find out which polling policies and polling groups are associated with a specific datapipe, go to the `*_Collection_Policies` file in the appropriate datapipe packages directory. The first field is the name of the polling policy and the 8th field is the name of the polling group.

4 Maintaining HP Performance Insight Installation

This chapter describes the administration tasks that you should perform at specific intervals to maintain your PI installation.

Post-Installation Administration Tasks

The day after you install PI, you must start performing the following administration tasks:

- Verify that PI is running.
- Schedule reports.
- Modify users, groups, and group filters.

Verifying that HP Performance Insight is Running

To verify that PI 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 `$SYBASE/$SYBASE_ASE/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 time-out errors.

Locating the Log File

The `trend.log` file contains messages for all of PI'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 PI is installed, the log files are located in the `log` folder under the PI 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 Click **Start** → **Settings** → **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 **Environment Variables**.
- 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. See the *HP 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 398 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 565.) 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.

The different report formats are: `.html`, `.srep`, `.csv`, and `.pdf`. You can email `.pdf` and `.csv` reports. 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, see the *HP 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 525.

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 select 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 20, Interpreting Log Files](#).

Weekly Administration Tasks

You should back up the database at least once a week. 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, do either of the following:

- **Offline backup:** To perform this backup, you must shut down Oracle and PI.

To shut down PI, stop the PI HTTP Server and PI Timer services. See [Starting and Stopping HP Performance Insight Processes](#) on page 103 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, see 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 the 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 not created automatically. For more information about creating a backup server, see [Creating a Backup Server on a UNIX System](#) on page 88 for details.
- On Windows systems, the backup server is not created automatically. For more information about creating a backup server, see [Creating a Backup Server on a Windows System](#) on page 91 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 92 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 describes how to create a backup 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, see 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. To create a backup server with a different name, use the `srvbuild` utility.

To create a new backup server:

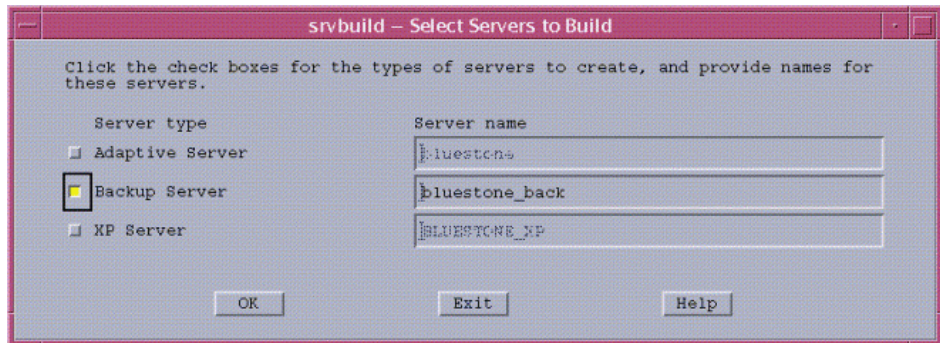
- 1 Log on to the system as `root`.
- 2 Open a terminal window.
- 3 Ensure that the `$SYBASE` variable points to the directory in which Sybase is installed.
- 4 Change to the directory in which Sybase is installed by entering the following command:

```
cd $SYBASE/bin
```

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

```
./srvbuild
```


The srvbuild - 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.

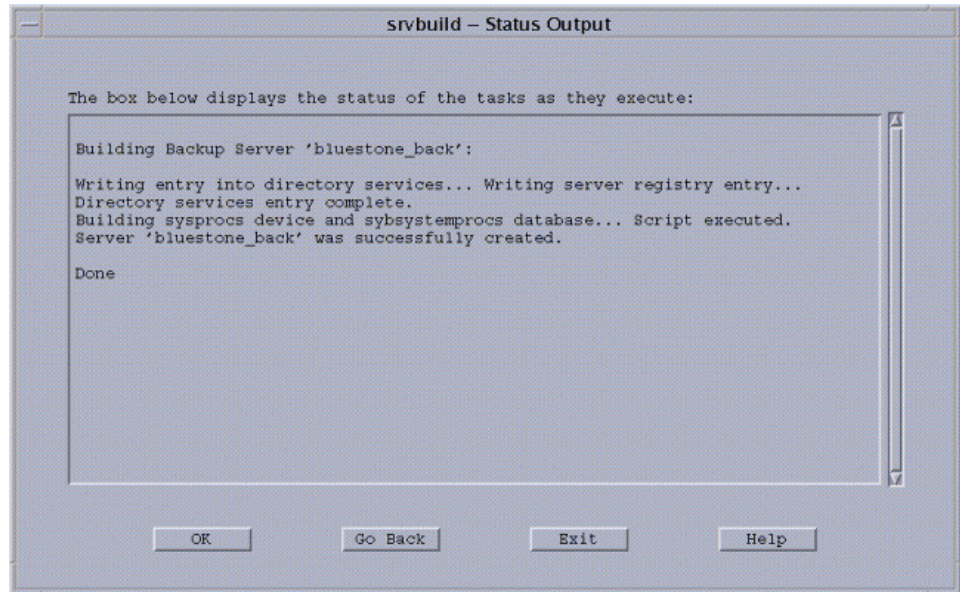
- 7 To accept the default name, click **OK**.

The srvbuild - Server Attribute Editor dialog box opens.



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

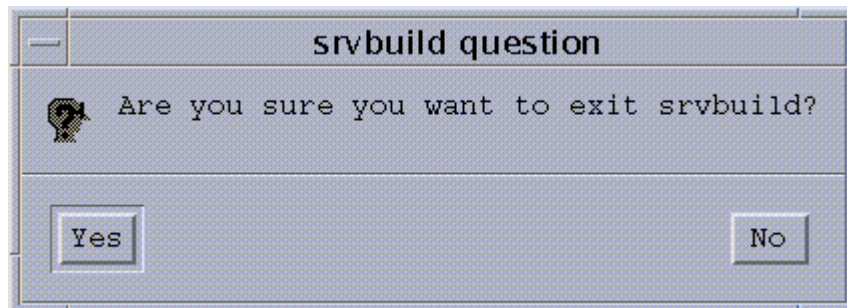
The Status Output window opens.



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

- Click **Exit**.

The `srvbuild` question dialog box opens.



- Click **Yes**.

After you create the new backup server, you need to drop the entry for the old backup server from the `syssservers` 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. 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 Click **Start** → **Programs** → **Sybase** → **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. 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 Type 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

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 98 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 Type 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 and *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 Click **Start** → **Settings** → **Control Panel**. The Control Panel window opens.
- 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 **Start** → **Programs** → **Sybase** → **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:

- *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 `sybsystemprocs`) 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 93.)

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 `sybsystemprocs` 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**, **sybsystemprocs**, 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 sybsystemprocs 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 show 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 `syssservers` 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 Click **Start** → **Settings** → **Control Panel**. The Control Panel window opens.
- 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 **Stop**.

Backing Up the Sybase SQL Server Using a Script

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 97.)



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 sybssystemprocs, to a file system 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/\$SYBASE_ASE/installerror.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 [Creating a Backup Server on a UNIX System](#) on page 88 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 <PI-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/
$SYBASE_ASE/install/error.log
date | tee -a $SYBASE/$SYBASE_ASE/install/error.log

# Start isql session to perform the backups
#
$SYBASE/$SYBASE_OCS/bin/bin/isql -Usa -P << EOFEOF | tee -a
$SYBASE/$SYBASE_ASE/install/error.log
checkpoint
go
dump database master to ${TAPE_DEVICE} file="master.dmp" with
init
go
dump database sybssystemprocs to ${TAPE_DEVICE}
file="sybssystemprocs.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/  
$SYBASE_ASE/install/error.log  
date | tee -a $SYBASE/$SYBASE_ASE/install/error.log  
echo " " | tee -a $SYBASE/$SYBASE_ASE/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 PI installation and database when the following occur:

- After your initial PI installation

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

- 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 PI 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 102 and [Backing Up Your Oracle Database](#) on page 87 or [Backing Up Your Sybase Database](#) on page 87 for details.
- Starting and stopping PI processes
- Relocating the PI 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 HP Performance Insight Processes

PI has the following processes:

- OVPI HTTP Server

This process controls the PI client's access to the PI 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 PI 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 565 for details.

- OVPI Timer

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

See the *HP Performance Insight Reference Guide* for more information about the `trendtimer` command.

Starting and Stopping Processes on Windows

When you installed PI, OVPI HTTP Server and OVPI Timer are set to start automatically as services. If they do not start 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:


- 1 Click **Start** → **Settings** → **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 104.

- 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 Click **Start** → **Settings** → **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 PI, 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 logged in as `root`.
- 3 Start the OVPI HTTP Server process by entering one of the following commands:
 - HP-UX:
`/sbin/init.d/ovpi_httpd start`
 - Linux:
`/etc/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`
 - Linux:
`/etc/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 HP Performance Insight Sybase Database

If you need to move the PI Sybase database from one host to another, create and use Sybase backup servers on both Host A (the source host) and Host B (the target host).

To relocate the PI 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 PI 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 the latest backup by starting an `isql` session and entering the following command:

```
isql -Usa -P  
1> dump database dpipe_db to 'path-to-diskspace'  
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.

Confirm that the database on Host B is complete, you can turn on **trendtimer** on Host B.

Verifying the Relocation of the HP Performance Insight Database

After you relocate the database, check the following items on Host B:

- Polling requests

Using Polling Policy Manager, 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 8, 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 were 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_dpipe -Pdsi_dpipe
1> update tl_groups set dsquery = Sybase-DSQUERY-value,
2> data_db = Sybase-DSQUERY-value
3> go
```

- 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 (.rep, .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 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 HP Performance Insight in an HTTP Proxy or Portal Environment

If you have an HTTP proxy server located between the PI's Web Access Server component and your real server, you may need to configure the appropriate authentication for PI's Web Access Server component after you install PI. To determine what kind of authentication is best, you need to know the username translation capabilities of your proxy server.

Ideally, it is best to use a proxy sever that can do either of the following:

- Take the inbound proxy or portal username and pass it directly with the password to PI using proper HTTP header encoding.

- Use some sort of translational methodology to convert the inbound name to a PI account name.



It is helpful to distinguish between portals and proxies. Proxies merely pass HTTP traffic from one system to another (largely for security compartmentalization) with traffic enforcement and logging potential. Portals serve a broader purpose of application integration, and offer many additional features and capabilities.

When an HTTP-based application is generating packets, it creates the HTTP header. If the proxy server does not have the ability to specify the username and password in the header (for example, because it relies on a sublayer to create the HTTP header), it may not be powerful enough to take a proxy username, like *igeorges*, and translate or even pass it in clear text to the PI's Web Access Server component using HTTP 1.1 basic authentication. To integrate with Service Information Portal (SIP), which relies on Tomcat to build its HTTP headers, you cannot use HTTP basic authentication. PI provides other authentication methods to pass a username and password to PI. These options are as follows:

- HTTP Basic. This method requires that you log on to the Web Access Server by supplying the username and password for your user account.
- URL. This method requires that you supply a username when you view reports using a URL.
- None. This method requires no authentication.



You cannot use both HTTP Basic and URL authentication at the same time. For example, suppose you want users in your Network Operations Center to have their web browsers started from Network Node Manager (NNM) and be able to connect to PI without having to log on; however, you want to retain standard PI security for every other user account.

You can configure this scenario by using two Web Access Server systems. One server can run with the HTTP Basic authentication option, while the other server can run with the None authentication option. In this case, it is necessary for the PI Administrator to redeploy reports to the second Web Access Server. This can be done using the **deploytool** utility. See the *HP Performance Insight Reference Guide* for information about using this utility.

To set up authentication for the Web Access Server:

- 1 Start your Web browser.
- 2 Type the URL in the following format with Address or Location field of your browser, and 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 type 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 Type your Administrator username and password, and then click **OK**.

PI creates a default Administrator account called **trendadm** when it is installed. Use this account or use another account that has Administrator privileges.

The Web Access Server opens.

- 5 From the Links bar, select **Administration**. The Administration page opens.
- 6 In the Administration tree, expand Service Configuration and then click **Web Access Server**. The Web Access Server Settings page opens.
- 7 In the Realm box, type a realm name for the server. It can contain up to 50 characters.

A **realm** is a database of usernames and passwords that identify valid users of a web application (or set of web applications) and provides an enumerated list of the roles associated with each valid user. Roles are similar to groups in UNIX operating systems, because access to specific web application resources is granted to all users possessing a particular role rather than enumerating the list of associated usernames. A

particular user can have any number of roles associated with his or her username. Typically, you only need to change the realm value in response to the request of a portal admin.

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 **Log on** on the Web Access Server.

- 8 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. The browser accepts the user input and adds the information to every subsequent HTTP header for the session to that server. The username and password are Base64 encoded.
 - **URL.** This method requires that you supply a username when you view reports using a URL.

This option is for integration with applications that either do not have the ability to encode the username in an HTTP header (a relatively primitive proxy server like SIP) or which do not have the concept of usernames (for example, Network Node Manager) where you just want to send a browser request for some generic network operation center account to the PI server and allow PI to do the catalog-view masking.

There is no password required with this option. Since arguments are passed using clear text on the command line (the URL), it does not make sense to send a password in clear text. Therefore, when you use URL authentication, it is assumed that the PI server is running in a trusted environment, because users will not be prompted for passwords, and anyone with a web browser can access any users data by changing the URL value of the `&ov_user` argument.

For more information about using URLs to view reports, see the *HP Performance Insight Guide to Building and Viewing Reports*.

- **None.** This method requires no authentication. Users will access the system with the same privileges as the user account specified in the Account box. This option does not prompt the user for log on information. To achieve a basic level of filtering and access control, the PI server allows you to specify a default PI account with the None log on using the Account ID 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 568.

- 9 In the Account ID box, type the name of a user account to associate with the None authentication option. All other user accounts will acquire the same privileges as this account when accessing the Web server.
- 10 Click **Apply**.
- 11 When you have finished your changes, stop and start the Web Access Server. See [Starting and Stopping HP Performance Insight Processes](#) on page 103 for this procedure.

Example: Setting Up an Apache Proxy Server

To set up an Apache proxy server, you will need to add some lines to the `httpd.conf` file. The `httpd.conf` file is located in the `conf` subdirectory of the directory into which you installed Apache.

To set up the server, 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.

There is no need to change the authentication in this case, because basic HTTP authentication is the default.

5 Using the PI Management Console

The PI Management Console lets you manage most of the features of your PI system. The various snapins in the PI Management Console are intuitive and provide ease of use to customers. PI provides another application for managing various features of the product called the Web Access Server.

You can use the Management Console to perform configuration, administration and monitoring of the Performance Insight Reporting platform.

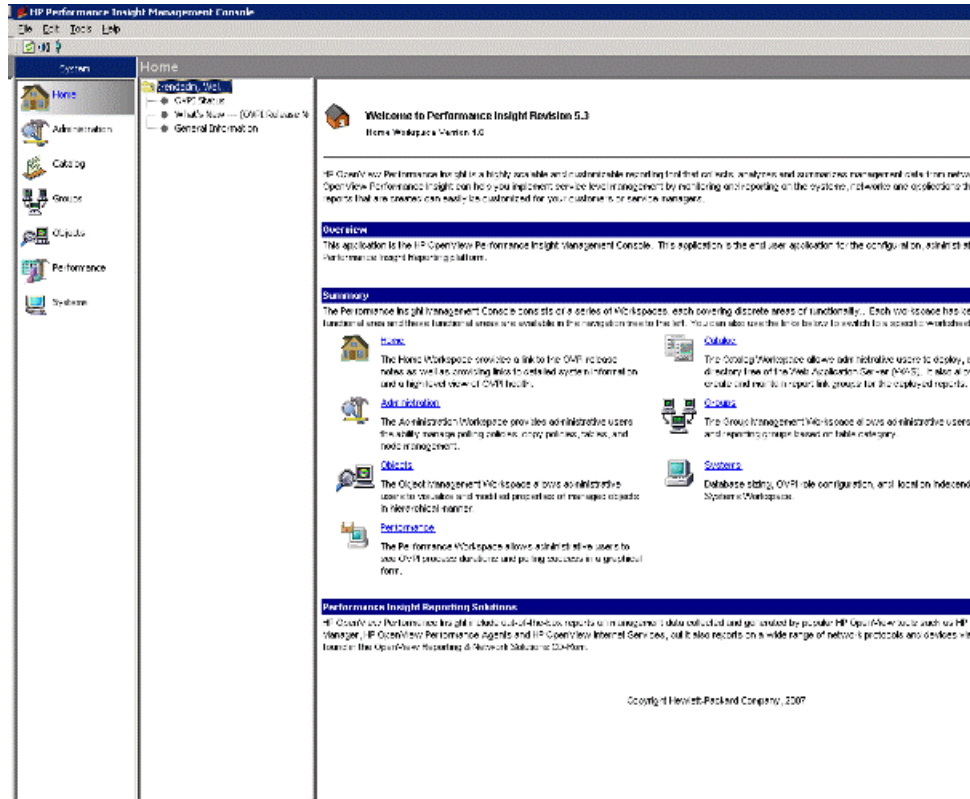
The Management Console consists of a series of workspaces, each covering discrete areas of functionality. Each workspace has been further sub-divided by functional area and these functional areas are available in the navigation tree to the left. Also you can use the given links on the right to switch to a specific worksheet.

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

- **Home:** The Home snapin provides a link to the PI release notes and detailed system information and a high level view of PI health.
- **Administration:** This function lets you add/delete/modify nodes, polling policies, and copy policies. It also lets you view the tables used in PI via the table management functionality. It lets you perform the following functions:
 - **Node Management:** You can manage nodes, SNMP Profiles and Community String Profiles.
 - **Tables:** You can view information about various types of tables in the system.
 - **Polling Policy:** Administrative users can create and maintain how, what, and frequency of data collection.

- **Copy Policy:** Administrative users can create, modify and delete copy policies on the central server to control the behavior of data copy functionality in PI. For more information, see [Chapter 17, Location Independent Reporting and Copy Policies](#).
- **Catalog:** The Catalog/Report Link Management function allows administrative users to deploy, delete, and move reports in the directory tree of the Web Application Server (WAS). It also allows the administrative user to create and maintain report link groups for the deployed reports.
- **Groups:** The Group Management function allows administrative users to create and maintain polling and reporting groups based on table category.
- **Objects:** The Object/Property Management function allows administrative users to visualize and modify properties of managed objects in hierarchical manner.
- **Performance:** The PI Performance Graphs function allows administrative users to see PI graphs for Process Durations, Polling Health and Row Counts. For more information, see [Chapter 9, Internal Monitoring of PI using Performance Graphs](#).
- **Systems:** The System/Network Administration function lets you view the System information related to Databases, Web Access Servers and Remote Pollers.

The following figure shows the HP Performance Insight Management Console window



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, see the *HP 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
```

See the *HP 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:

- Click **Start** → **Programs** → **HP Software** → **Performance Insight** → **Management Console**.
- From a Windows Command Prompt window, type the following command at the MS-DOS command prompt, and then press **ENTER**:

```
piadmin
```

See the *HP 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 PI.

- b Double-click `piadmin.exe`.

Logging On

To log on to the Management Console:

- 1 Start the Management Console.
The Login dialog box opens.
- 2 Type your username, your password, and then select the system on which you want to use the application.
- 3 Click **Login**.

The Management Console Interface



The figure on [page 117](#) shows the HP Performance Insight Management Console window.

The HP Performance Insight Management Console comprises a menu bar, a toolbar, a system pane, a navigation pane and a viewing pane.

After you open any of the applications available from the Management pane or open the MIB Browser, which is available from the Tools menu, the application window displays 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 Management Console has the following menus: File, Edit, Tools and Help.



The menu bar changes according to the application that is opened.

File Menu

Table 2 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 figure on page 127 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 3 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 4 Management Console Tools Menu Options

Menu Option	Description
Database Tuner	The Database Tuner opens the Install Wizard for Database Configuration.
CertificateMgmt UI	The Certificate Management UI is used when you collect data from HP Agents that support HTTPS protocol. Lets you view and manage the security certificates.
NNM-PI Integration Wizard	Opens the NNM-PI Integration Wizard. You can use the wizard to add, modify, or delete the details of the NNM servers you want to integrate with PI. See the <i>HP Network Node Manager and HP Performance Insight Integration User Guide</i> for more information.

Table 4 Management Console Tools Menu Options

Menu Option	Description
MIB Browser	Displays the MIB Browser application, which lets you analyze, query, and troubleshoot the 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.



In a client-only installation, you cannot use the applications that are normally available from the Tools menu following a complete PI installation. You cannot register a new data table, property table, or automatically add a collection to the database.

Help Menu




Table 5 Management Console Help Menu Options

Menu Option	Description
Help Topics	Provides access to the PI online documentation.
About Performance Insight...	Displays version and copyright information for PI.

Toolbar

The buttons that appear on the toolbar vary according to the application you are using.

Table 6 Management Console Toolbar Buttons

Button	Menu Option	Description
	Refresh	Refreshes the display.
	Change Server	Displays the Select Admin Server dialog box, which lets you select another system to manage.
	Logout/ Login	Lets you log off or to log on the Management Console.

The Home Snapin

The Home snapin provides easy access to PI applications and lets you view the status of PI services and polling. It also shows the Connection and System information and lets you read the release notes. The monitoring capability simplifies your task of maintaining PI systems. The systems are constantly monitored to record performance data and trace potential problems, thereby eliminating the need to look through voluminous log files for information about the health of the systems.

The Home snapin provides the following options in the navigation pane:

- **OVPI Status:** Provides basic statistical information about the selected system. For more information, see [Viewing HP Performance Insight Status](#) on page 125.
- **What's New...(OVPI Release Notes):** Brings up the PI Release Notes document. For more information, see [Viewing HP Performance Insight Release Notes](#) on page 127.
- **General Information:** Provides PI Connection and System information system information. For more information, see [Viewing HP Performance Insight General Information](#) on page 127.

Viewing HP Performance Insight Status

The PI Status option helps you understand the following aspects of the PI environment:

- PI System Services and System Configuration
- PI Polling Status

Opening the HP Performance Insight Status Page

- 1 In the navigation tree provided in the Home snapin, click the **OVPI Status** leaf node.

The OVPI Status Page opens and contains data for the system that you have logged into.

The OVPI Status Page contains the following information:

- **OVPI Status:** provides basic statistics about the selected system.

- **Services:** Lists all possible services related to PI and their current status.
 - PI services that are available are marked with a green check mark.
 - PI services that are down are marked with a red cross symbol.
 - If a service is not applicable to PI, the service name is grayed out and the status is indicated as *not installed*.



You cannot start or stop a service from the OVPI Status page.

- **Configuration:** Provides a limited view of key data storage areas.

The statistics are determined in the following manner:

- For the collect home process, the graph displays the current fill percentage (%) based on used space/total space.
- For the database, if one tablespace is filled above 90%, the status is marked as High. If a table space fill is above 80%, the label is Warning. For all other conditions, the status is marked as Acceptable. Database fill statistics can be queried from the system's `dsi_segment_v` table instead of from the RDBMS system. This table should be used to provide a consistent view of how PI sees a full database.

If you click the **More** button, you are taken to Individual Tablespace Fill Statistics page.

- **Polling Status:** Displays PI polling statistics for the related system. The data is collected from the CSV files that are generated by `mw_collect`. The data is rendered in the form of tables for three intervals: current hour, last hour, and yesterday.
 - **Percentage of Node Connectivity:** This displays the percentage of nodes that PI is connected to. Errors in this value indicate downed nodes or agents and invalid community strings. These are basically due to any non-zero exit status.
 - **Percentage of Nodes Collected:** This displays the percentage of nodes that PI successfully connected to and received data from. Nodes that returned zero rows but exited with a zero exit status would subtract from 100% success in this case.

- **Rows Collected:** This displays the total rows collected for all polling policies.
- **Number of ERRORS from child collector:** This displays the summary of all collection errors.
- **Number of Processes Exceeding the Interval:** This displays the number of collections that exceeded the collection interval. This is a query against the `mw_collect` entries.

Viewing HP Performance Insight Release Notes

The OVPI Release Notes section contains the following information:

- New features introduced in this version
- Installation Notes
- Software Requirements, Documentation Updates, Known Problems, Limitations, and Workarounds
- Support
- Legal Notices

To open the Release Notes, click the **What's New ... (OVPI Release Notes)** leaf node.

Viewing HP Performance Insight General Information

The General Information option on the navigation tree in the Home snapin has a snapshot view of various system parameters. This includes the following:

- **Connection Information:** This gives the name of the system and the database to which you are connected.
- **System Information:** This gives information about the Operating System, System load and PI Application related to the selected system.

To open the General Information option, click the General Information leaf node in the navigation tree.

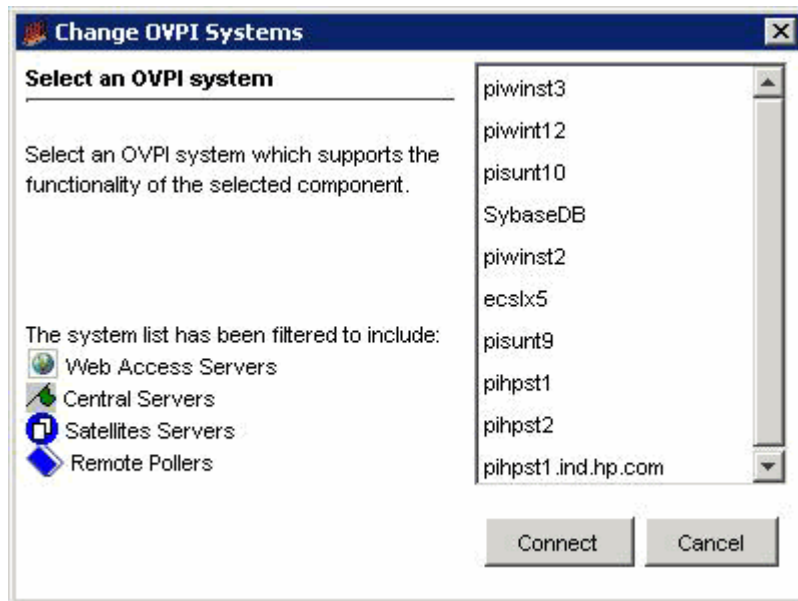
Changing Database Servers

PI can read and store data on any PI 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.

To change the server:

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

The Change OVPI Systems dialog box opens, which displays the available PI servers.



- 2 Select a server name, and then click **Connect**.

The Management Console displays information from the database of the server that you selected.

To verify that you are accessing the new database server, from the Home Snapin icon, click the **General Information** option on the navigation tree. The Connection Information pane identifies the current database.

Saving User Preferences

PI lets you maintain a unique set of user preferences for the following applications:

- Report Builder
- Report Viewer
- Management Console

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

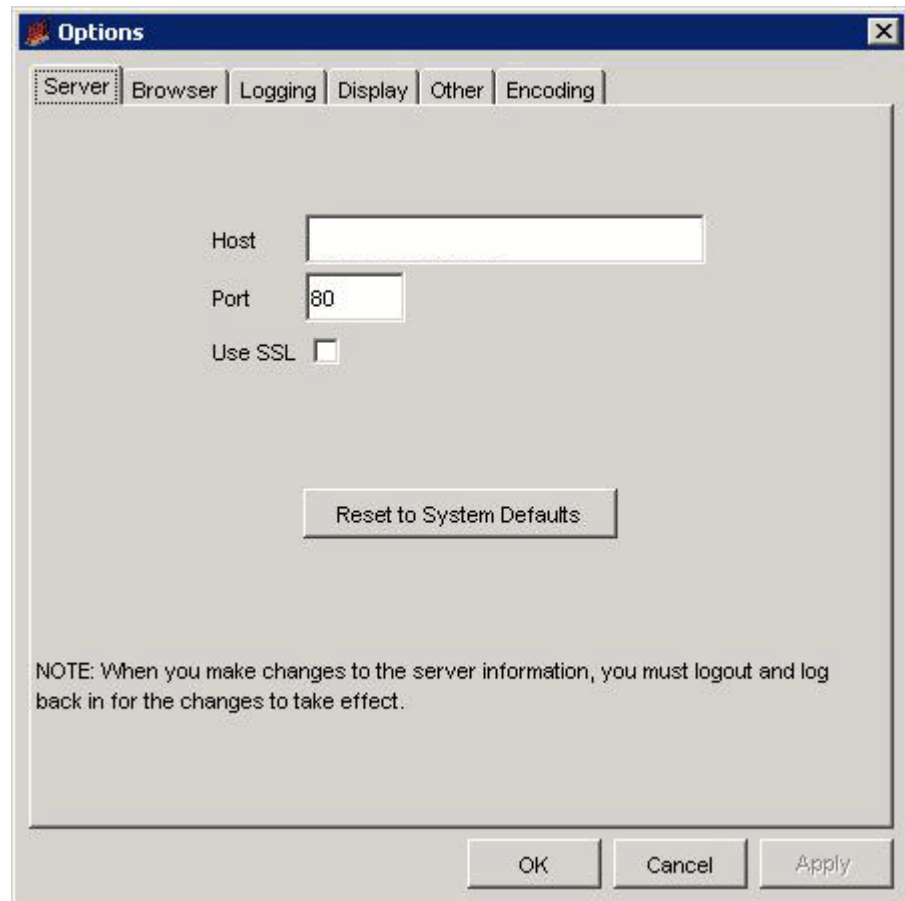
For example, you can specify a directory in which PI should save Management Console log files, and that setting will be saved in your `.ini` file for the Management Console application. Your `.ini` files are stored in your home directory, using the following naming convention:

```
.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 **Server** 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 PI. 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 PI.


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:

- Type 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.

6 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 Performance Insight (PI) 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 *HP Performance Insight Guide to Building and Viewing Reports*.

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 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. 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 transforms the raw data into a form that can be used by report packs.

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

Report Pack and Datapipeline 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. 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. 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

Windows Desktop

On the Windows taskbar, click **Start**, point to Programs, HP Software, Performance Insight, and select **Package Manager**. The Welcome window opens..

HP Performance Insight 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 Software, Performance Insight, and select **Management Console**.
The Management Console window opens.
 - b From the Tools menu, select **Package Manager**.
The Welcome window opens (see figure on [page 139](#)).
- UNIX
 - a Log on as trendadm.
 - b Type the following command:
piadmin
The Management Console window opens.
 - c From the Tools menu, select **Package Manager**.
The Welcome window opens (see figure on [page 139](#)).

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 opens (see figure on [page 139](#)).

- UNIX

a Log on as trendadm.

b Type the following at the command line:

packagemgr

The Welcome window opens (see figure on [page 139](#)).

Installing Packages

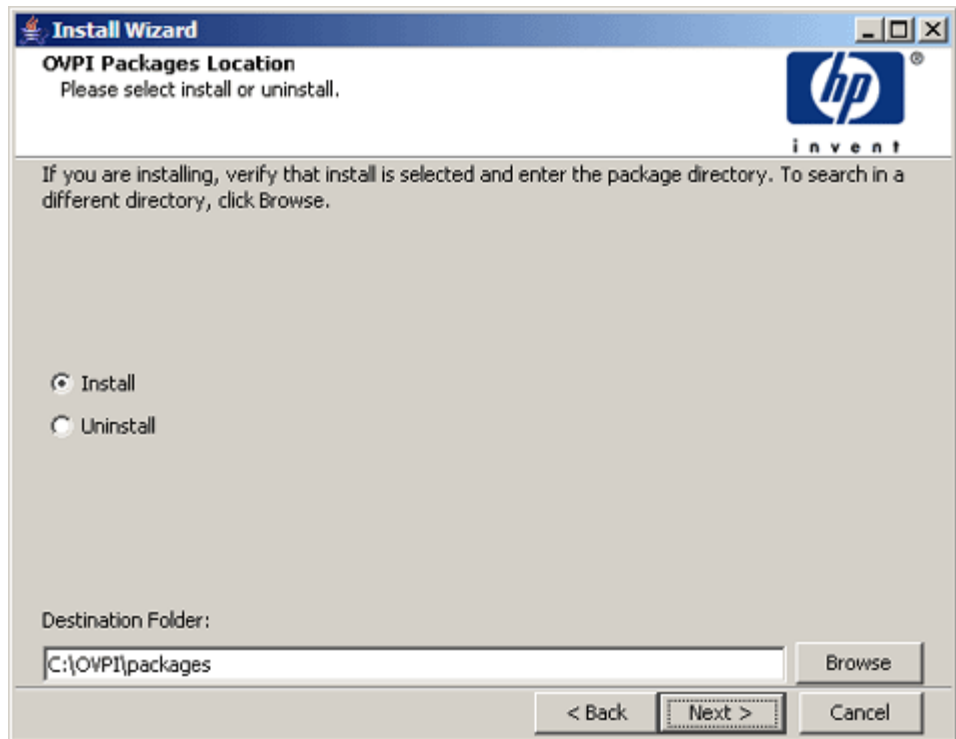
To install packages:

- 1 Start Package Manager. (See [Starting Package Manager](#) on page 137).

The Welcome window opens.



- 2 Click **Next**. The OVPI Packages Location window opens.




3 Follow these steps:

- a Select **Install**, and enter the directory in which you saved the package or packages.

This should be the `packages` folder of the directory into which you installed PI. For example, on a Windows system, this directory might be `C:\OVPI\packages`.

- b Click **Next**. The OVPI Report Deployment window opens.



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. See the *HP Performance Insight Guide to Building and Viewing Reports* for details.


- 4 Follow these steps:
 - a Select the **Deploy Reports** check box to deploy the reports to the Web Access Server and then follow these steps:

 You must 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.

- 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.

b Select the **Use SSL** check box if the Web Access Server is running on HTTPS. This is mandatory if SSL is enabled, otherwise you will not be able to proceed with the installation. When you select the check box, the port number for OVPI Application Server automatically changes to 443.

PI 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 525.

 If your configuration has multiple Web Access Server systems, use the `deploytool` command after you have installed the packages on your Central Server to deploy reports to the other Web Access Server systems. See the *HP Performance Insight Reference Guide* for more information about the `deploytool` command.

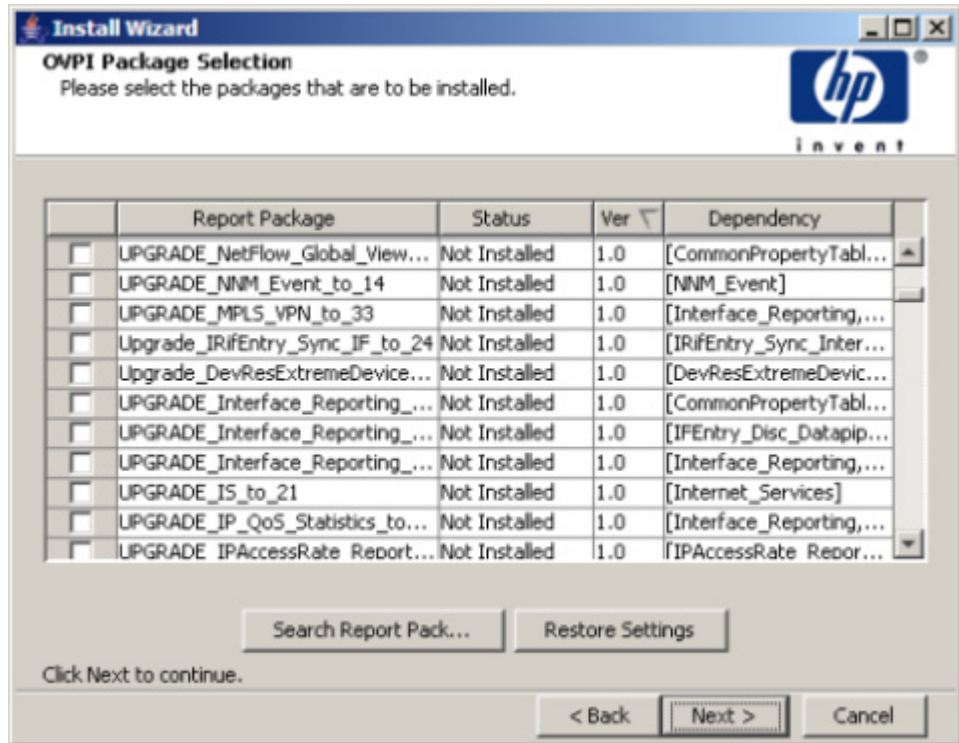
Alternatively, you can use the Deployment wizard to deploy reports one at a time to these systems. See the *HP Performance Insight Guide to Building and Viewing Reports* for more information about the Deployment wizard.

5 Click **Next**.



If the packages you are installing cannot be installed on the current database, a message will appear, 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.



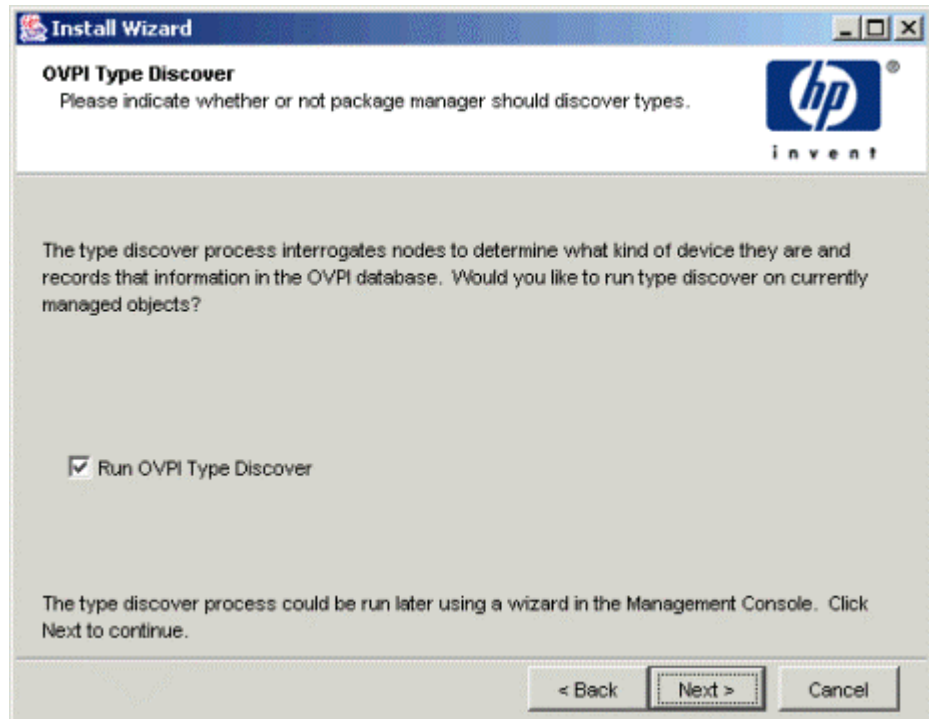
This window lists the following:

- **Report Package:** Lists the report pack packages and their upgrades (if any) that are currently located in the packages folder.

Click **Report Package** in the heading bar to sort the listed reports in alphabetical order. By default the report packs and their upgrades are listed together.

- **Status:** Lists whether or not the report pack package is installed. To sort the report packs based on their status click **Status** in the heading bar.
 - **Version:** Lists the version number of the report pack package. To list the versions in ascending or descending order, click **Ver** in the heading bar.
 - **Dependency:** Lists 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.
 - **Search Report Pack:** Click **Search Report Pack** to search for a particular report pack from the listed report packs.
 - **Restore Settings:** Click **Restore Settings** to reset the report packs in the default order.
- 6 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.
- 7 Click **Next**.

The OVPI Type Discover window opens.



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. If you know which types you have in your environment, you may want to run the type discovery process after you install the packages.

▶ If the database has less than 10% of free space, a warning message appears. You can ignore this message and continue installation.

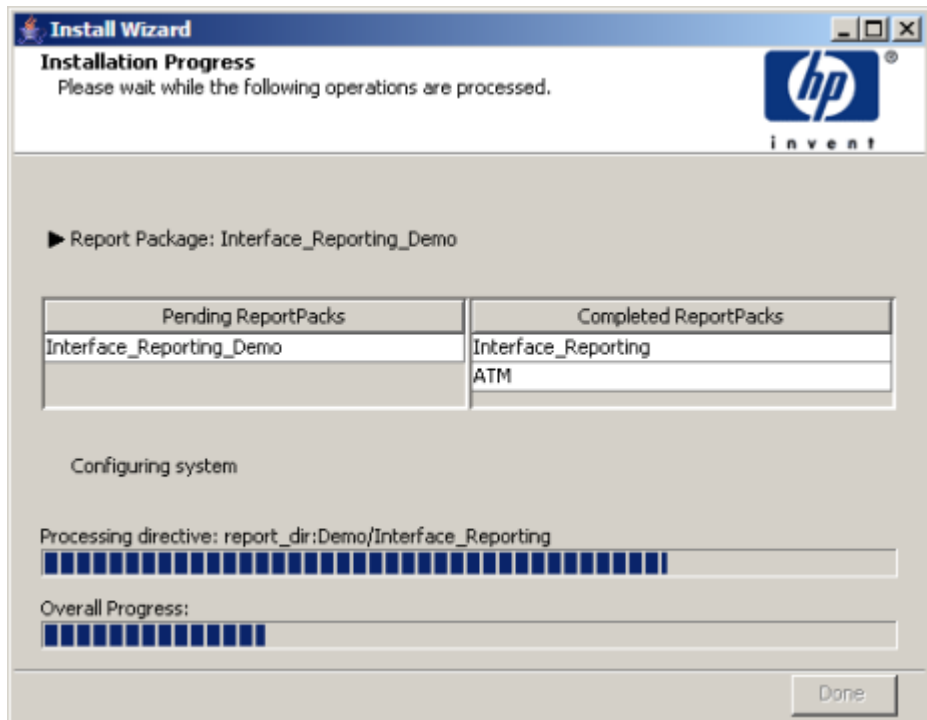
- 8 Do *one* of the following:
 - α To run Type Discovery after the package installation, click **Next**.

- 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 152), and then run type discovery (page 159).

The Selection Summary window opens

- 9 Verify the information listed in the Selection Summary window, and 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 **Back** to return to any of the previous screens and make corrections. When finished, return to the Selection Summary window, and click **Install**. The Installation Progress window opens.



Under the Pending ReportPacks column, selected report packs, yet to be installed, are listed. The installed report packs are listed under the Completed ReportPacks column.

If the installation of a package fails, an error message appears.

To continue with the installation, do one of the following:

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

Select the **Don't ask me again** check box if you do not want any error message to appear for further failures. After the installation completes, click Details to check for any error messages.

- 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.

10 Click **Done**.

Uninstalling Packages

To uninstall packages:

- 1 Log on to the system as the PI Administrator.
- 2 Start Package Manager (see [Starting Package Manager](#) on page 137). The Welcome window opens (see figure on [page 139](#)).
- 3 Click **Next**.

The OVPI Packages Location window opens (see figure on [page 140](#)).

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

The OVPI Report Undeployment window opens.

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.
 - Type your Administrator user name and password. You must have an account with administrator privileges to undeploy reports.
 - PI 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 Packages Selection window opens.

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.

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 **Back** 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, showing the individual and overall progress of the uninstall.

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

7 Discovering the Network

This chapter describes how to retrieve information about the nodes on your network. Specifically, it describes how to use the PI 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 [The Discovery Process](#) on page 152 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 159 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. It also lets you control additional features of the discovery process by specifying different options.

See the *HP 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.

▶ The Discovery and Type wizards, the `trend_discover` command, and the `node_manager` command are only available on a PI server system.

If you have systems of the Network Node Manager (NNM) and want to add nodes from NNM to your system, you can use the NNM-PI Integration wizard. See the *HP Network Node Manager and Performance Insight Integration User Guide* for more information about this wizard.

The Discovery Process

The discovery process pings a range of IP addresses, one at a time, and analyzes and identifies the responses it gets. It uses that information to populate or update tables in the PI 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. Use the `trend_discover -1` 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 pings each IP address (except beginning subnet addresses [0] and broadcast addresses [255]) in that IP address range. One of the following occurs:

- If there is no device at the IP address, there is no response.
- If there is a device at the 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 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 118 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 525.

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

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

SNMP Discovery Wizard

Enter the IP address range and subnet mask for the network to be discovered.

Starting IP Address
16.116.4.0

Ending IP Address
16.116.4.255

Subnet Mask
255.255.255.0

Click Next to Continue.

< Back Next > Cancel

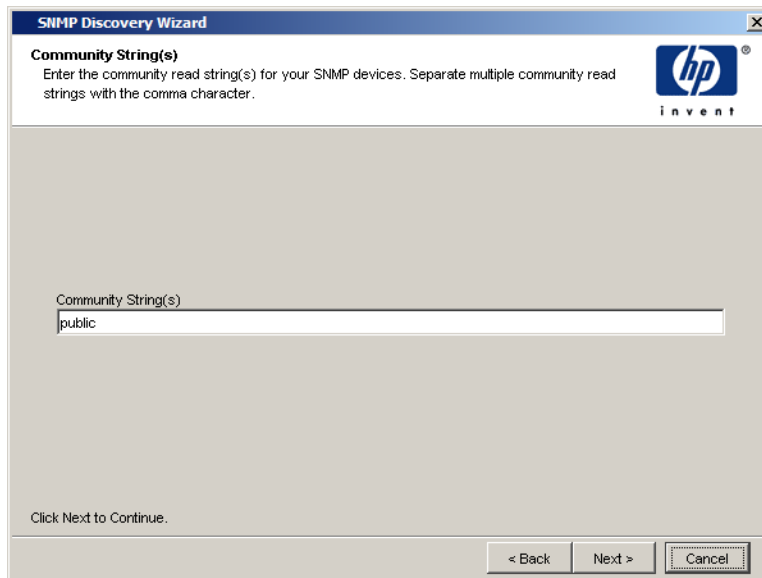
- 3 Do the following:
 - a Specify the minimum value for the IP address in the Starting IP Address box.

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

The following window opens.

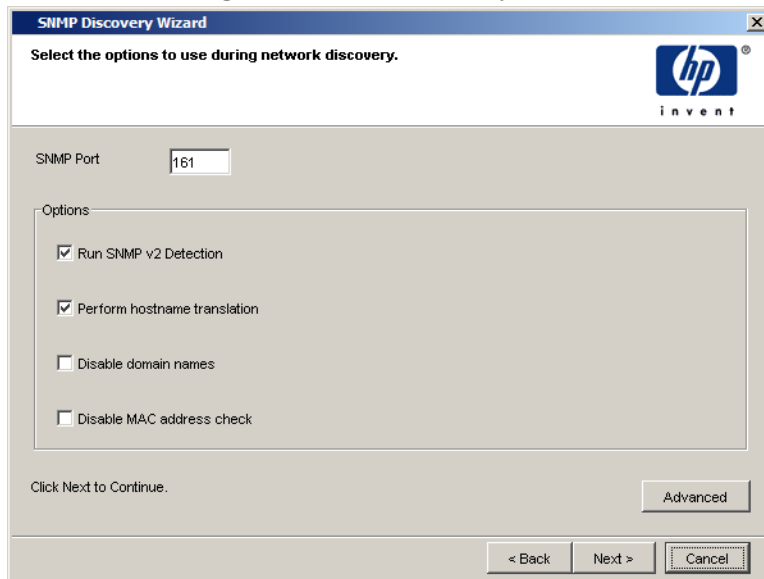


When the SNMP Discovery wizard finds a device at an IP address, it sends an SNMP GET to determine whether or not that device is SNMP manageable. This SNMP GET includes a read community string. The targeted device will respond to the SNMP Discovery wizard only if its community string matches the string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

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

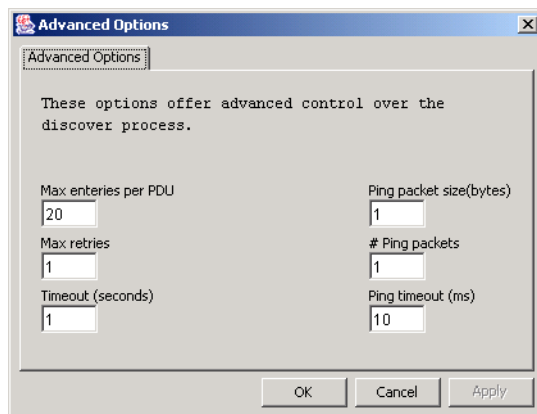
Host names are typically textual names that are used to identify devices. When you select this option, the wizard attempts to locate the host name for the IP address of each node that it finds.

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.

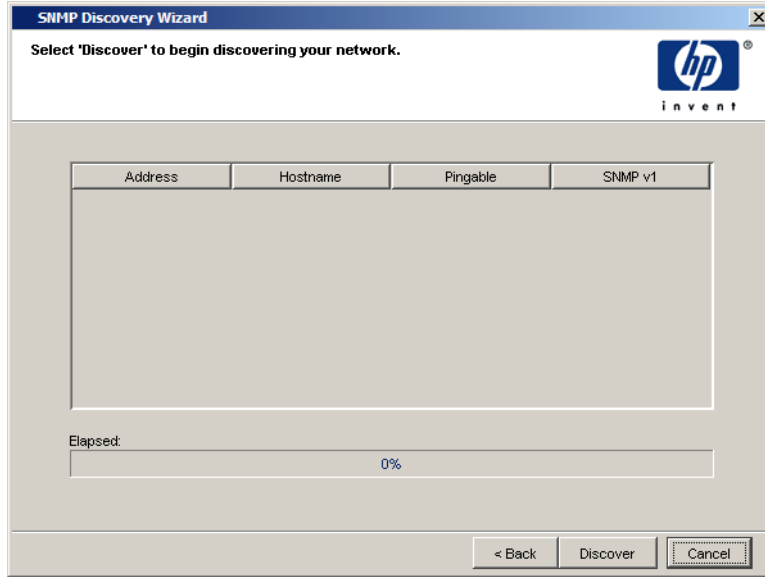
- e Select or click to clear the **Disable MAC address** check box.
- f Click **Advanced**. The Advanced Options dialog box opens, which lets you specify additional options for the discovery process.



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 SNMP Discovery Wizard 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 records this information in the database. This information lets PI 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 (see below).
- Package Manager. Start the type discovery process during report pack installation. See [step 8](#) on page 145 for details.
- Command line. Use the `trend_discover -t` 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.

The Type Discovery Wizard

To start the Type Discovery Wizard:

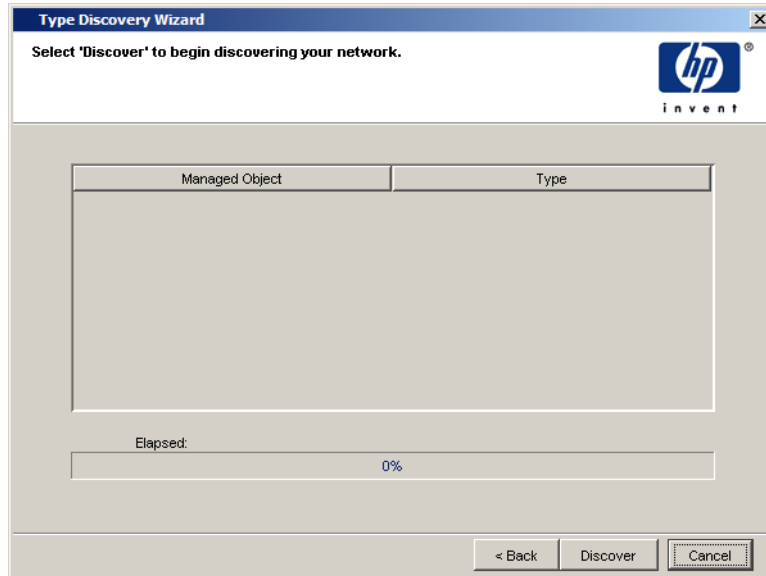
- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 for details.
You must log on to the Management Console using an account that has Administrator privileges. For more information about user accounts, see [Maintaining User Accounts](#) on page 525.
- 2 From the Tools menu, select **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.

8 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 171](#)).
- Create new collections ([page 173](#)).
- Create new polling policies ([page 252](#)).
- Edit existing polling policies ([page 257](#)).
- Delete existing polling policies ([page 259](#)).
- Import nodes ([page 275](#)).
- Manage polling groups ([page 260](#)).
- Manage nodes ([page 271](#)).
- Manage community string profiles ([page 279](#)).
- Manage SNMP profiles ([page 282](#)).
- Manage Datapipe installations ([page 285](#)).
- Modify property table values ([page 288](#)).

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 datapipe packages using Package Manager, polling strategies are provided but you might want to change them to suit your collection needs.

For example, if the Device Resource Cisco Router Datapipe is installed, it configures three polling policies that collect every 15 minutes and one polling policy that collects every hour. These policies will collect from the group of Cisco routers. However, if you want to collect at 5 minute intervals instead of at 15 minute intervals and only run the hourly collection once a day, you should modify the default policies. You might also only want to collect for a subset of the Cisco routers on each of your satellite systems; thus, you would want to modify the grouping too.

Some datapipes allow you to configure directed instance polling. Directed instance polling allows data to only be collected from a subset of interfaces. See [Directed Instance Polling](#) on page 166 for details.

Polling Groups

PI 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 affect 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 PI 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. PI polling is aimed at providing a baseline to support the troubleshooting activity and at identifying patterns generally occurring over hours, days, and weeks.

Directed Instance Polling

You may not want to poll all the components for a particular device. This situation is common when in the case of interfaces.

For example, suppose you have a switch with 100 interfaces, which are all connected to different PCs and servers. If you were to collect data from all of these interfaces using standard SNMP polling, you would collect a row of data for every interface, which would result in 100 rows of data.

In many cases, you may not be interested in all the data for each interface. For example, perhaps you are not concerned about the drops to individual PCs. You are only interested in the connection to the parent switch or the connections that feed a server. In this case, you may only want to poll four out

of 100 connections. Directed-instance polling allows data to be collected from a subset of interfaces. The subset is identified either by rules or by specific provisioning.

The advantage of directed instance polling is that it lets you manage a bigger network, because you are not collecting data from every individual interface when you do not care about most of them anyway. This also reduces the collection load and the network traffic. The only disadvantage is that you do not have the data for every interface; however, if you have decided to use directed instance polling, you have already determined that you do not want all of the data.

The Interface Reporting and FrameRelay Service report packs have the following datapipe that you can use to set up directed instance polling:

- ATM ifEntry datapipe. See the *ATM ifEntry Datapipe User Guide* for details.
- Frame Relay CPE datapipe. See the *Frame Relay CPE Datapipe User Guide* for details.
- Interface Reporting ifEntry datapipe. See the *Interface Reporting ifEntry Datapipe User Guide* for details.

Starting Polling Policy Manager

To start Polling Policy Manager:

- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 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 525.

- 2 Click the **Administration** icon in the navigation pane. The Administration snapin opens.
- 3 In the navigation tree provided in the Administration snapin, click **Polling Policy**.

The Polling Policy Management window opens.

The screenshot shows the HP OpenView Performance Insight Management Console. The Administration snapin is active, and the Polling Policy tree is expanded. The main window displays a table of polling policies with the following columns: Name, Datatype, Collector, Table Alias Name, Group Name, Poll From, Server, Interval, User, and Description.

Name	Datatype	Collector	Table Alias Name	Group Name	Poll From	Server	Interval	User	Description
Flagged_I...	dpipe_s...	MW	Raw Interface by I...	discover-by-If...	piwint7	piwint7	15 Mnu...	iren...	Discover inte...
Flagged_J...	dpipe_s...	MW	Raw Interface by I...	discover-by-If...	piwint7	piwint7	15 Mnu...	iren...	Discover inte...
Flagged_J...	dpipe_s...	MW	Raw Interface by I...	discover-by-If...	piwint7	piwint7	15 Mnu...	iren...	Discover inte...
Flagged_J...	dpipe_s...	MW	Raw Interface by I...	discover-by-If...	piwint7	piwint7	15 Mnu...	iren...	Discover inte...
2007-03-...	dpipe_s...	MW	xDPVpnInterfaceP...	MPLS_VPN	piwint7	piwint7	15 Mnu...	iren...	
2007-03-...	dpipe_s...	MW	xDPVpnDeviceStats	MPLS_VPN	piwint7	piwint7	15 Mnu...	iren...	
2007-03-...	dpipe_s...	MW	xDPVpnInterfaceC...	MPLS_LSR	piwint7	piwint7	1 Hour	iren...	
2007-03-...	dpipe_s...	MW	xDPVpnInterfaceC...	MPLS_VPN	piwint7	piwint7	1 Hour	iren...	
SysReso...	dsi_sr	SR	SysResOVPA_CPU	OVPA	piwint7	piwint7	1 Hour	iren...	The CPU rela...
AI_IfInde...	dpipe_s...	MW	Raw Interface by I...	index-by-IfIndex	piwint7	piwint7	1 Day	iren...	Discover inte...
AI_IfAlias...	dpipe_s...	MW	Raw Interface by I...	index-by-IfAlias	piwint7	piwint7	1 Day	iren...	Discover inte...
AI_IfDesc...	dpipe_s...	MW	Raw Interface by I...	index-by-IfDescr	piwint7	piwint7	1 Day	iren...	Discover inte...
AI_IfName...	dpipe_s...	MW	Raw Interface by I...	index-by-IfName	piwint7	piwint7	1 Day	iren...	Discover inte...
Device Sy...	dpipe_s...	MW	Raw MB-II System...	GENMB2F_Ds...	piwint7	piwint7	1 Day	iren...	System grou...
ATM_IfAl...	dpipe_s...	MW	V_V_ATMEntry_IfA...	IfAlias-ATM	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
ATM_IfDe...	dpipe_s...	MW	V_V_ATMEntry_IfD...	IfDescr-ATM	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
ATM_IfIn...	dpipe_s...	MW	V_V_ATMEntry_IfI...	IfIndex-ATM	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
ATM_IfNa...	dpipe_s...	MW	V_V_ATMEntry_IfN...	IfName-ATM	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
IfAlias_Co...	dpipe_s...	MW	Raw IfAlias Data V...	IfAlias-all	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
IfDescr_C...	dpipe_s...	MW	Raw IfDescr Data ...	IfDescr-all	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
IfIndex_C...	dpipe_s...	MW	Raw IfIndex Data ...	IfIndex-all	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...
IfName_C...	dpipe_s...	MW	Raw IfName Data ...	IfName-all	piwint7	piwint7	15 Mnu...	iren...	Collect IfEntr...

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 are 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 .
- Select **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 various filters. You can use the following filters to view only a subset of the available polling policies:

- **User.** View polling policies filtered by username.
- **Interval.** View polling policies filtered by polling interval.
- **Group Name.** View polling policies filtered by group name.
- **Poll From.** View polling policies filtered by the machine that is doing the polling.
- **Datapipe.** View polling policies filtered 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:

The screenshot shows the HP OpenView Performance Insight Management Console. The left sidebar contains navigation options: Home, Administration, Catalog, Groups, Objects, Performance, and Systems. The main window is titled 'Administration' and features a tree view on the left with 'Polling Policy' selected. The main area displays a table of polling policies with filters for User, Interval, Group Name, and Poll From. The table columns are Name, Datapipe, Collector, Table Alias Name, Group Name, Poll From, Server, Interval, User, and Description.

Name	Datapipe	Collector	Table Alias Name	Group Name	Poll From	Server	Interval	User	Description
Flagged_..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	15 Mn...	trendadm	Discover info...
All_Inds..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	1 Day	trendadm	Discover info...
Flagged_..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	15 Mn...	trendadm	Discover info...
All_Inds..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	1 Day	trendadm	Discover info...
Flagged_..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	15 Mn...	trendadm	Discover info...
All_Inds..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	1 Day	trendadm	Discover info...
Flagged_..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	15 Mn...	trendadm	Discover info...
All_Inds..._dope_smo	MV	Raw Interface by I...	discover-by-l...	discover-by-l...	pollr7	pollr7	1 Day	trendadm	Discover info...
ATM_Inf..._dope_smo	MV	YV_ATMEntry_Inf...	Infos-ATM	Infos-ATM	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
ATM_Inf..._dope_smo	MV	YV_ATMEntry_Inf...	Infos-ATM	Infos-ATM	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
ATM_Inf..._dope_smo	MV	YV_ATMEntry_Inf...	Infos-ATM	Infos-ATM	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
Device S..._dope_smo	MV	Raw MIB4 System...	CGMIB4F_D...	CGMIB4F_D...	pollr7	pollr7	1 Day	trendadm	System grou...
Infos_C..._dope_smo	MV	Raw Infos Data	Infos-Inf...	Infos-Inf...	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
InfDesc..._dope_smo	MV	Raw InfDesc Data	InfDesc-Inf...	InfDesc-Inf...	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
InfName..._dope_smo	MV	Raw InfName Data	InfName-Inf...	InfName-Inf...	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
InfMAs_C..._dope_smo	MV	Raw InfMAs Data	InfMAs-Inf...	InfMAs-Inf...	pollr7	pollr7	15 Mn...	trendadm	Collect Inf...
2007-03..._dope_smo	MV	rdPmInterfaceC...	MPLS_LSR	MPLS_LSR	pollr7	pollr7	1 Hour	trendadm	
2007-03..._dope_smo	MV	rdPmInterfaceC...	MPLS_VPN	MPLS_VPN	pollr7	pollr7	1 Hour	trendadm	
2007-03..._dope_smo	MV	rdPmDeviceStats	MPLS_VPN	MPLS_VPN	pollr7	pollr7	15 Mn...	trendadm	
2007-03..._dope_smo	MV	rdPmVpnVrt	MPLS_VPN	MPLS_VPN	pollr7	pollr7	15 Mn...	trendadm	
SysRes..._ds_sr	SR	SysResOVPA_CPU	OVPA	OVPA	pollr7	pollr7	1 Hour	trendadm	The CPU res...

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 PI
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 Collections

The Collection Definition wizard lets you create a collection for a polling policy. A collection identifies the table in the database that will contain the data collected data from your devices. You can define a collection based on data from any of the following:

- MIB. See [Creating a Collection from a MIB](#) on page 173.
- HP Performance Agent. See [Creating a Collection from the HP Performance Agent or the HP Operations Agent](#) on page 197.
- CSV file. See [Creating a Collection from a CSV File](#) on page 230.
- User Defined. See [Creating a User Defined Collection](#) on page 241

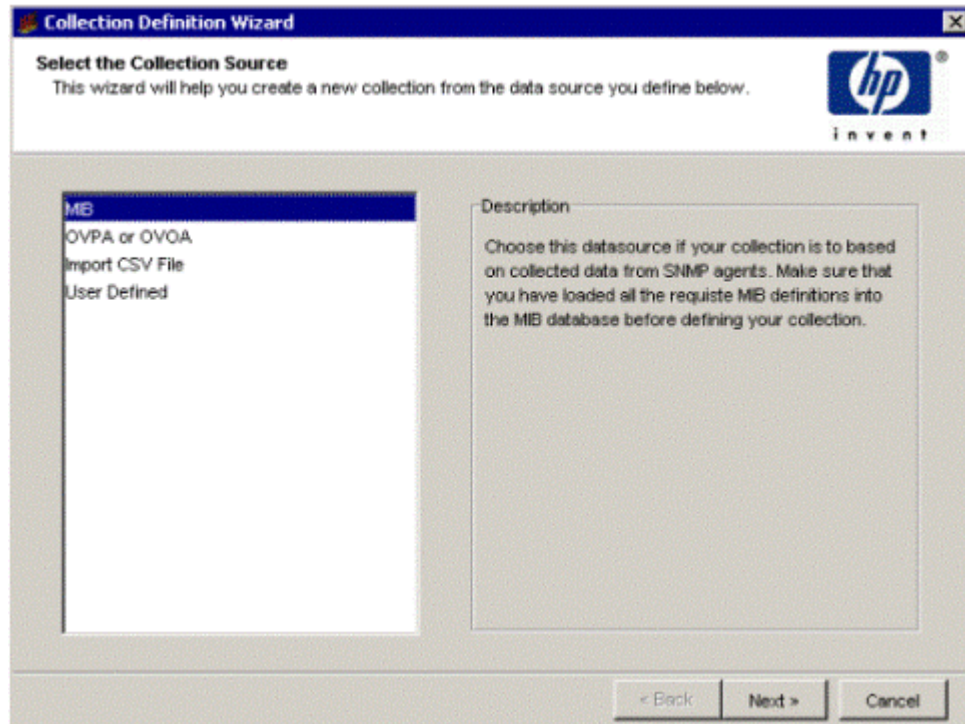
Creating a Collection from a MIB

To create a collection from a MIB, complete the following procedure:

Task 1: Start the New Collection Wizard.

- 1 From the Management Console's File menu, select **Create Collection**.

The Select the Selection Source window opens.



- 2 Select **MIB**.
- 3 Click **Next**.

Task 2: Specify a collection name, collection category, and a collection description.

The Define Collection window opens.

Collection Definition Wizard

Define Collection

Enter a collection name that will be used to create the set of tables in the database. The name must be 30 characters or less. The collection category can be selected from the existing list or you can enter a new name to be used for grouping your collections. You may optionally enter a description up to 255 characters.

Collection Name:

Collection Category:

Collection Description:

< Back Next > Cancel

- 1 In the Collection Name box, type an alias name for the data table. This name is stored in the PI data repository.



The wizard will assign a name to the collection based on the alias name you provide.

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 (_).

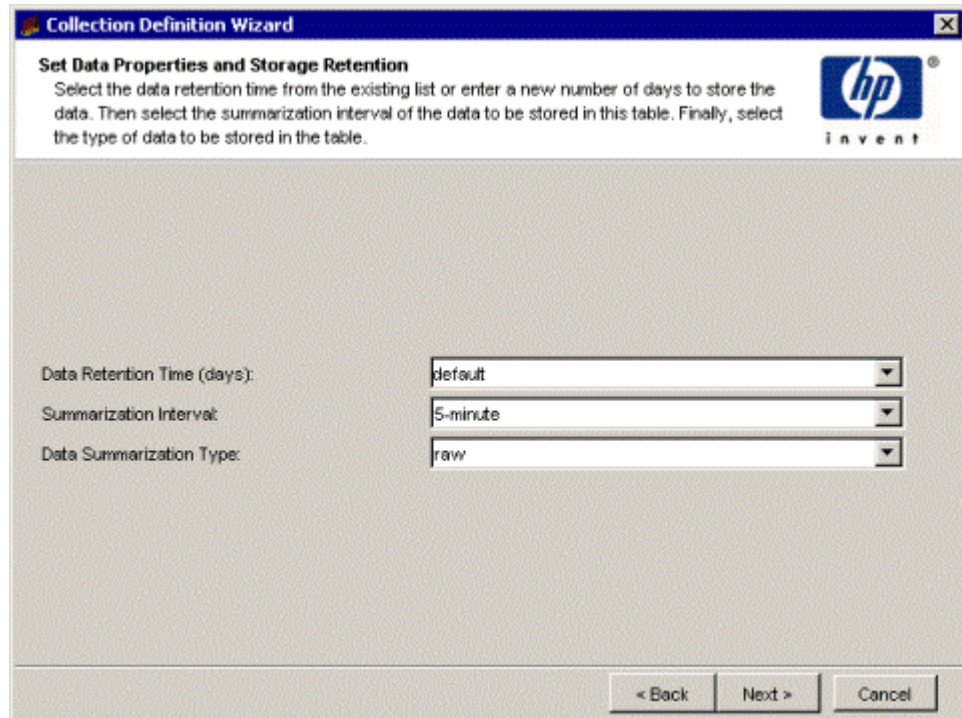
- 2 In the Collection Table Category box, type or select the category for the collection table.

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.

- 3 In the Collection 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: Set the data properties and storage retention for the collection table.

The Set Data Properties and Storage Retention window opens.



The screenshot shows a window titled "Collection Definition Wizard" with a sub-header "Set Data Properties and Storage Retention". The window includes the HP logo and the text "i n v e n t". Below the header, there is a paragraph of instructions: "Select the data retention time from the existing list or enter a new number of days to store the data. Then select the summarization interval of the data to be stored in this table. Finally, select the type of data to be stored in the table." The main area contains three dropdown menus: "Data Retention Time (days):" with "default" selected, "Summarization Interval:" with "5-minute" selected, and "Data Summarization Type:" with "raw" selected. At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

- 1 In the Data Retention Time box, type or select the number of days you want to keep the data in the table.
- 2 In the Summarization Interval box, select a time characteristic for the data in the table. For example, a 5 minute table contains data for a 5 minute interval. Data collected at 5 minute intervals is called rate data.
- 3 In the Data Summarization Type box, select a table type.


The table type specifies the class of data contained in the table. The table type options are as follows:

- **rate**
Specifies that the table contains rate data, which is delta data computed from a pair of rows from the raw input data.
- **raw**
Specifies that the table contains raw (as-pollled) data.

4 Click **Next**.

Task 4: Select the OID Objects for the collection.

The Select OID Objects for the Collection window opens.







1 In the MIB Database box, expand a MIB's folder to display its objects.

- 2 Select the objects you want to add to the Selected MIB Variables box. Each selected object becomes a column in the table.

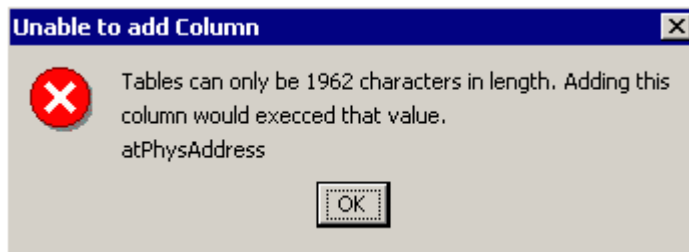


You can only have a total of 250 columns there are 13 default columns you can add. Thus, you can add 237 additional columns; otherwise, an error message displays.


Do any of the following:

- Select one object in the MIB Database box, and click .
 - To select consecutive objects, select the first item, press and hold down the **Shift** key, select the last object, and then click .
 - To select nonconsecutive objects, press and hold down the **CTRL** key, click each object, and then click .
- 3 If necessary, to remove an object from the Selected MIB Variables box, use one of the selection methods in of this procedure, and then click .

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 box. The follow message displays:

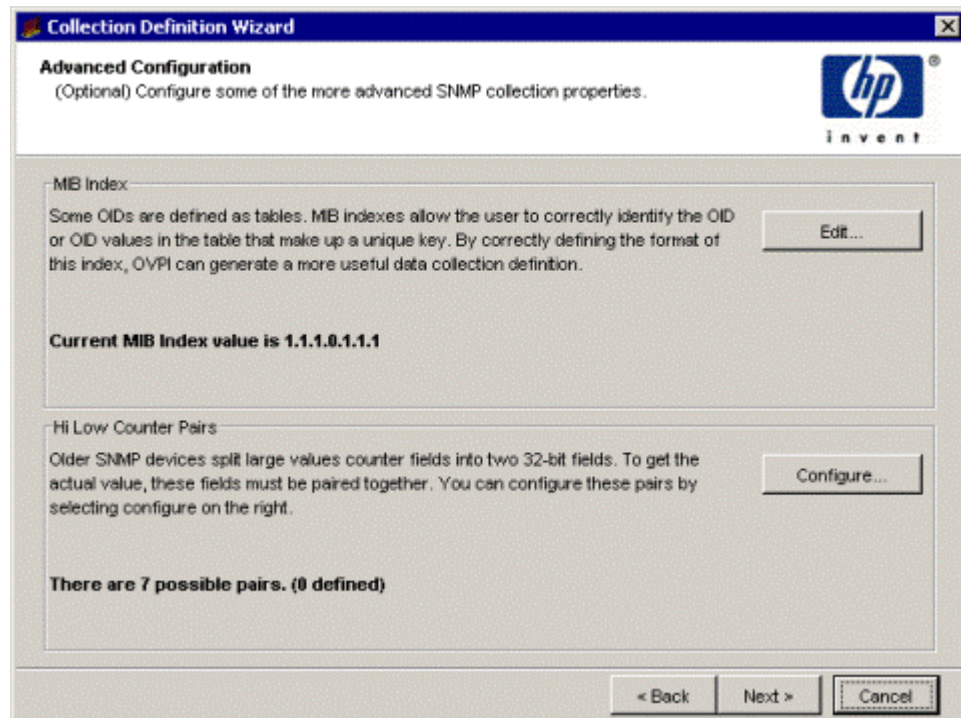


If you select objects that are not valid as columns, such as a folder, the Unable to add new column error message displays. It will list the objects. These objects will not appear in the Selected MIB variables box.



4 Click **Next**.

Task 5: Configure advanced SNMP properties.



You can configure the following:

- MIB index. Some OIDs are defined as a table. Using a MIB index, you can identify the OID or OID values that comprise a unique key.

- **Hi Low Counter Pairs.** Older SNMP devices split counter fields with large values into two 32-bit fields. To get the actual values, you must pair these fields together.
- 1 Do any of the following, if necessary:
 - Click **Edit** to define the format of a MIB index. See [Determining the Format of a MIB Index](#) on page 180 and [Specifying a MIB Index Format](#) on page 183 for details.
 - Click **Configure** to define a 64-bit counter. See [Defining a 64-bit Counter](#) on page 185.
 - 2 If you do not need to use either of these features, click **Next** and see [Task 6](#) on page 186.

Determining the Format of a MIB Index

The following example shows how to determine the MIB index format for the `frCircuitEntry` object in the `frCircuitTable` group in the `rfc2115-FRAME-RELAY-DTE` MIB.

- 1 Open the MIB file. Locate the entry and note that its index is `frCircuitIfIndex`, `frCircuitDlci`, which is shown in the following figure:

```

frCircuitEntry OBJECT-TYPE
    SYNTAX      FrCircuitEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "The information regarding a single Data Link
        Connection. Discontinuities in the counters contained
        in this table are indicated by the value in
        frCircuitCreationTime."
    INDEX ( frCircuitIfIndex, frCircuitDlci )
    ::= { frCircuitTable 1 }


FrCircuitEntry ::=
    SEQUENCE {
        frCircuitIfIndex      InterfaceIndex,
        frCircuitDlci         DLCI,
        frCircuitState        INTEGER,
        frCircuitReceivedFECNs Counter32,
        frCircuitReceivedBECNs Counter32,
        frCircuitSentFrames   Counter32,
        frCircuitSentOctets   Counter32,
        frCircuitReceivedFrames Counter32,
        frCircuitReceivedOctets Counter32,
        frCircuitCreationTime TimeStamp,
    }

```

The index comprises one integer value (frCircuitIfIndex) followed by another integer value (frCircuitDlci).

When viewed with the MIB Browser, this indicates that the index will be formatted using an integer value, separated by a period, followed by a second integer value (for example, *n.n*).

- 2 Open the MIB using MIB Browser and navigate to the frCircuitEntry object in the frCircuitTable group, which is shown in the following figure:



When you select the object and perform an SNMP Get request, the object is polled. The Query Results box shows the values of that object for all of its instances in that table. For example, for the `frCircuitIfIndex` instance, the MIB index format is 1.1, because 3 and 20 are each single integers. You would type 1.1 in the MIB index box in [step 1](#) on page 183 for the `frCircuitEntry` object.

After you specify the MIB index format in [step 1](#) on page 183, do the following:

- Create the new property table columns. It is suggested that you give the tables the same names as the index components ([Task 7](#)).
- Make sure you set the correct datatypes and, if applicable, the length ([Task 8](#)).
- Specify each new column as an Object-By and Collection-By variable ([Task 8](#)).

- After you specify the columns, click to clear the Object-By and Collection-By options for dsi_table_key ([Task 8](#)).



If you do not perform the steps above, your collection definition will not work.


Specifying a MIB Index Format

This step is optional. Some OIDs are defined as a table. Using a MIB index, you can identify the OID or OID values that comprise a unique key.

To specify the format of a MIB index value:

- 1 To enter a MIB index value, click Edit in the Advanced Configuration window ([page 179](#)).

The MIB Index value window opens.



The MIB index applies to tabled entries only. Non-tabled entries will have an index of zero.

- 2 Type the MIB index format. See your MIB to determine this.

The MIB index format would be determined as follows:

- In the example shown above, the MIB index format would be 1.

- When there are two indexes, the MIB index format would be 1.1, indicating two indexes that are each one integer in length.
- If you had three entries that were `ifIndex` (integer), `atfIndex` (integer), and `atNetAddress` (an IP address, which is character string value), the MIB index format would be `1.1.maxstringlength`.

In this instance, suppose the IP address is 118. 4.20. 7. The maximum string length for the address would be four sets of three integers (12 total) delimited by three periods, making the maximum length 15. In this case, the MIB index format is 1.1.15.



Note that the maximum length for any character string value cannot exceed 255 characters.

- 3 Click **OK** in the Enter a MIB Index value window.
- 4 Click **Next** in the Advanced Configuration window to continue.

When you have non-tabled entries, the following occurs:

- Non-parsed index will be inserted into `dsi_table_key`.
- If there are multiple index values, they will be concatenated and decimal delimited (for example, 1.2).

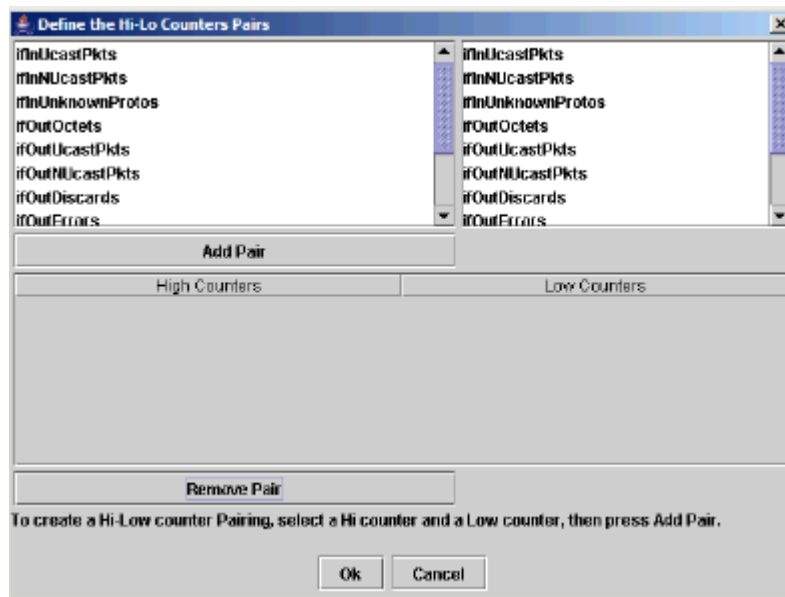
When you supply a MIB index for tabled entries, the following occurs:

- The `mw_collect` command will insert each component associated with the index into the property table column(s) that you create.
- After `mw_collect` inserts the components, it will take each component and underscore delimit it and insert it into the table key (for example, 1_2).

Defining a 64-bit Counter

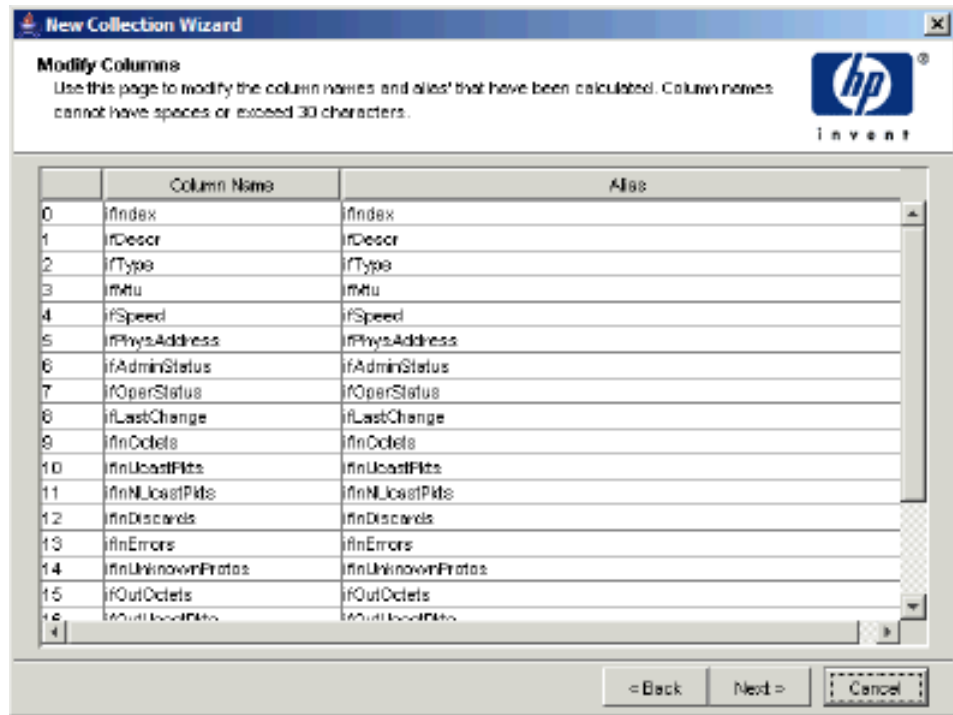
This step is optional. You only need to perform this task if you are using 32 bit counters. You would need a high and low pair to get a full 64 bit counter value. You create an artificial 64-bit counter by ANDing two 32-bit counters.

- 1 To define a 64-bit counter, click **Configure** in the Advanced Configuration window (page 179). The Define Hi-Lo Counter Pairs window opens.



- 2 Select a counter from each column and click **Add Pair**.
Continue in this manner to add pairs as necessary to define all the desired counters for the new collection table.
The wizard does not determine the appropriate pairings. You have to know the high and low pairs that you want to pair.
- 3 If you add a pairing by mistake, select it and click **Remove Pair**.
- 4 Click **OK** in the Define the Hi-Lo Counter Pairs window.
- 5 Click **Next** in the Advanced Configuration window to continue.

Task 6: Modify column and alias names.



If needed, you can modify the column and alias names.

- 1 Click in any Column Name or Alias box to edit its text.
- 2 Click **Next**.

Task 7: Specify the property table associated with the data table.

The Associate with a Property Table window opens.

- 1 Do one of the following:
 - Select **Create a new property table**, and specify a name for the property table that describes the data it will collect. The name must be unique and cannot be greater than 30 characters. Valid characters are alphabetic characters, numeric digits, and underscores (_). Go to [Task 8](#) on page 187.

- Select **Reuse existing property table**, and then select a property table from the list. Go to [Task 9](#) on page 194.



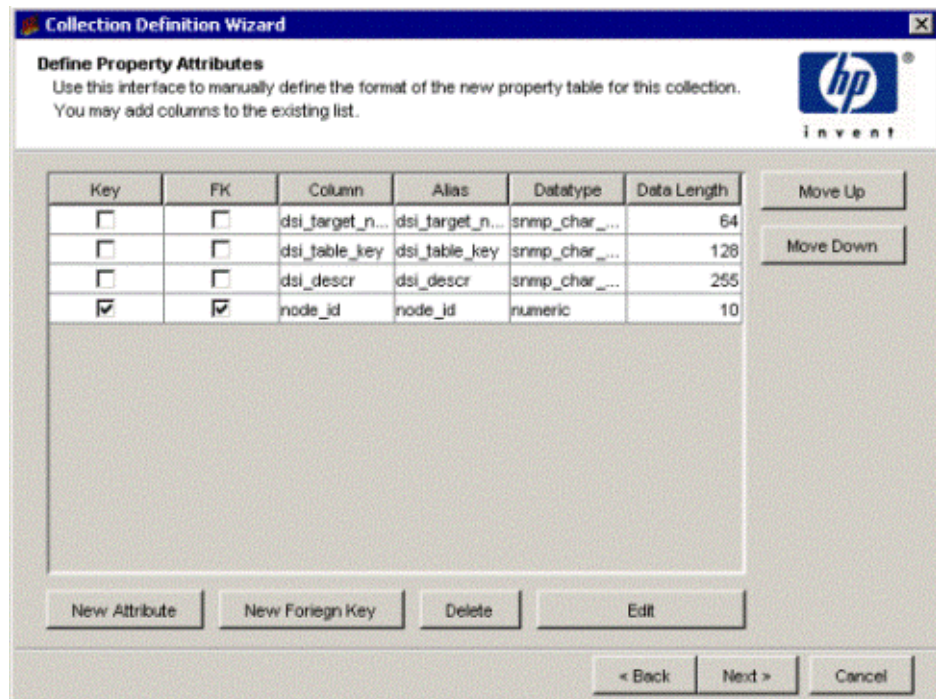
If you reuse an existing property table, you may corrupt data that is already stored in that table.

2 Click **Next**.

Task 8: Define the property attributes.

You can create a property table with the existing columns or you can add more columns.

The Define Property Attributes window opens.



The following options are provided:

- Add a column or add a column with a foreign key to a property table and set up the column map directives in TEEL.

See [Adding a Column](#) on page 188 and [Adding a Column with a Foreign Key](#) on page 191.

- Change the order of columns.
See [Changing the Order of Columns](#) on page 193.
- Delete columns.
See [Deleting a Column](#) on page 193.
- Edit columns.
See [Editing a Column](#) on page 193.

Adding a Column

To add a column to the property table, follow these steps:

- 1 To add a column, click **New Attribute**.

The Property Attribute - New dialog box opens.

- 2 Click the Name list, and select a name for the column. By default, the names of the OID objects you selected will appear.
- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 In the Datatype list, click a datatype value.

Following is the list of available datatype options and their corresponding database storage type:

PI Data Type	Database 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

- 7 In the Data Length list, click a length value.

Valid data types for length values are `char_string`, `hex_string`, `numeric`, `octet_string`, `snmp_char_string`, `snmp_hex_string`, or `snmp_octet_string`.

- 8 Select or click to clear the **Can contain NULL data box**.

When you select this option, it indicates that the column can contain NULL values. If the **Is an identifying attribute** box is checked, this option will automatically be disabled.

- 9 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

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 **Is an identifying attribute** box must not be selected and there should be no value in the Default Value box.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#) on page 193.

10 Click **Apply**.

The column is added to the property table.

11 Repeat the procedure to add more columns, or click **Next**. to continue.

[Adding a Column with a Foreign Key](#)

To add a column with a foreign key:

1 To add a column, click **New Foreign Key**.

The New Foreign Key dialog box opens.

The screenshot shows a dialog box titled "Property Attribute - New". It has a standard Windows-style title bar with a close button. The dialog contains the following fields and controls:

- Name:** A text box containing "Column_4".
- Alias:** A text box containing "Column_4".
- Description:** An empty text box.
- Default Value:** An empty text box.
- Foreign Key Table:** A dropdown menu with "ksi_user" selected.
- Is an identifying attribute:** An unchecked checkbox.
- Buttons:** "Apply" and "Cancel" buttons at the bottom right.

- 2 Type a name for the column in the Name box.

The column name cannot exceed 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.

- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 Click the Foreign Key Table list and select the property table that has the key.
- 7 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#) on page 193.

- 8 Click **Apply**.
The column is added to the property table.
- 9 Repeat the procedure to add more columns, or click **Next** to continue.

Changing the Order of Columns

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 **Move Up**.
- To move the column down, click **Move Down**.

Deleting a Column

To delete a column from the property table, click **Delete**.

Editing a Column

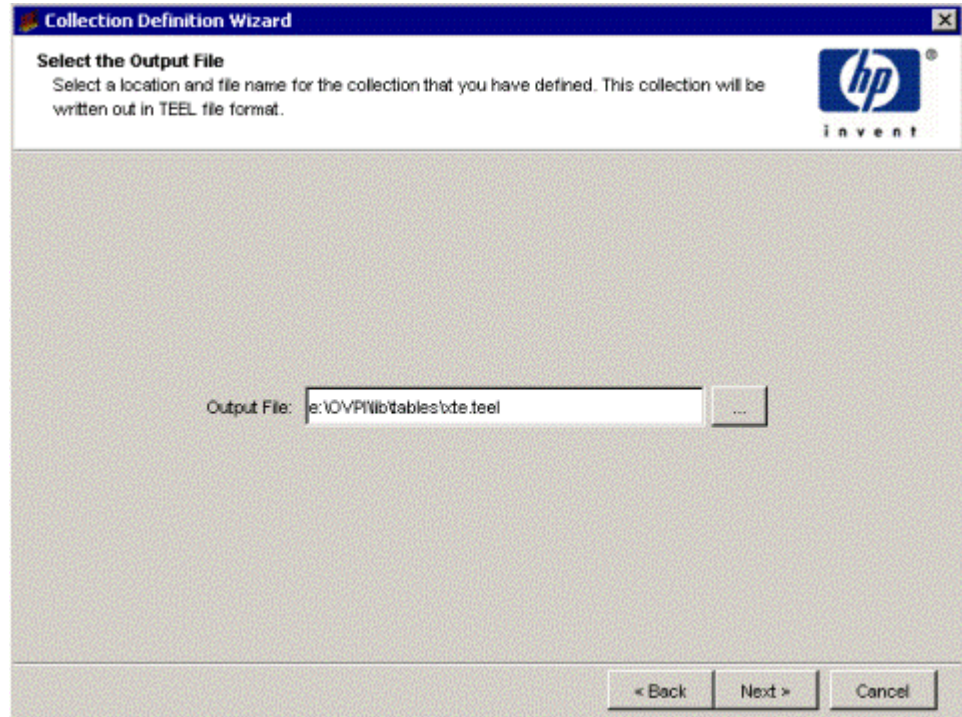
To edit a column:


- 1 Double-click the column in the Define Property Attributes window.
The Property Attribute Edit dialog box opens, which is identical to the Property Attribute New or Property Attribute New Foreign Key dialog boxes.
- 2 Make any changes.
- 3 Click **Apply**.

Repeat the procedure, or click **Next** to continue.

Task 9: Specify an output file for the collection.

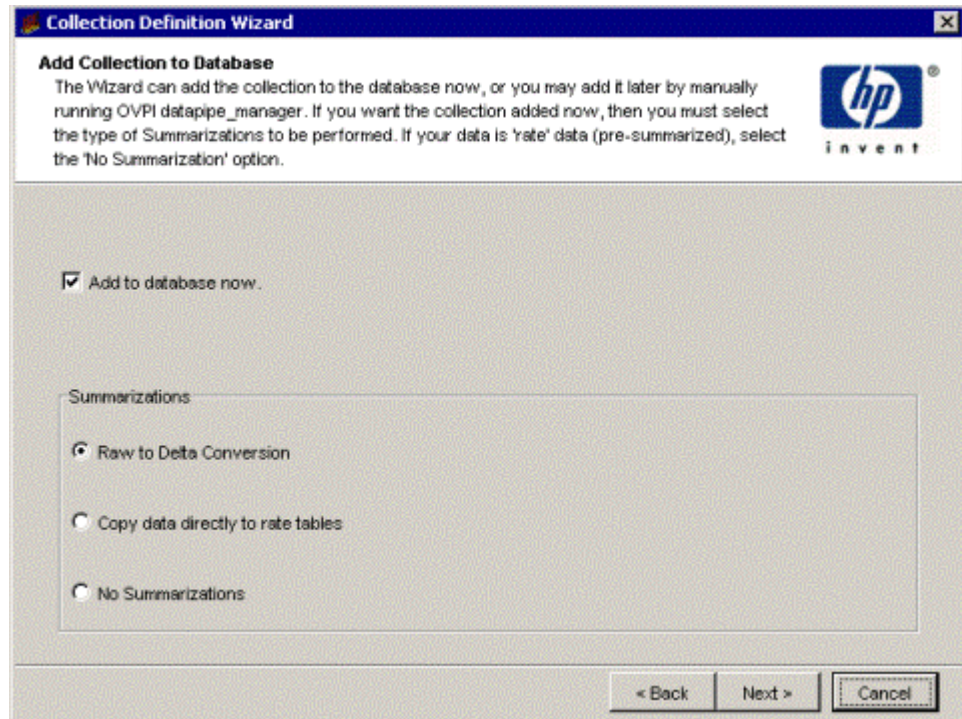
The Select the Output File page opens.



- 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 10: Register the collection with the database.

The Add Collection to Database window opens.



- 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 type 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.

If you need additional information about the options for the **datapipe_manager** command, see the *HP Performance Insight TEEL Reference Guide* and the *HP Performance Insight Reference Guide*.

- 2 In the Summarizations box, select either **Copy Data Directly to Rate Tables** or **No Summarizations**.

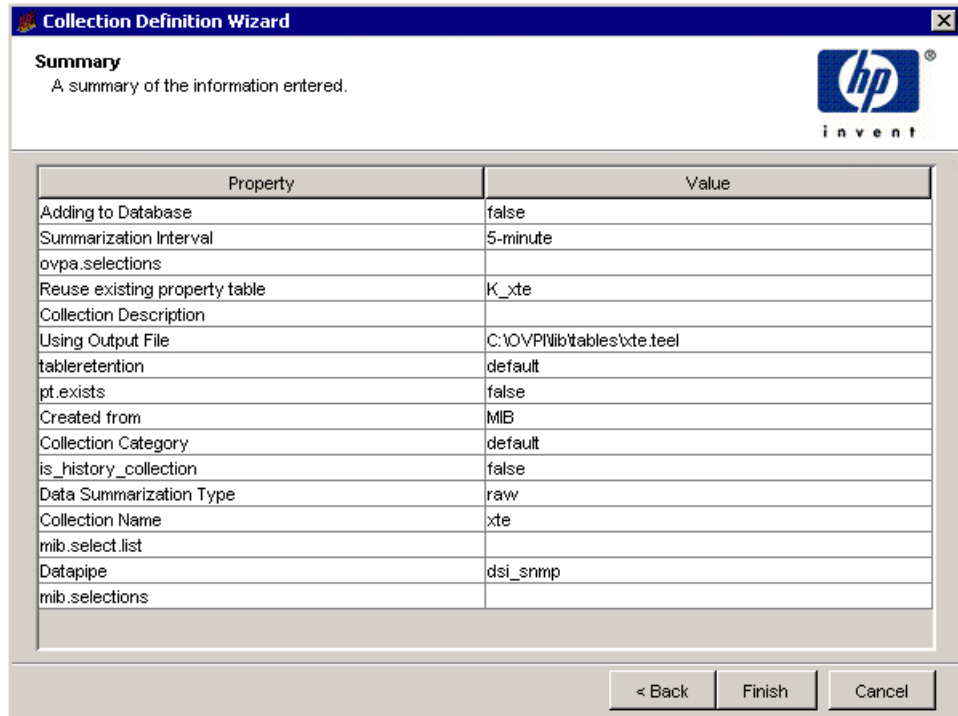


Additionally, if you want to do hourly or weekly summarizations, you must create a roll-up file (.sum file) and execute it using the `trend_sum` command. See the *HP Performance Insight Reference Guide* for details about using the `trend_sum` command.

- 3 Click **Next**.

Task 11: 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 PI creates the TEEL file and stores it in the designated location.

- 3 Create a polling policy for the collection. See [Creating Polling Policies](#) on page 252 for details.

Creating a Collection from the HP Performance Agent or the HP Operations Agent

This section includes the following topics:

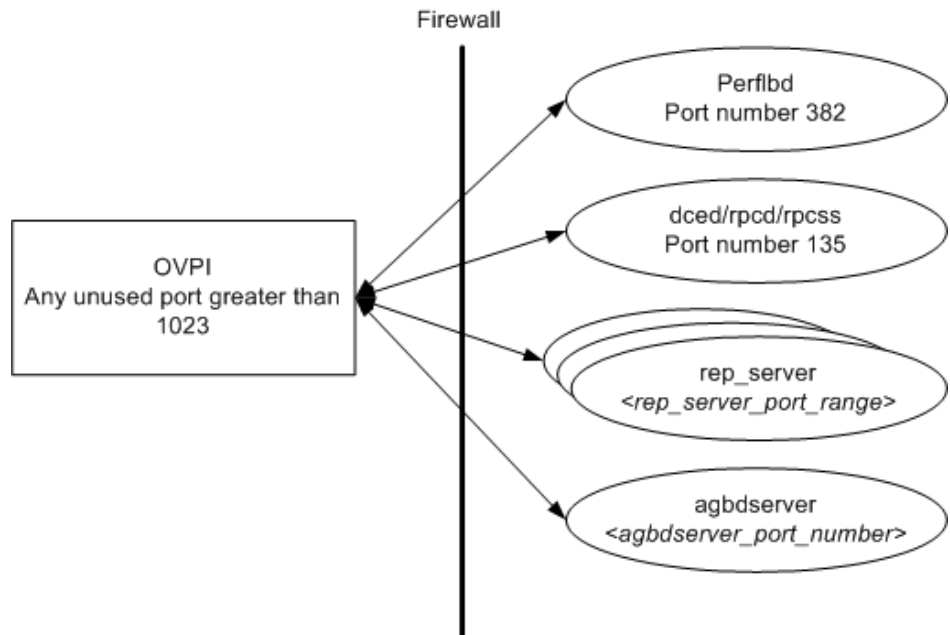
- Communicating in a firewall environment.
- Collecting Data from Agents with HTTPS Access.
- Create Perl script for data transformation on row basis.

Communicating in a Firewall Environment

This section describes the steps to perform to allow communication between PI and HP Performance Agent across a firewall.

Configuring the Performance Agent 3.x Ports

The basic communication model between PI and HP Performance Agent 3.x across a firewall is shown in the following figure.



To configure the port settings for the `agbdserv` and `rep_server` services on the Performance Agent system, you must log on using an account that has administrator privileges and follow these steps:

- 1 Locate the `services` file, which is stored in the following location:
 - UNIX: `/etc`
 - Windows: `%windir%\system32\drivers\etc`
- 2 Open the file using a text editor and add the following parameters at the end of the file.

```
agbdserv <port_number>/tcp
```


```
rep_server <port_number>/tcp
```

If there are multiple data sources configured in the `perflbd.rc` file, the first `rep_server` will use the specified port number. All other `rep_server` services add one to the last used port number.

- 3 Save the file.
- 4 Restart the agent services.

Configuring the Performance Agent 4.x Ports

The basic communication model between PI and HP Performance Agent 4.x across a firewall is shown in the following figure.



By default, CODA uses any unused port in addition to port 383 used by the BBC Communication Broker. To allocate a port number of your choice to CODA on the Performance Agent system, you must log on using an account that has administrator privileges and run the following commands at the command prompt:

```
ovconfchg -ns coda.comm -set SERVER_PORT <port_number>
```

```
ovc -restart
```

To enable single-port communication, you can configure the port settings for CODA to share the same port used by BBC Communication Broker. To do so, type the following commands at the command prompt:

```
ovconfchg -ns coda.comm -set SERVER_BIND_ADDR localhost
```

```
ovc -restart
```

Creating a Collection from the HP Performance Agent or the HP Operations Agent

You can create a collection that identifies the metrics you want to collect from the HP Performance Agent or the embedded performance component (EPC) of the HP Operations agent. These agents collect, summarize, time-stamp, and log performance data for a system that is used to monitor the health of a system.



Before you can create a collection, you must discover systems that have agents installed on them. The best way to discover and type agent systems is to install the SR_OVPA datapipe and OVPA_Collection_Datapipe. See the *System Resource OVPA Datapipe Release Statement*, which is available when you install the System Resources Report Pack.

The `pa_discovery` command is a utility that allows you to discover an HP Performance Agent or an HP Operations agent on a PI system. For complete information on `pa_discovery` command see the *HP Performance Insight Reference Guide*.

After the data is collected through `pa_discovery`, you need to run the following script:

```
trend_proc -f {DPIPE_HOME}/scripts/SR_OVPA_Daily.pro
```

A collection is created using the Polling Policy Manager application.


On a front end only UNIX machine, there is no `trendadm` user and hence PI specific environment variables are not created. If you want to create a collection TEEL file from an HP Performance Agent Datasource on a front end only machine, you will have to manually set the following library search path environment variables, if they already don't exist:

Operating System	Environment variable	Required Paths
HPUX	SHLIB_PATH	/opt/OV/lib and \$DPIPE_HOME/lib
Linux and Solaris	LD_LIBRARY_PATH	/opt/OV/lib and \$DPIPE_HOME/lib
Windows	PATH	C:\Program Files\HP OpenView\bin and %DPIPE_HOME%\bin

To create a collection, complete the following procedure:

Task 1: Start the Collection Definition wizard.

- 1 On the Management Console click the **Administration** snapin. The Administration navigation window opens.
- 2 Select **Polling Policy** from the navigation tree.
- 3 On the Menu bar, click **File** and then select **Create Collection**.
Select the Collection Source window opens.



- 4 Select **OVPA or OVOA**.
- 5 Click **Next**.

Task 2: Specify a collection name, collection category, and a collection description.

The Define Collection window opens.

Collection Definition Wizard

Define Collection

Enter a collection name that will be used to create the set of tables in the database. The name must be 30 characters or less. The collection category can be selected from the existing list or you can enter a new name to be used for grouping your collections. You may optionally enter a description up to 255 characters.

Collection Name:

Collection Category:

Collection Description:

< Back Next > Cancel

- 1 In the Collection Name box, type an alias name for the data table. This name is stored in the PI data repository.



The wizard will assign a name to the collection based on the alias name you provide.

The name must be unique and can be up to 255 characters. Valid entries include uppercase and lowercase alphabetic characters, all numeric digits, and the special characters hyphen (-) and underscore (_).


- 2 In the Collection Table Category box, type or select the category for the collection table.

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.

- 3 In the Collection 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: Set the data properties and storage retention for the collection table.

The Set Data Properties and Storage Retention window opens.



- 1 In the Data Retention Time box, type or select the number of days you want to keep the data in the table.
- 2 In the Summarization Interval box, select a time characteristic for the data in the table. For example, a 5 minute table contains data for a 5 minute interval. Data collected at 5 minute intervals is called rate data.



When defining a collection for HP Performance Agent or HP Operations agent, the only data type available is raw.

- 3 In the Data Summarization Type box, select a table type.

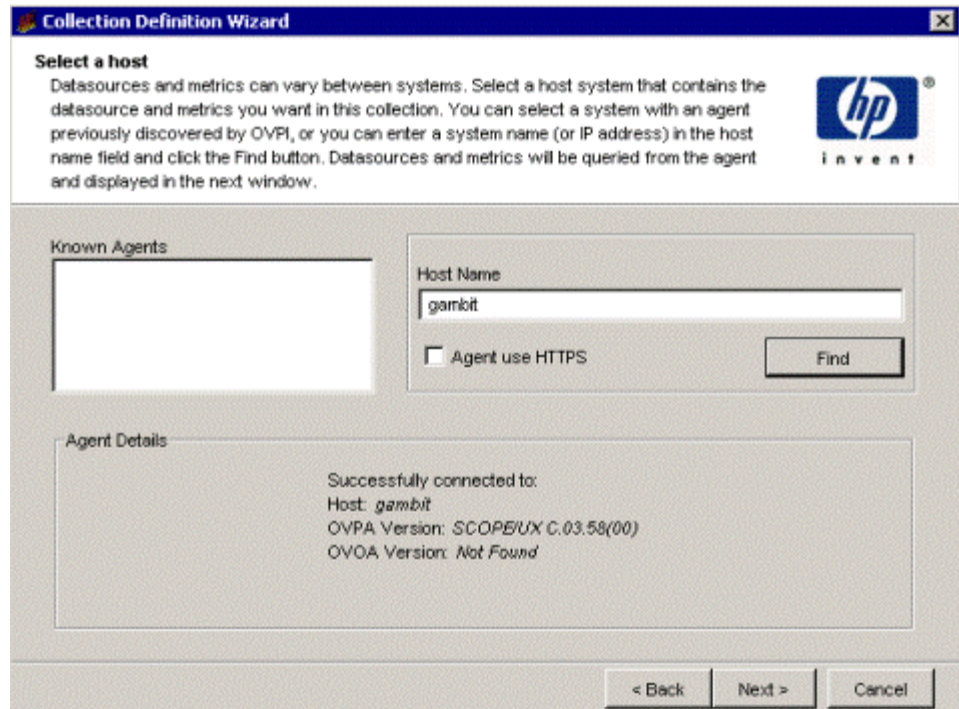
The table type specifies the class of data contained in the table. The table type options are as follows:

- **rate**
Specifies that the table contains rate data, which is delta data computed from a pair of rows from the raw input data.
- **raw**
Specifies that the table contains raw (as-pollled) data.

4 Click **Next**.

Task 4: Select the host to retrieve the metrics.

The Host Selection window opens.




- 1 Select the host on which the agent is installed by doing either of the following:
 - In the Known Agents box, select a host.

- In the Hostname box, type the hostname of the system on which the agent is installed, and then click **Find**.
- ▶ Known agents are located in the `K_pdatasources` property table. This table is populated by the SQL script provided by `SR_OVPA` and `OVPA_Collection` datapipes.

2 Click **Next**.

Task 5: Select the metrics that you want to use in the collection.




The Metric Selection window opens.



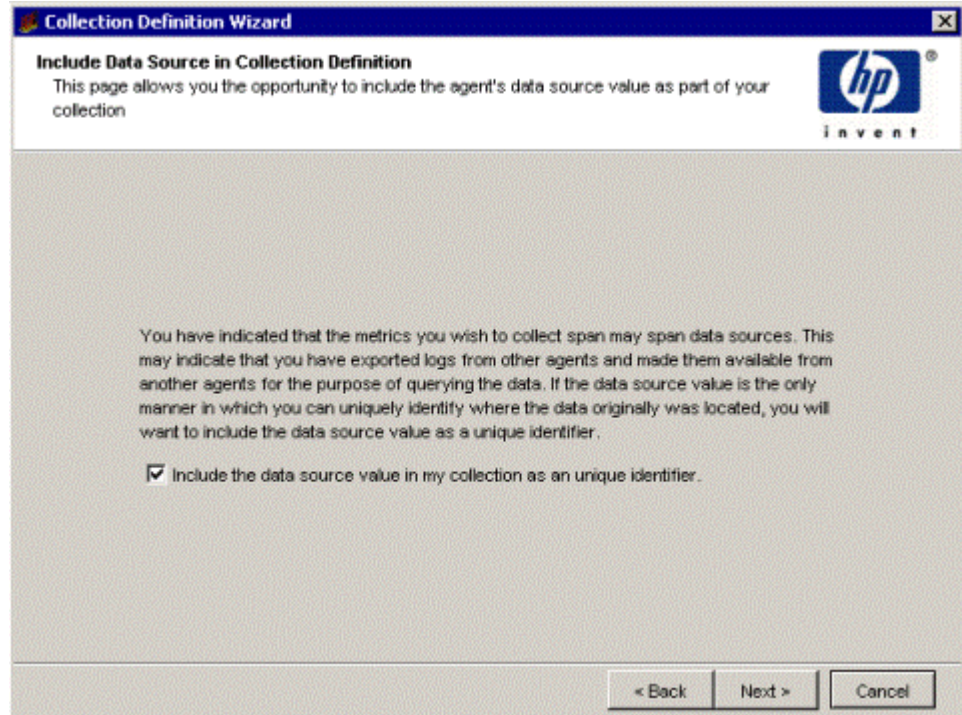
1 Expand the list of metrics in the Available Metrics box.

- ⚠ Ensure that the agent metric names do not have spaces, otherwise the data collection might fail.

If there are spaces in the metric names, rename the metrics on the agent side.

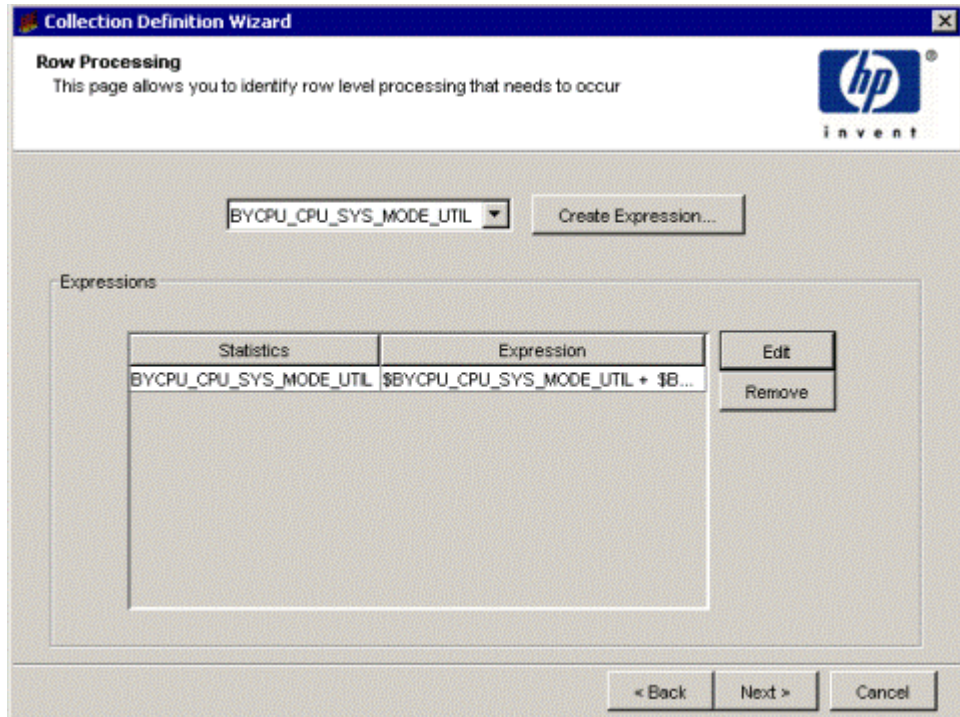
- 2 Do any of the following:
 - Select one metric in the Available Metrics box, and click .
 - To select consecutive objects, select the first item, press and hold down the **Shift** key, select the last object, and then click .
 - To select nonconsecutive objects, press and hold down the **CTRL** key, click each object, and then click .
- 3 Select either or both of the following options:
 - Select the **Selected metrics occur in more than one datasource** check box when either of the following is true:
 - Data to be collected spans multiple agent data sources.
 - Data for metrics selected in the Selected Metrics window is to be gathered from all mounted agent data log files.
 - Data is to be collected from either SCOPE or CODA data sources.

The Include Data Source in Collection Definition Wizard window opens.



- Select the **Selected Metrics need Row Processing** check box if you want to perform Row processing on the selected metrics. Row processing can help you alter the collected data to suite your requirement.

The Row Processing window opens.



- 4 Click **Next** and do one of the following:
 - If you have selected the **Selected metrics occur in more than one datasource** check box in [step 3](#), you should complete the data transformation procedure using Perl script. Go to [Task 6](#).
 - If you have not selected the **Selected metrics occur in more than one datasource** check box in [step 3](#), go to [Task 7](#).

Task 6: Create Perl script for data transformation on row basis.

Agent Metrics Imported to Another System

The data collected on an agent system is periodically extracted from the originating system and imported into another agent on a different system. This allows you to query the metric data and perform analysis without hampering the performance of the originating system. Thus, when data is collated from multiple systems, the only way to uniquely identify the originating system is by the data source identifier in the metric tuple.

Use INTERVAL Metric as Delta_Time Value

Some DSI data sources as well as the standard PROCESS metrics are not stored in the agent with summarizations available because of the nature of the data. PI must be able to retrieve this data and run the raw to delta processing on it. To ensure that the delta processing is correct, the data must include a valid `delta_time` column. Typically, the agent contains this information within its INTERVAL metric. Therefore, the child collector must be able to use the agent created metric of INTERVAL for use as `delta_time` in the agent collection.

`dpipe_pa` will log metrics automatically only in these scenarios:

- If the data source is SCOPE, and the metric class for which data is collected is either PROCESS or CONFIGURATION.
- If the data source is SCOPE, and the summarization interval is 0, the INTERVAL metric of that class is logged.
- If the data source is CODA and the summarization interval is 0, the GBL_INTERVAL metric from the CODA global class is logged.

Data transformation is useful in the following scenarios:

Raw to Delta Processing for Non-summarized Metrics

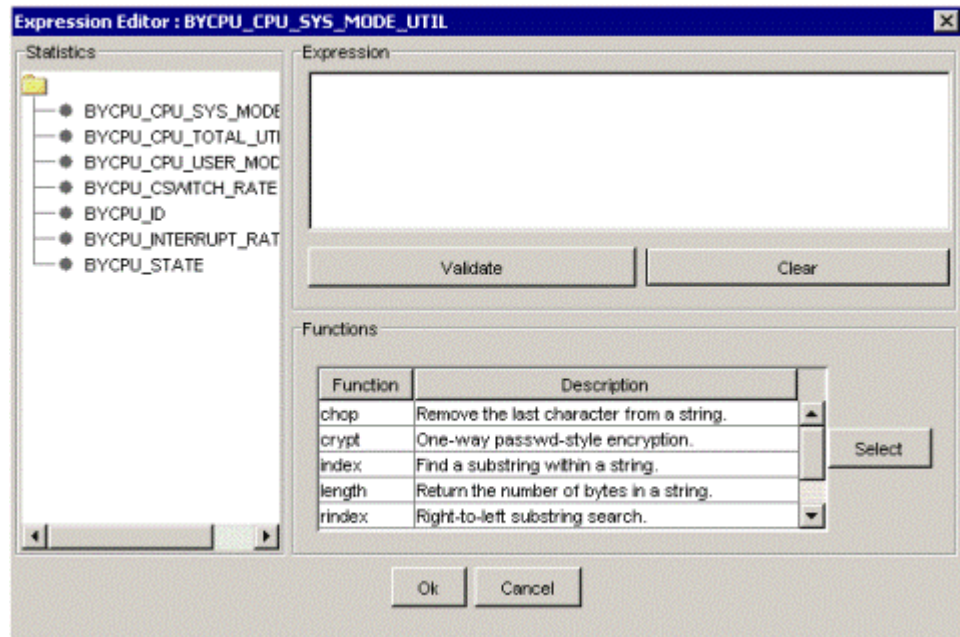
For non-summarized metrics, if you want to perform a Raw to Delta processing, you need to have the delta time between two sets of metrics collected from the agent. The Row Processing Perl scripts can be used to generate the Delta time value from the collected metrics.

Support Non-Standard Record Date

Custom DSI data sources may store data in a proprietary format and this data may be added into the agent system for subsequent reporting. When the data is added, the timestamp associated with each record is no longer the time value when the record was written, but instead it is the timestamp when the record was added to the agent system. This gives an invalid skew to the data which is then shown incorrectly within PI. The **dpipe_pa** collector therefore must be able to construct a `ta_period` from metric data instead of using only the timestamp from the agent. To be able to handle this, the collector must support the Row Process2 directive to be able to set the `ta_period` column equal to the result of an expression performed upon the remaining columns.

To create the Perl expression for transforming data on row basis, perform the following steps:

- 1 In the Include Data Source in Collection Definition window, select the check box if you want to include the agent data source name in the collection. The data source name will serve as the unique identifier.
- 2 Click **Next**.
The Row Processing window opens.
- 3 Select a metric from the drop down list and click **Create Expression**.
The Expression editor window opens.



The following information appears in the Expression Editor window:

- | | |
|------------|---|
| Statistics | The selected metrics in the collection. |
| Expression | Edit pane for creating the Perl expression. |
| Functions | List of commonly used Perl functions. |

- 4 From the Functions list, select the required Perl function and then click **Select**. The Perl function is displayed in the Expression pane. Edit the expression as required.

When you create a collection with row process directive, the datapipe name will be `dsi_tablename`.

The Perl script generated by the GUI takes care of PI internal table structure details. If you are writing your own Perl script you must be aware of PI internal structure details.

An example Perl script is similar to the following:

```
sub processRecord {
    $record = shift(@_);
    $separator = ' ';
    $Offset = 9;
    #setting array
    @MyCols = split( /$separator/, $record );
    #setting variables
    $BYCPU_CPU_CLOCK = $MyCols[$Offset + 0];

    #writing expressions
    $BYCPU_CPU_CLOCK_new = length( $BYCPU_CPU_CLOCK );

    #putting all back to array
    $MyCols[$Offset + 0] = $BYCPU_CPU_CLOCK_new;

    #constructing resulting record
    $record_new = $MyCols[0];
    for($count = 1; $count <= $#MyCols; $count++) {
        $record_new =
$record_new.$separator.$MyCols[$count];
    }

    if(length($record_new) <= (length($record) * 3)) {
        return $record_new;
    };
    return $record;
}
```

- 5 Click **Validate** to check the syntax of the Perl expression.

- If the syntax is correct, click **OK** in the confirmation message box, to return to the Row Process window and the expression will be listed in the table.

The metric for which the expression was created is displayed in the Expressions list. Use the **Edit** and **Remove** options to update or remove the expression.

- If the syntax is incorrect, click **OK** in the message box to return to the Expression editor. Do one of the following:
 - Modify the expression to correct the syntax. Revalidate the expression.
 - Click **Clear** to remove the script. and to recreate the expression, repeat [step 4](#) and [step 5](#).

6 To create Perl expressions for other metrics, repeat [step 3](#) through [step 5](#).

7 Click **Next** and go to [Task 7](#).



You can manually create a Perl script to process data without using the Collection Definition wizard. Follow the guidelines given below while creating the script:

- Define a function in the Perl script. The name of the function must be `processRecord`. Note that the name is case-sensitive.
- The datapipe name for the associated TEEL file should be unique.
- Save the Perl script.
- Add the `RowProcees2=<full path to the Perl script>` directive to the TEEL file.

The TEEL file and the Perl script is by default stored in `OVPI_HOME/lib/tables` directory.

Task 7: [Specify the property table associated with the data table.](#)

The Associate with a Property Table window opens.

- 1 Do one of the following:
 - Select **Create a new property table**, and specify a name for the property table that describes the data it will collect. The name must be unique and cannot be greater than 30 characters. Valid characters are alphabetic characters, numeric digits, and underscores (`_`). Go to [Task 8](#) on page 213.

- Select **Reuse existing property table**, and then select a property table from the list. Go to [Task 9](#) on page 220.



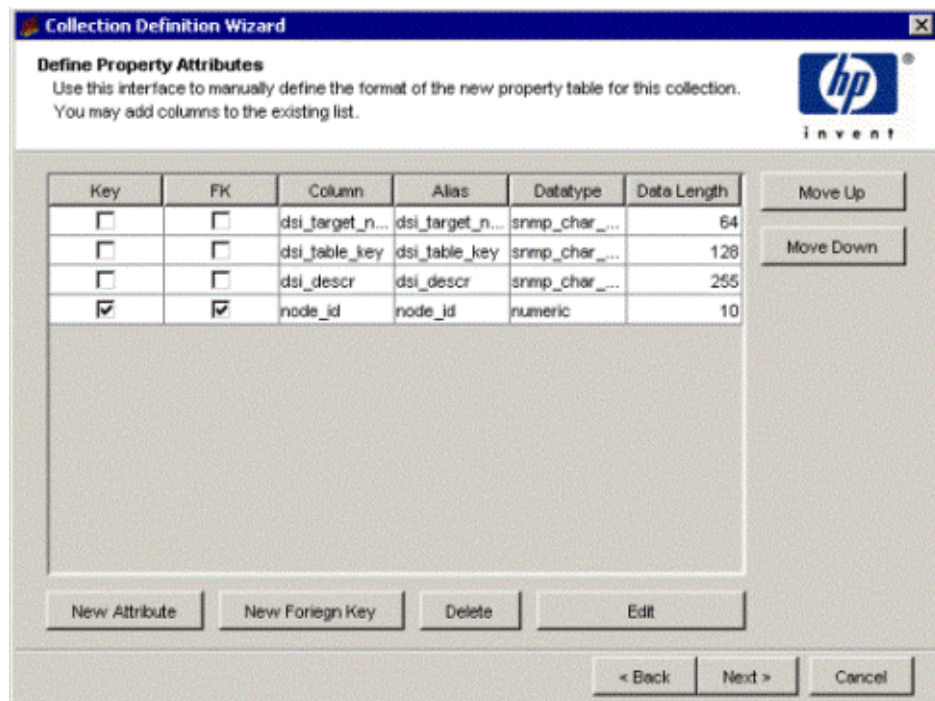
If you reuse an existing property table, you may corrupt data that is already stored in that table.

2 Click **Next**.

Task 8: Define property attributes.

You can create a property table with the existing columns or you can add more columns.

The Define Property Attributes window opens.



The following options are provided:

- Add a column or add a column with a foreign key to a property table and set up the column map directives in TEEL.

See [Adding a Column](#) on page 214 and [Adding a Column with a Foreign Key](#) on page 217.

- Change the order of columns.
See [Changing the Order of Columns](#) on page 219.
- Delete columns.
See [Deleting a Column](#) on page 219.
- Edit columns.
See [Editing a Column](#) on page 219.

Adding a Column

To add a column to the property table, follow these steps:

- 1 To add a column, click **New**.
The New Attribute dialog box opens.

- 2 Click the Name list, and select a name for the column. By default, the names of the OID objects you selected will appear.
- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 In the Datatype list, click a datatype value.

Following is the list of available datatype options and their corresponding database storage type:

PI Data Type	Database 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

- 7 In the Data Length list, click a length value.

Valid data types for length values are `char_string`, `hex_string`, `numeric`, `octet_string`, `snmp_char_string`, `snmp_hex_string`, or `snmp_octet_string`.

- 8 Select or click to clear the **Can contain NULL data box**.

When you select this option, it indicates that the column can contain NULL values. If the **Is an identifying attribute** box is checked, this option will automatically be disabled.

- 9 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

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 **Is an identifying attribute** box must not be selected and there should be no value in the Default Value box.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#).

10 Click **Apply**.

The column is added to the property table.

11 Repeat the procedure to add more columns, or click **Next**. to continue.

[Adding a Column with a Foreign Key](#)

To add a column with a foreign key, follow these steps:

1 To add a column, click **New**.

The New Foreign Key dialog box opens.

The screenshot shows a dialog box titled "Property Attribute - New". It contains the following fields and controls:

- Name:** Text box containing "Column_4".
- Alias:** Text box containing "Column_4".
- Description:** Empty text box.
- Default Value:** Empty text box.
- Foreign Key Table:** Dropdown menu with "ksi_user" selected.
- Is an identifying attribute:** Unchecked checkbox.
- Buttons:** "Apply" and "Cancel" buttons at the bottom right.

- 2 Type a name for the column in the Name box.

The column name cannot exceed 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.

- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 Click the Foreign Key Table list and select the property table that has the key.
- 7 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#).

- 8 Click **Apply**.
The column is added to the property table.
- 9 Repeat the procedure to add more columns, or click **Next** to continue.

Changing the Order of Columns

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 **Move Up**.
- To move the column down, click **Move Down**.

Deleting a Column

To delete a column from the property table, click **Delete**.

Editing a Column

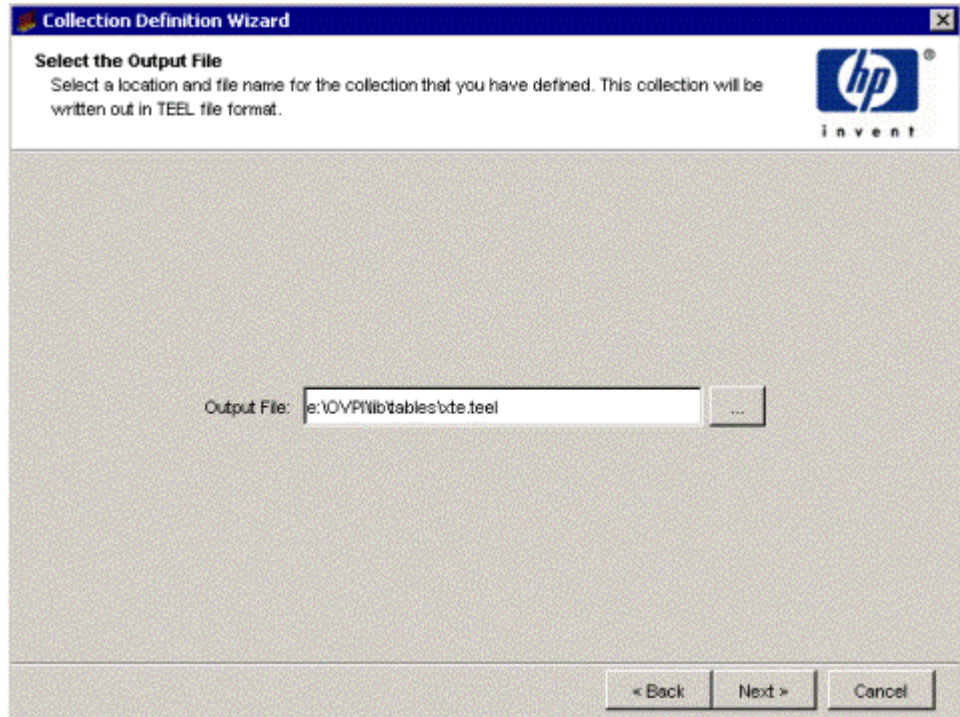
To edit a column, follow these steps:

- 1 Double-click the column in the Define Property Attributes window.
The Property Attribute Edit dialog box opens, which is identical to the Property Attribute New or Property Attribute New Foreign Key dialog boxes.
- 2 Make any changes.
- 3 Click **Apply**.


- 4 Repeat the procedure, or click **Next** to continue.

Task 9: Specify an output file for the collection.

The Select the Output File window opens.



- 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:

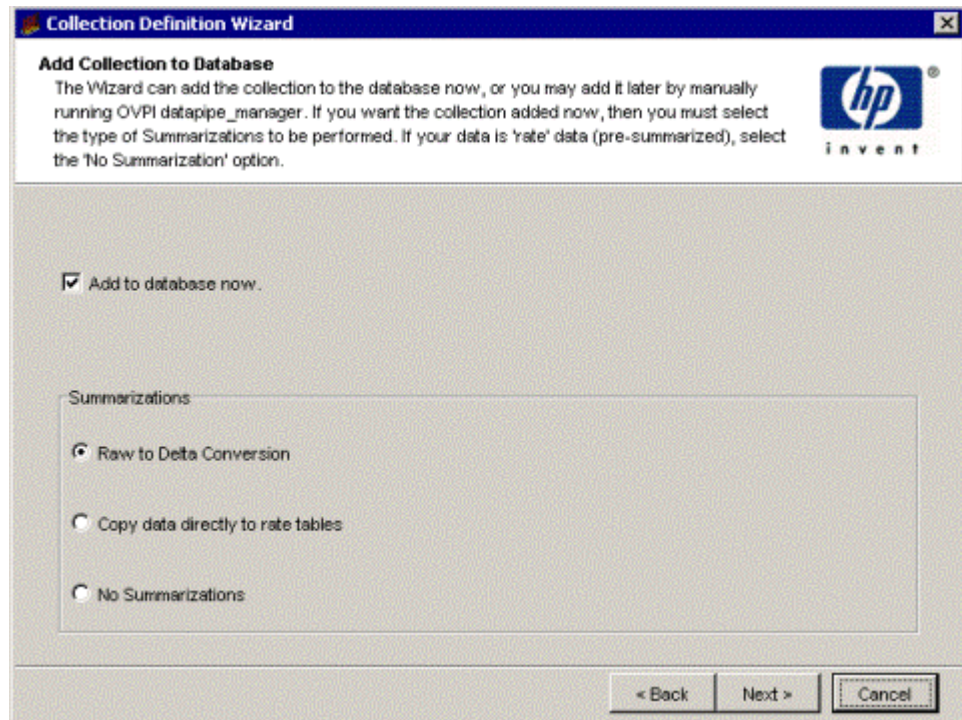
- Type the new information in the Output File box.
- Click . The Open dialog box opens.

Navigate to the folder where you want to store the output file, type the file name in the File name box, and then click **Open**. The Output File box is updated with the new file information.

- 3 Click **Next**.

Task 10: Register the collection with the database.

The Add Collection to Database window opens.



- 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 type 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.

If you need additional information about the options for the **datapipe_manager** command, see the *HP Performance Insight TEEL Reference Guide* and the *HP Performance Insight Reference Guide*.

- 2 In the Summarizations box, select either **Copy Data Directly to Rate Tables** or **No Summarizations**.

 Raw to delta conversions are not available for HP Performance Agent collections.

Additionally, if you want to do hourly or weekly summarizations, you must create a roll-up file (.sum file) and execute it using the `trend_sum` command. See the *HP Performance Insight Reference Guide* for details about using the `trend_sum` command.

- 3 Click **Next**.

Task 11: Review the Summary window.

The Summary dialog box opens.

- 1 Verify that all the collection property values are correct.


If necessary, click **Back** to return to a previous window to make a change. After you make a change, click **Next** until you return to the Summary window.

- 2 Click **Finish** to create the data table.

The data table creation wizard closes. PI creates a Perl script and a TEEL file and stores it in the designated location.

Collecting Data from Agents with HTTPS Access

PI supports collection of data from HP Performance Agents and HP Operations Agents that use HTTPS protocol for secure communication.

 HP Performance Agents and HP Operations agents will be collectively referred to as **agents** in this guide unless there is a specific need to qualify the agent by its name.

For HTTPS communication, the agents must support CODA version 10; otherwise, HTTP or DCE method will be used. Because HTTPS communication is certificate-based, certificates must be installed on the PI system (including remote pollers) that collects data from agents. PI will utilize the client certificates already installed for the agents. The PI system acts as a certificate client and the certificate server (certificate authority) is provided by the HP management server or HP Performance Agent agent systems. The client certificates must be exchanged to establish HTTPS communication.

For more information about HTTPS communication, see the *HP Operations HTTPS Agent Concepts and Configuration Guide*. This guide is available with the HP Software documentation set and is available at the following URL:

<http://h20230.www2.hp.com/selfsolve/manuals>

In some network configurations, multiple certificate servers may be providing certificates for the agents. In this scenario, you may obtain trusted certificates from all the certificate servers for the respective agents. This can be performed using the Certificate Management application provided in the PI Management Console. For more information, see [Manage Certificates](#) on page 223.



If SSL_SECURITY option is set to ALL or REMOTE in the [coda] namespace on HP Performance Agent systems, HTTP communication fails. Only HTTPS is supported.

Prerequisites to Collect Data

- HTTPS communication must be enabled on the network PI is monitoring.
- Certificates must be installed on the PI system that polls data from the agents. For more information about installing a certificate, see [Install a Certificate](#) on page 226.
- The name of the certificate server must be known to the PI administrator. The administrator will be required to provide the certificate server name while managing certificates.
- You must update the *PI_installation_directory/data/pa_rpt.cnfg* configuration file. Open the file using a text editor and enable the `https` option. This will allow data collection from HP Performance Agent systems. If HTTPS communication fails, HTTP will be used. This is the default CODA behavior.

Manage Certificates

You can use the Certificate Management application to install and manage certificates on the PI system that collects data from the agents.

Start the Certificate Management Application

You can start the Certificate Management application from the [PI Management Console](#) or the [Command Line](#).



The Certificate Management application is available in the Management Console only when remote pollers or the PI Database Schema and Performance Manager components are installed.

PI Management Console

To start Certificate Management application, perform the following steps:

- 1 Start the PI Management Console.



You must log on to the Management Console using an account that has Administrator privileges.

- 2 From the Tools menu, select **Certificate Management UI**.

The Certificate Management window ([page 226](#)) opens.



If you are using Windows system, Certificate Management application can be started from the Windows Start menu.

Command Line

Do one of the following depending on your operating system:

- Windows

Type the following at the command line:

```
certmgmt
```

- UNIX

Log on as **trendadm**, and type the following at the command line:

```
certmgmt
```


The Certificate Management window opens.

The screenshot shows a window titled "Certificate Management" with a blue title bar. The window is divided into two main sections: "OV Configuration" and "Certificates".

OV Configuration:

- Contains two text input fields: "OV-Core-Id" and "CERTIFICATE AUTHORITY (CA)".
- A "Set CA" button is located to the right of the "CERTIFICATE AUTHORITY (CA)" field.
- The "Certificate Status:" is set to "Undefined: Check if BBC is running".

Certificates:

- Contains a "Node Certificates" list box with three empty rows and scroll arrows.
- Below the list box is a "Certificate Management Actions:" dropdown menu with the selected option "I need a new client certificate".
- A "Go" button is located to the right of the dropdown menu.
- Below the dropdown menu is a "Trusted Certificates" list box with three empty rows and scroll arrows.

At the bottom of the window, there are two buttons: "Refresh" and "Close".

The information in the Certificate Management window is as follows:

OV-Core-Id	The unique identifier value for each system hosting an HTTPS agent. The value is automatically assigned during installation of HP Software on that system. The value cannot be modified. Example of an OvCoreId value: d498f286-aa97-4a31-b5c3-806e384fcf6e
Certificate Authority (CA)	Represents the name of the Certificate Authority. The name must be a valid certificate server name identifier.
Certificate Status	Shows the status information.
Node Certificates	Shows the list of available node certificates.
Certificate Management Actions	Provides available actions for certificate management.
Trusted Certificates	Shows the list of available trusted certificates.

Install a Certificate

To install a certificate on the PI system using Certificate Management application, perform the following steps:

- 1 Start the Certificate Management application. See [Start the Certificate Management Application](#) on page 224.
- 2 In the Certificate Authority (CA) box, type the name of the certificate authority and then click **Change CA**.
- 3 From the Certificate Management Actions list, select **I need a new client certificate**.
- 4 Click **Go**. A confirmation dialog box appears.
- 5 Click **OK** to confirm, and then click **Refresh**.

The certificate status information appears in the OV Configuration section.

If the following status information appears, OV Communication Broker is not running.

Undefined: Check if BBC is running.

To start the OV Communication Broker, log on as a user with administrative privileges and type the following at the command line:

```
ovc -start
```

If you see the certificate requesting pending status, it implies that the certificate server is in the process of servicing the certificate request. You will have to wait for this process to complete. Click **Refresh** to see the updated status.

Import a Certificate

To import a certificate, perform the following steps:

- 1 Start the Certificate Management application.
- 2 From the Certificate Management Actions drop down list, select one of the following:
 - If you are importing a trusted certificate, select **Import a trusted certificate from disk**.
 - If you are importing a node certificate, select **Import a node certificate from disk**.
- 3 Click **Go**. The Import From dialog box opens.
- 4 Navigate to the folder that contains the certificate you want to import, select the certificate, and then click **OK**.

If you are importing a trusted certificate go to [step 6](#); otherwise, go to [step 5](#).

- 5 In the Enter Password dialog box, do the following:
 - a Type your password and click **OK**.
The Re-Enter Password dialog box opens.
 - b Retype the password and click **OK**.
- 6 Click **OK** to confirm and then click **Refresh**.

Export a Certificate

To export a certificate, perform the following steps:

- 1 Start the Certificate Management wizard.
- 2 Do one of the following:
 - If you are exporting a trusted certificate, select a trusted certificate and then select **Export a trusted certificate to disk** from the Certificate Management Actions list.
 - If you are exporting a node certificate, select a node certificate and then select **Export a node certificate to disk** from the Certificate Management Actions list.
- 3 Click **Go**. The Export To dialog box opens.
- 4 Navigate to the folder where you want to store the certificate.
- 5 Type the file name in the Enter file name box, and click **OK**. If you are exporting a trusted certificate go to [step 6](#); otherwise, go to [step 7](#).
- 6 In the Enter Password dialog box, do the following:
 - a Type your password and click **OK**.
The Re-Enter Password dialog box opens.
 - b Retype the password and click **OK**.
- 7 Click **OK** to confirm and then click **Refresh**.

Update Trusted Certificates

Multiple certificate servers may be providing certificates for agents in distributed fashion. In this scenario, you may obtain trusted certificates from all the certificate servers for the respective agents.

To do so, perform the following steps:

- 1 Start the Certificate Management wizard.
- 2 Select **Update trusted certificates** from the Certificate Management Actions list.
- 3 Click **Go**. A confirmation dialog box appears.
- 4 Click **OK** to confirm and then click **Refresh**.
The updated certificates will appear in the Trusted Certificates list.

Remove a Certificate

To remove a node certificate or a trusted certificate, follow these steps.

- 1 Start the Certificate Management application.
- 2 Select a certificate that you want to remove from the Node Certificates list or Trusted Certificates list.
- 3 Select **Remove the selected certificate** from the Certificate Management Actions list.
- 4 Click **Go**. A Remove Certificate dialog box opens.
- 5 Click **Yes** to confirm and then click **Refresh**.

Collecting Data from Agents through Proxy

It is possible to configure PI to collect data from the agents through a proxy. You can do this by configuring the following command in the PI server:

```
ovconfchg -ns <namespace> -set <attr> <value>
```

where:

-ns <namespace> sets a namespace for following options.

-set <attr> <value> sets an attribute (proxy) and values (port and addresses) in current namespace.

For example:

```
ovconfchg -ns bbc.http -set PROXY  
"web-proxy:8088-(*.mycom.com)+(*.a.mycom.com;*)"
```

defines which proxy and port to use for a specified hostname.

Format:

```
proxy:port +(a)-(b);proxy2:port2+(a)-(b); ...;
```

a: list of hostnames separated by a comma or a semicolon, for which this proxy shall be used.

b: list of hostnames separated by a comma or a semicolon, for which the proxy shall not be used.

The first matching proxy is chosen.

It is also possible to use IP addresses instead of hostnames so 15.*.*.*

or 15:*:*:*:*:* would be valid as well, but the correct number of dots or colons MUST be specified. IP version 6 support is not currently available but will be available in the future.

```
PROXY=web-proxy:8088- (*.hp.com) + (*.a.hp.com;*)
```

The proxy web-proxy is used with port 8088 for every server (*) except hosts that match *.hp.com, for example www.hp.com. If the hostname matches *.a.hp.com, for example, merlin.a.hp.com the proxy server will be used.

Creating a Collection from a CSV File

To create a collection from a CSV file:


Task 1: Start the New Collection Wizard.

- 1 From the Management Console's File menu, select **Create Collection**.
The Select the Selection Source window opens. See figure on [page 174](#).
- 2 Select **Import CSV File**.
- 3 Click **Next**.

Task 2: Specify a collection name, collection category, and a collection description.

The Define Collection window opens. See figure on [page 175](#).

- 1 In the Collection Name box, type an alias name for the data table. This name is stored in the PI data repository.

 The wizard will assign a name to the collection based on the alias name you provide.

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 Table Category box, type or select the category for the collection table.

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.

- 3 In the Collection 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: Set the data properties and storage retention for the collection table.

The Set Data Properties and Storage Retention window opens. See figure on page 176.

- 1 In the Data Retention Time box, type or select the number of days you want to keep the data in the table.
- 2 In the Summarization Interval box, select a time characteristic for the data in the table. For example, a 5 minute table contains data for a 5 minute interval. Data collected at 5 minute intervals is called rate data.
- 3 In the Table Type box, select a table type.

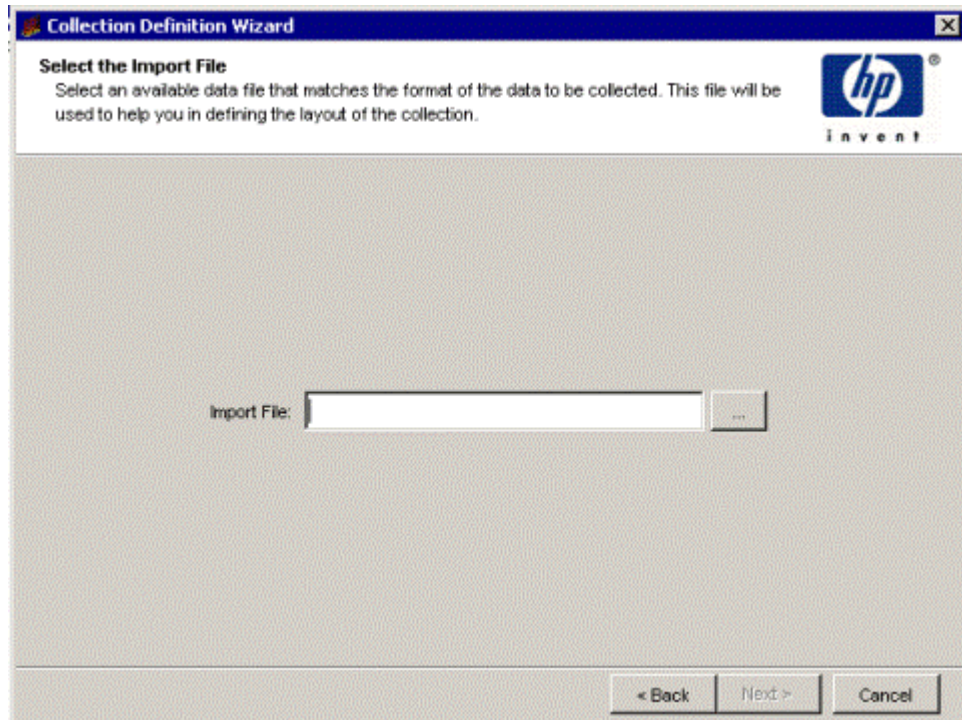
The table type specifies the class of data contained in the table. Following is the list of available table type options:

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.

- 4 Click **Next**.

Task 4: Select the CSV file.


The Select the Import File window opens.



- 1 Do one of the following:
 - Type the path to the CSV file you want to import.
 - Click **Browse** to navigate to the file using the Open dialog box. Select the file and click **OK**.
- 2 Click **Next**.

Task 5: Map the data fields from the file to the statistic you want to use in the collection.

The Map the Data Fields to Statistics window opens.



The window is populated with the data fields from the file, allowing you to map the fields to the statistic the field should use in the collection.

- 1 Select a data field, and then do the following:
 - a Click to clear the Include check box if you do not want to include the selected field.
 - b Click the Column box to type the column name for the statistic. The column name can have a maximum length of 20 characters.
 - c Click the Alias box to select an alias name for the statistic. The alias name can have a maximum length of 255 characters.
 - d Click the Datatype box to select the column name for the statistic.
 - e Click the Length box to type the length of the statistic. The length can be a number from 1 to 255.
 - f Click to clear the Nullable check box if the data field is allowed to have no data.

- g Click the Default box to type the value that should be inserted into the database for the statistic if no data is supplied.

If you leave this as null, no data is inserted. The length of this value is determined by the value specified in the Length field. For example, if the value for length is 5 and the statistic is numeric, the default value can be an integer up to 99999.

- h Click the Description box to type a description of the statistic.
- 2 Repeat [step 1](#) for each data field.
 - 3 Click **Next**.

Task 6: [Specify the property table associated with the data table.](#)

The Associate with a Property Table window opens.

- 1 Do one of the following:
 - Select **Create a new property table**, and specify a name for the property table that describes the data it will collect. The name must be unique and cannot be greater than 30 characters. Valid characters are alphabetic characters, numeric digits, and underscores (_). Go to [Task 7](#) on page 234.
 - Select **Reuse existing property table**, and then select a property table from the list. Go to [Task 8](#) on page 240.




If you reuse an existing property table, you may corrupt data that is already stored in that table.

- 2 Click **Next**.

Task 7: [Define Property Attributes](#)

You can create a property table with the existing columns or you can add more columns.

The Define Property Attributes window opens.



The following options are provided:

- Add a column or add a column with a foreign key to a property table and set up the column map directives in TEEL.

See [Adding a Column](#) on page 236 and [Adding a Column with a Foreign Key](#) on page 237.

- Change the order of columns.

See [Changing the Order of Columns](#) on page 239.

- Delete columns.

See [Deleting a Column](#) on page 239.

- Edit columns.

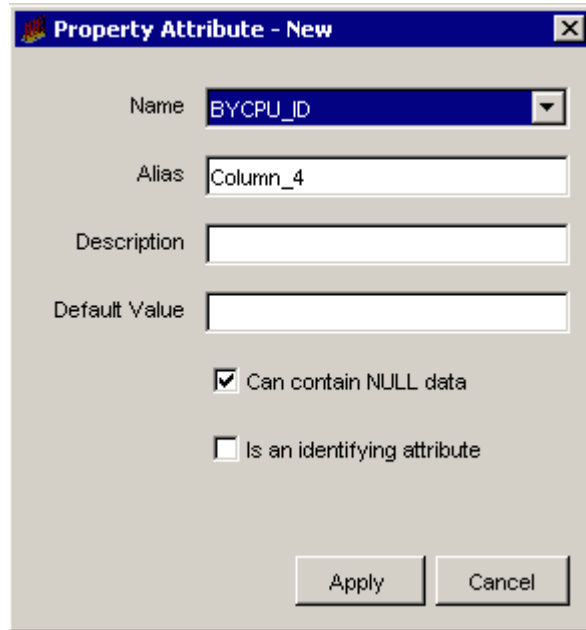
See [Editing a Column](#) on page 239.

Adding a Column

To add a column to the property table, follow these steps:

- 1 To add a column, click **New**.

The New Attribute dialog box opens.



The screenshot shows a dialog box titled "Property Attribute - New". It has a blue title bar with a close button. The dialog contains the following fields and options:

- Name:** A dropdown menu with "BYCPU_ID" selected.
- Alias:** A text box containing "Column_4".
- Description:** An empty text box.
- Default Value:** An empty text box.
- Can contain NULL data
- Is an identifying attribute
- Buttons:** "Apply" and "Cancel" buttons at the bottom.

- 2 Click the Name list, and select a name for the column. By default, the names of the OID objects you selected will appear.
- 3 Type an alias name for the column in the Alias box.
The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.
- 4 Type a description for the table in the Description box.
A description is optional. If you specify a description, it cannot exceed 255 characters.
- 5 If the column should have a default value, type a value in the Default Value box.
A default value is optional. 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.

- 6 In the Datatype list, click a datatype value.

For the available datatype options and their corresponding database storage types, see [step 6](#) on page 189.

- 7 In the Data Length list, click a length value.

Valid data types for length values are `char_string`, `hex_string`, `numeric`, `octet_string`, `snmp_char_string`, `snmp_hex_string`, or `snmp_octet_string`.

- 8 Select or click to clear the **Can contain NULL data box**.

When you select this option, it indicates that the column can contain NULL values. If the **Is an identifying attribute** box is checked, this option will automatically be disabled.

- 9 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

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 **Is an identifying attribute** box must not be selected and there should be no value in the Default Value box.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#) on page 239.

- 10 Click **Apply**.

The column is added to the property table.

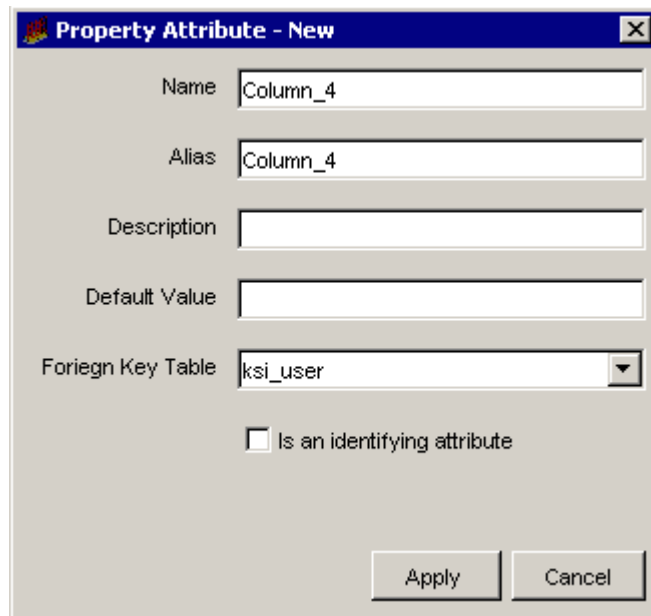
- 11 Repeat the procedure to add more columns, or click **Next**. to continue.

[Adding a Column with a Foreign Key](#)

To add a column with a foreign key, follow these steps:

- 1 To add a column, click **New**.

The New Foreign Key dialog box opens.



The screenshot shows a dialog box titled "Property Attribute - New". It contains the following fields and controls:

- Name:** Text box containing "Column_4".
- Alias:** Text box containing "Column_4".
- Description:** Empty text box.
- Default Value:** Empty text box.
- Foreign Key Table:** Dropdown menu with "ksi_user" selected.
- Is an identifying attribute:** Unchecked checkbox.
- Buttons:** "Apply" and "Cancel" buttons at the bottom right.

- 2 Type a name for the column in the Name box.

The column name cannot exceed 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.

- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 Click the Foreign Key Table list and select the property table that has the key.
- 7 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#) on page 239.

- 8 Click **Apply**.
The column is added to the property table.
- 9 Repeat the procedure to add more columns, or click **Next** to continue.

Changing the Order of Columns

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 **Move Up**.
- To move the column down, click **Move Down**.

Deleting a Column

To delete a column from the property table, click **Delete**.


Editing a Column

To edit a column, follow these steps:

- 1 Double-click the column in the Define Property Attributes window.
The Property Attribute Edit dialog box opens, which is identical to the Property Attribute New or Property Attribute New Foreign Key dialog boxes.
- 2 Make the changes and click **Apply**.
Repeat the procedure, or click **Next** to continue.

Task 8: Specify an output file for the collection.

The Select the Output File window opens. See figure on [page 194](#).

- 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.

Navigate to the folder where you want to store the output file, type the file name in the File name box, and then click **Open**. The Output File box is updated with the new file information.

- 3 Click **Next**.

Task 9: Register the collection with the database.

The Add Collection to Database window opens. See figure on [page 195](#).

- 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 type 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.

If you need additional information about the options for the **datapipe_manager** command, see the *HP Performance Insight TEEL Reference Guide* and the *HP Performance Insight Reference Guide*.

- 2 In the Summarizations box, select either **Copy Data Directly to Rate Tables** or **No Summarizations**.



Additionally, if you want to do hourly or weekly summarizations, you must create a roll-up file (.sum file) and execute it using the `trend_sum` command. See the *HP Performance Insight Reference Guide* for details about using the `trend_sum` command.

- 3 Click **Next**.

Task 10: Review the Summary window.

The Summary dialog box opens. See figure on page [page 196](#).

- 1 Verify that all the collection property values are correct.
If necessary, click **Back** to return to a previous window to make a change. After you make a change, click **Next** until you return to the Summary window.
- 2 Click **Finish** to create the data table.
The data table creation wizard closes, and PI creates a TEEL file and stores it in the designated location.
- 3 Create a polling policy for the collection. See [Creating Polling Policies](#) on page 252 for details.

Creating a User Defined Collection

You can use the collection wizard to create a collection when an import file is not available, but you know the format that is used. This method should be used for file-based collections only.

Task 1: Start the New Collection Wizard.

- 1 From the Management Console's File menu, select **Create Collection**.
The Select the Selection Source window opens ([page 174](#)).
- 2 Select **User Defined**.
- 3 Click **Next**.

Task 2: Specify a collection name, collection category, and collection description.

The Define Collection window opens ([page 175](#)).

- 1 In the Collection Name box, type an alias name for the data table. This name is stored in the PI data repository.



The wizard will assign a name to the collection based on the alias name you provide.

The name must be unique and can be up to 30 characters. Valid entries include uppercase alphabetic characters and lowercase alphabetic characters, all numeric digits, and the special character underscore (_).

- 2 In the Collection Table Category box, type or select the category for the collection table.

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.

- 3 In the Collection 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: Set the data properties and storage retention for the collection table.

The Set Data Properties and Storage Retention window opens ([page 176](#)).

- 1 In the Data Retention Time box, type or select the number of days you want to keep the data in the table.
- 2 In the Summarization Interval box, select a time characteristic for the data in the table. For example, a 5 minute table contains data for a 5 minute interval. Data collected at 5 minute intervals is called rate data.
- 3 In the Table Type box, select a table type.

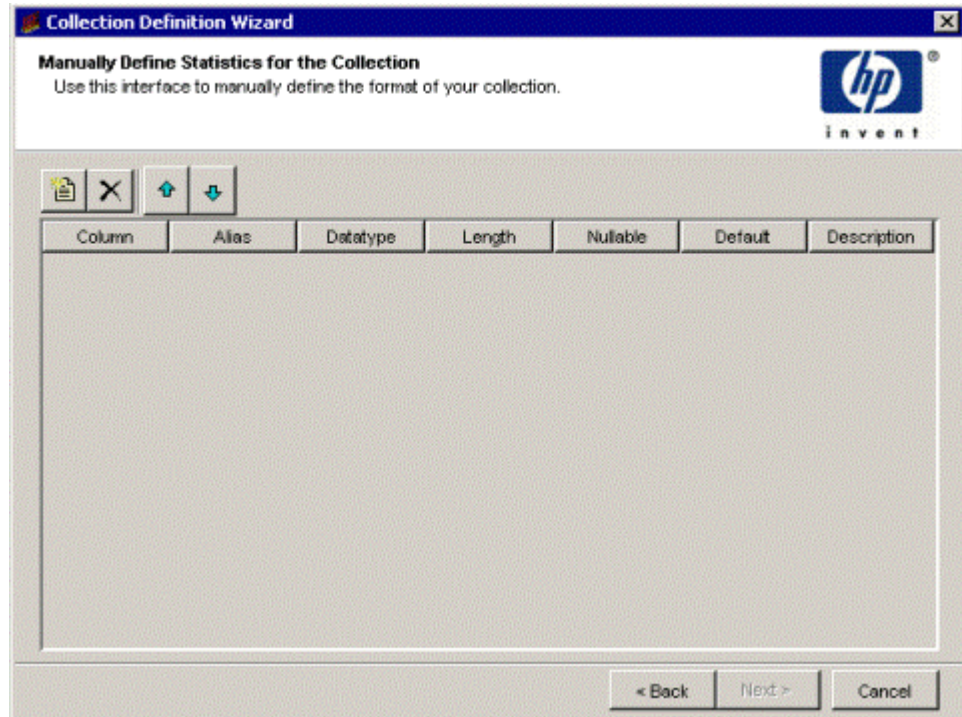
The table type specifies the class of data contained in the table. Following is the list of available table type options:


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.

- 4 Click **Next**.

Task 4: Define the statistics for the collection.

The Manually Define Statistics for the Collection window opens.



- 1 Click  to create a column to add to the table.
- 2 To change the column name, click in the Column box, and type a new name for the statistic.
The name cannot exceed 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.
- 3 To change the alias name of the statistic, click in the Alias box, and type a new alias name.
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.
For the available datatype options and their corresponding database storage types, see [step 6](#) on page 189.




- 5 To change the length, double-click in the Length box, and type a new value.

Valid data types for length values are `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 To change the default, click in the Default box, and type a new value.
The length of the value cannot exceed 255 characters.
- 8 To change the description, click in the Description box, and type a new description.

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 window opens.

- 1 Do one of the following:
 - Select **Create a new property table**, and specify a name for the property table that describes the data it will collect. The name must be unique and cannot be greater than 30 characters. Valid characters are alphabetic characters, numeric digits, and underscores (`_`). Go to [Task 6](#) on page 245.
 - Select **Reuse existing property table**, and then select a property table from the list. Go to [Task 7](#) on page 250.



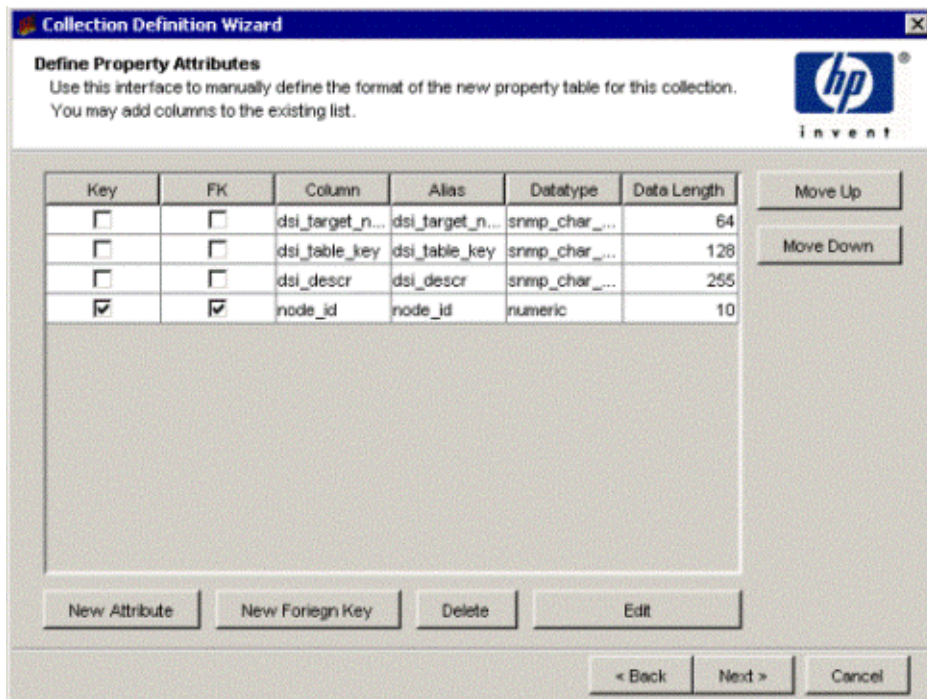
If you reuse an existing property table, you may corrupt data that is already stored in that table.

2 Click **Next**.

Task 6: Define Property Attributes

You can create a property table with the existing columns or you can add more columns.

The Define Property Attributes window opens.



The following options are provided:

- Add a column or add a column with a foreign key to a property table and set up the column map directives in TEEL.
See [Adding a Column](#) on page 246 and [Adding a Column with a Foreign Key](#) on page 248.
- Change the order of columns.
See [Changing the Order of Columns](#) on page 249.
- Delete columns.
See [Deleting a Column](#) on page 249.

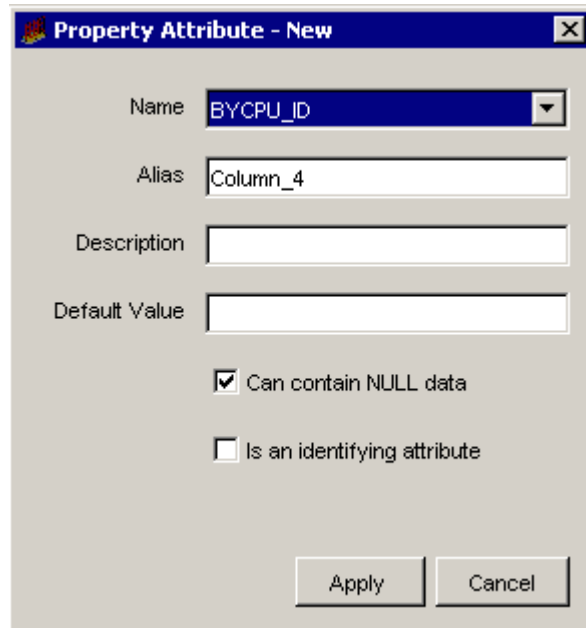
- Edit columns.
See [Editing a Column](#) on page 249.

Adding a Column

To add a column to the property table, follow these steps:

- 1 To add a column, click **New**.

The New Attribute dialog box opens.



The screenshot shows a dialog box titled "Property Attribute - New". It has a blue title bar with a close button. The dialog contains the following fields and options:

- Name:** A dropdown menu with "BYCPU_ID" selected.
- Alias:** A text box containing "Column_4".
- Description:** An empty text box.
- Default Value:** An empty text box.
- Can contain NULL data:** A checked checkbox.
- Is an identifying attribute:** An unchecked checkbox.
- Buttons:** "Apply" and "Cancel" buttons at the bottom.

- 2 Click the Name list, and select a name for the column. By default, the names of the OID objects you selected will appear.
- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 In the Datatype list, click a datatype value.

For the available datatype options and their corresponding database storage types, see [step 6](#) on page 189.

- 7 In the Data Length list, click a length value.

Valid data types for length values are `char_string`, `hex_string`, `numeric`, `octet_string`, `snmp_char_string`, `snmp_hex_string`, or `snmp_octet_string`.

- 8 Select or click to clear the **Can contain NULL data box**.

When you select this option, it indicates that the column can contain NULL values. If the **Is an identifying attribute** box is checked, this option will automatically be disabled.

- 9 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

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 **Is an identifying attribute** box must not be selected and there should be no value in the Default Value box.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#) on page 249.

- 10 Click **Apply**.

The column is added to the property table.

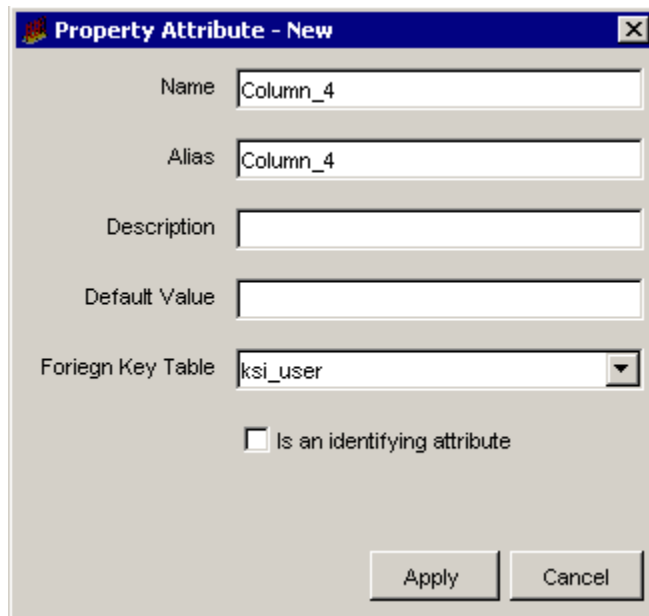
- 11 Repeat the procedure to add more columns, or click **Next**. to continue.

Adding a Column with a Foreign Key

To add a column with a foreign key:

- 1 To add a column, click **New**.

The New Foreign Key dialog box opens.



The screenshot shows a dialog box titled "Property Attribute - New". It has a close button in the top right corner. The dialog contains the following fields and controls:

- Name:** A text box containing "Column_4".
- Alias:** A text box containing "Column_4".
- Description:** An empty text box.
- Default Value:** An empty text box.
- Foreign Key Table:** A dropdown menu with "ksi_user" selected.
- Is an identifying attribute:** An unchecked checkbox.
- Buttons:** "Apply" and "Cancel" buttons at the bottom right.

- 2 Type a name for the column in the Name box.

The column name cannot exceed 30 characters and can contain the following characters: alphabetic, numeric, or an underscore.

- 3 Type an alias name for the column in the Alias box.

The name cannot exceed 255 characters and can contain the following characters: alphabetic, numeric, an underscore, or a hyphen.

- 4 Type a description for the table in the Description box.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 If the column should have a default value, type a value in the Default Value box.

A default value is optional. 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.

- 6 Click the Foreign Key Table list and select the property table that has the key.
- 7 Select the **Is an identifying attribute** box if you want the column to be a by-variable.

If you specify any other column other than `dsi_target_name` as a by-variable, you may not specify `dsi_table_key` as a by-variable. However, you may specify `dsi_table_key` as a by-variable only if you specify `dsi_target_name` as the first 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. To change the order of the columns in a property table, see [Changing the Order of Columns](#) on page 249.

- 8 Click **Apply**.
The column is added to the property table.
- 9 Repeat the procedure to add more columns, or click **Next** to continue.

Changing the Order of Columns

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 **Move Up**.
- To move the column down, click **Move Down**.

Deleting a Column

To delete a column from the property table, click **Delete**.

Editing a Column

To edit a column:

- 1 Double-click the column in the Define Property Attributes window.

The Property Attribute Edit dialog box opens, which is identical to the Property Attribute New or Property Attribute New Foreign Key dialog boxes.


- 2 Make any changes.
- 3 Click **Apply**.

Repeat the procedure, or click **Next** to continue.

Task 7: Specify an output file for the collection.

The Select the Output File window opens (page 194).

- 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.

Navigate to the folder where you want to store the output file, type the file name in the File name box, 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 window opens (page 195).

- 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 type the following command at the command line:

```
datapipe_manager -p create -a teal_file
```

In this instance, *teal_file* is the name of the output file from.

If you need additional information about the options for the **datapipe_manager** command, see the *HP Performance Insight TEEL Reference Guide* and the *HP Performance Insight Reference Guide*.

- 2 In the Summarizations box, select either **Copy Data Directly to Rate Tables** or **No Summarizations**.



Additionally, if you want to do hourly or weekly summarizations, you must create a roll-up file (.sum file) and execute it using the `trend_sum` command. See the *HP Performance Insight Reference Guide* for details about using the `trend_sum` command.

- 3 Click **Next**.

Task 9: Review the Summary window.

The Summary dialog box opens ([page 196](#)).

- 1 Verify that all the collection property values are correct.

If necessary, click **Back** to return to a previous window to make a change. After you make a change, click **Next** until you return to the Summary window.

- 2 Click **Finish** to create the data table.

The data table creation wizard closes, and PI creates a TEEL file and stores it in the designated location.

- 3 Create a polling policy for the collection. See [Creating Polling Policies](#) on page 252 for details.

Creating Polling Policies

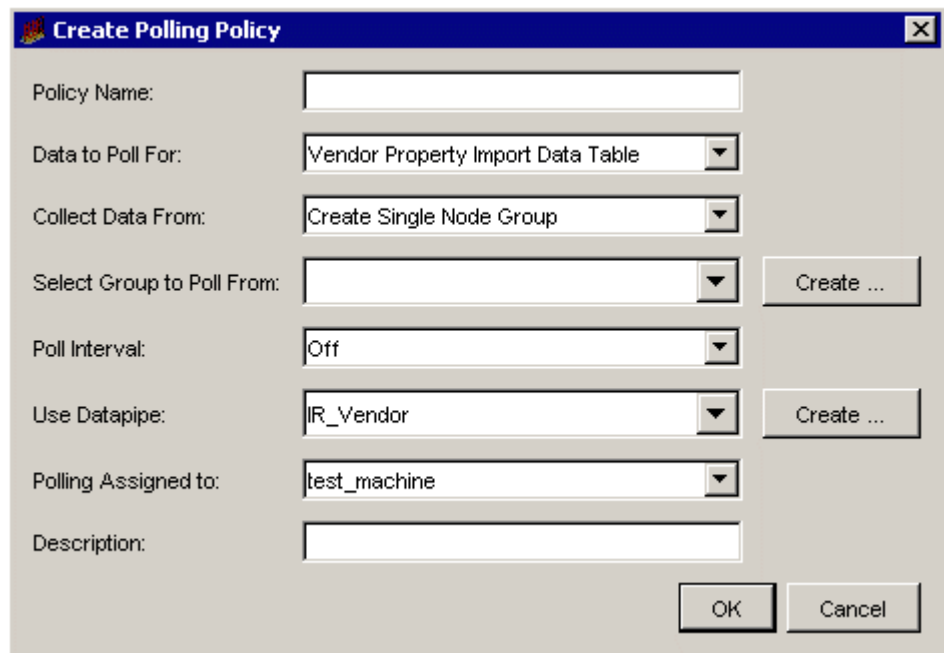
When you install a report pack, it provides the polling policies that it needs; however, there may be situations where you may need to create a new polling policy.

Each polling policy requires a *collection* and a *polling group*. The collection identifies the table in the database that will contain the collected data. The polling group identifies the managed object you want to poll. A polling group comprises a list of nodes or a list of instances. Each node requires an SNMP profile and a community string profile.

To create a polling policy, follow these steps:

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

The Create Polling Policy dialog box opens.



- 2 Type 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. This specifies the collection.

The tables that appear in this list are located in the database. The name of the current database appears in the navigation pane of the Polling Policy Management window.

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 128 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.

If you want to create a new collection, see [Creating Collections](#) on page 173 for details.

- 4 From the Collect Data From list, select one of the following from which to collect the data:
 - **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 260 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.
 - 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 285 for details.

- 9 In the Description box, type a text description of the polling policy.
- 10 Click **OK**.

The polling policy is created.

Example: Creating a Polling Policy for an HP Performance Agent Collection

Suppose you created an HP Performance Agent collection called `ovpa_northeast`. The following example shows how to create a polling policy for that collection.

To create a polling policy, follow these steps:

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

The Create Polling Policy dialog box opens.

- 2 Type a distinct and unique name for the polling policy in the Policy Name box like `ovpa_northeast_policy`.
- 3 From the Data to Poll For list, select `ovpa_northeast`.



The policy name and the collection name do not have to be the same.

- 4 From the Collect Data From list, select **All Nodes of the Same Type**.

This option 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.

- 5 Select **OVPA** from the Select Group to Poll From list.
- 6 From the Poll Interval list, select the polling interval. This specifies the amount of time between polling requests.
- 7 From the Use Datapipe list, select the `dsi_sr`. This datapipe should be the default. This specifies the method by which PI will collect the data.
- 8 From the Polling Assigned to list, select the server that will collect the data.
- 9 In the Description box, type a text description of the polling policy.
- 10 Click **OK**.

The polling policy is created.

Verifying that the Policy is Polling

You can verify that your policy is polling by doing the following:

- 1 Execute the `pa_collect` command. For example, to execute the entries in the polling policy with an interval value of 5, you would type the following:

```
pa_collect -i 5
```

See the *HP Performance Insight Reference Guide* for more information about the `pa_collect` command.

- 2 Use Table Manager to locate the property table for the collection and see if any rows are present.

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, follow these steps:

- 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 select **Polling Policy** from the Edit menu.

The Edit Polling Policy dialog box opens.

Edit Polling Policy

Policy Name: 2007-04-10_19:44:31:468_833

Data to Poll For: xDPVpnInterfaceConf

Collect Data From: All Nodes of the Same Type

Select Group to Poll From: MPLS_VPN Create ...

Poll Interval: 1 Hour

Use Datapipe: dpipe_snmp Create ...

Polling Assigned to: test_machine

Description:

OK Cancel

- 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 260 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.
 - 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 285 for details.
- 3 Click **OK**.

Deleting a Polling Policy

To delete one or more polling policies, follow these steps:

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

See [Selecting Items](#) on page 170 for details.

- 2 From the Edit menu, select **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.

Managing Polling Groups

To create, edit, or delete a polling group, follow these steps:

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

The Edit Polling Groups dialog box opens.

Edit Polling Policy

Policy Name: 2007-04-10_19:44:31:468_833

Data to Poll For: xDPVpnInterfaceConf

Collect Data From: All Nodes of the Same Type

Select Group to Poll From: MPLS_VPN

Poll Interval: 1 Hour

Use Datapipe: dpipe_snmp

Polling Assigned to: test_machine

Description:

- 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 261 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 270.
 - 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.

The name of the polling group must be unique across all types of groups. You cannot create two types of polling groups with the same name.

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.

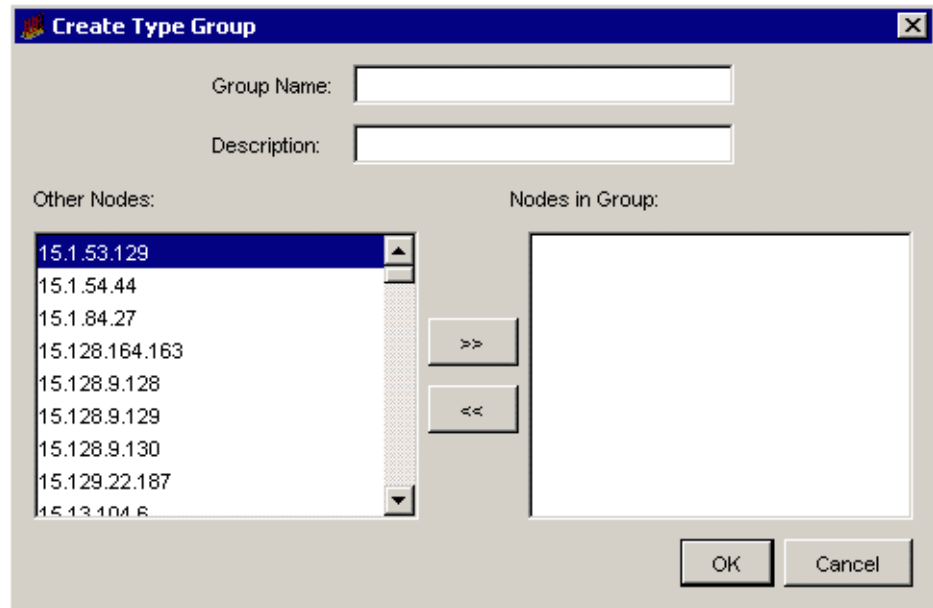


To create a single node group, follow these steps:


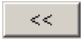
- 1 Type the name of the group in the Group Name box. This name can be a maximum of 30 characters.
- 2 Type a description for the group in the Description box. The length of the description can be 255 characters.
- 3 Select one node from the Select A Node box.
- 4 Click **OK**.

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.



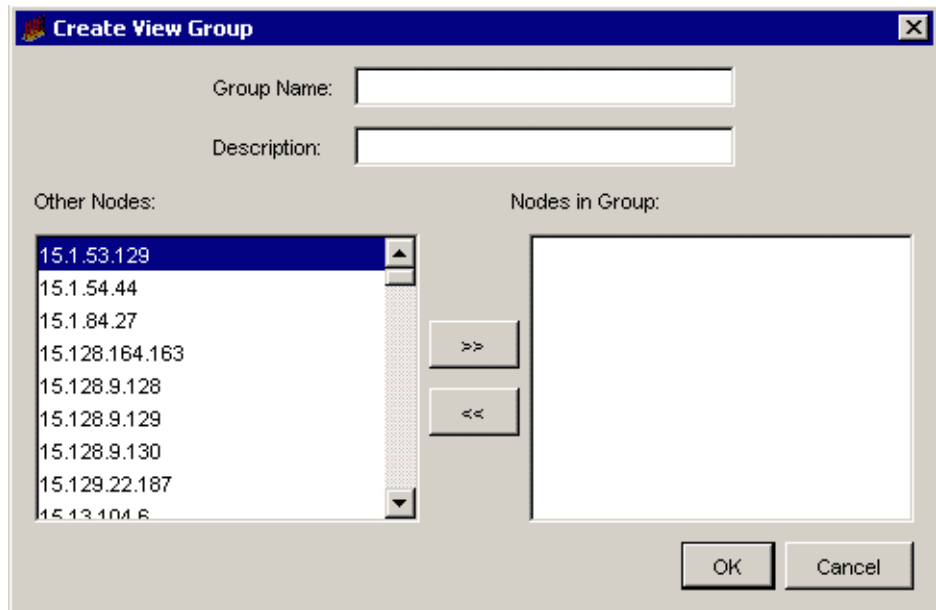
To create a group that contains nodes of the same type, follow these steps:

- 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.

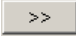
- Double-click a node in the Nodes in Group box. The node moves to the Other Nodes box.
- 5 When you are finished, click **OK**.


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.



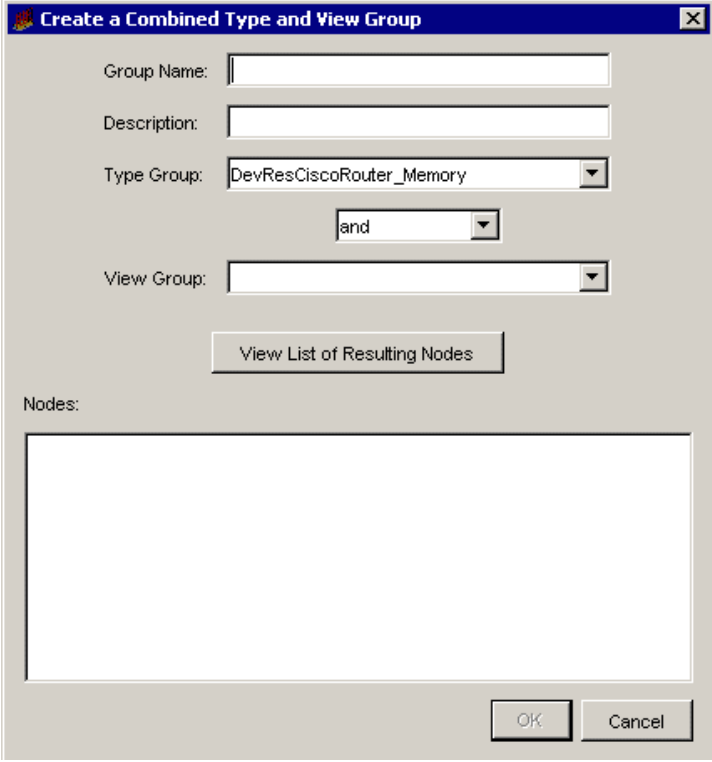
To create a group that has all nodes in the same view, follow these steps:

- 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 a 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.



Create a Combined Type and View Group

Group Name:

Description:

Type Group:

View Group:

Nodes:

To create a group of nodes that comprises a combination of type and view node groups, follow these steps:

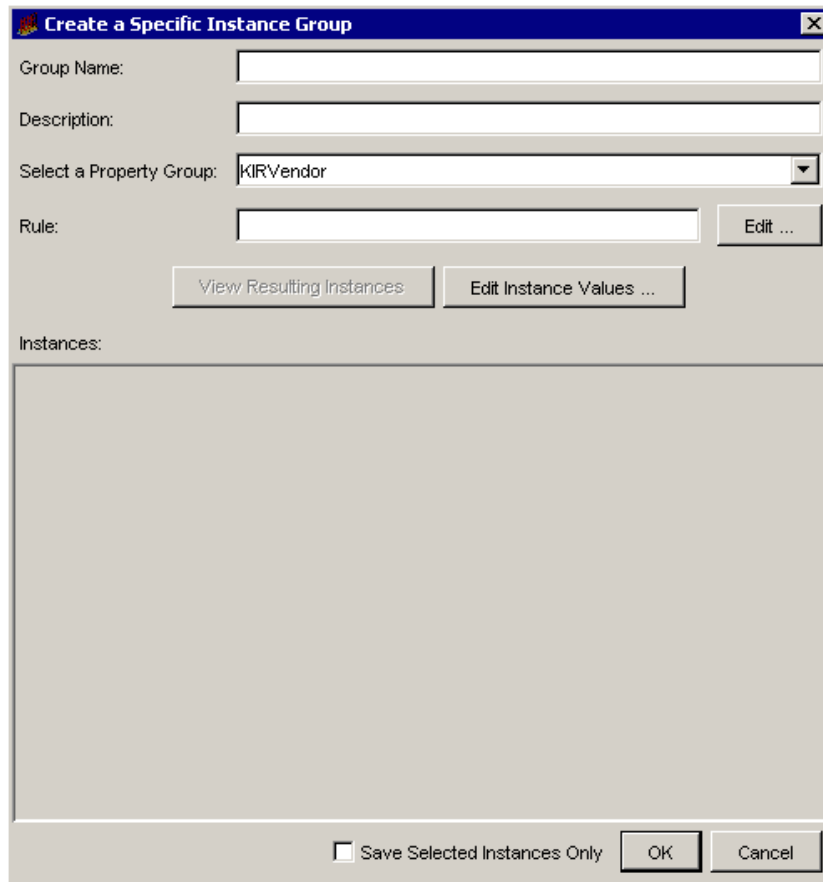
- 1 Type the name of the group in the Group Name box. This name can be a maximum of 30 characters.
- 2 Type a description for the group in the Description box. The length of the description can be 255 characters.
- 3 Select a group of nodes from the Type Group list.
- 4 Select a group of nodes from the 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, 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, 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, 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, the resulting group would have nodes 4, 5.
- 6 Click **View List of Resulting Nodes** 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 dialog box titled "Create a Specific Instance Group" features the following elements:

- Group Name:** A text input field.
- Description:** A text input field.
- Select a Property Group:** A dropdown menu with "KIRVendor" selected.
- 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 for displaying instance data.
- Footer:** A checkbox labeled "Save Selected Instances Only" and "OK" and "Cancel" buttons.

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

To create a specific instance group, follow these steps:

- 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.

The screenshot shows a dialog box titled "Create Rule" with a close button (X) in the top right corner. The main area is titled "Define Rule". It contains two rows of input fields. The first row has "Select Column:" followed by a dropdown menu showing "DSI_KEY_ID", an equals sign operator dropdown, and an empty text box. Below this is an "and" operator dropdown. The second row is identical to the first. At the bottom right are "OK" and "Cancel" buttons.

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 170 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 170 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 box.

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

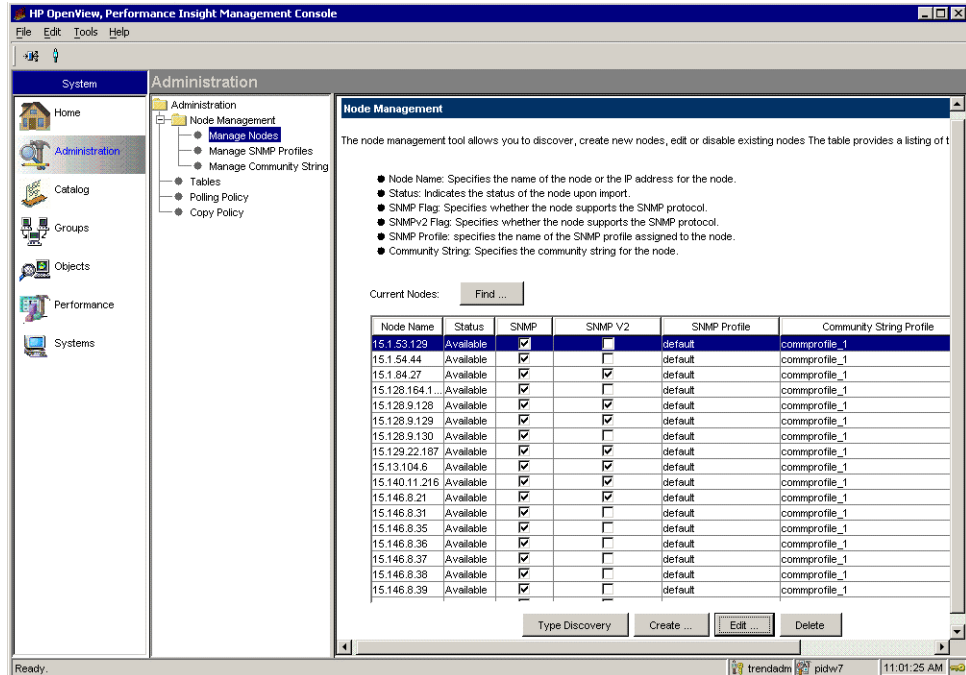
When the Edit a Specific Instance Group dialog box opens, the list of instances does not appear until you click **View Resulting Instances** unless you saved the list of instances by clicking the **Save Selected Instances Only** box and not the rule. If there are a lot of instances, you may have to wait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

Managing Nodes

To create, edit, or delete nodes, follow these steps:

- 1 In the navigation tree provided in the Administration snapin, expand the Node Management folder.
- 2 Click **Manage Nodes**.

The Node Management window 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.

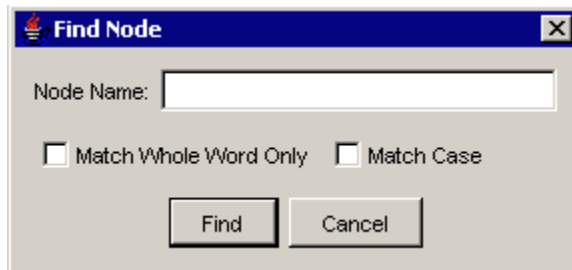
- **Community String Profiles.** This is the name of the Community String profile for the node.
- 3 To create a new node, see [Creating Nodes](#) on page 272.
 - 4 To edit an existing node, see [Editing Nodes](#) on page 274.
 - 5 To delete an existing node, see [Deleting Nodes](#) on page 274.
 - 6 When you are finished, click **Close**.

Finding a Node

To locate an existing node from the list, follow these steps:

- 1 In the Node Management window, click **Find**.

The Find dialog box opens.



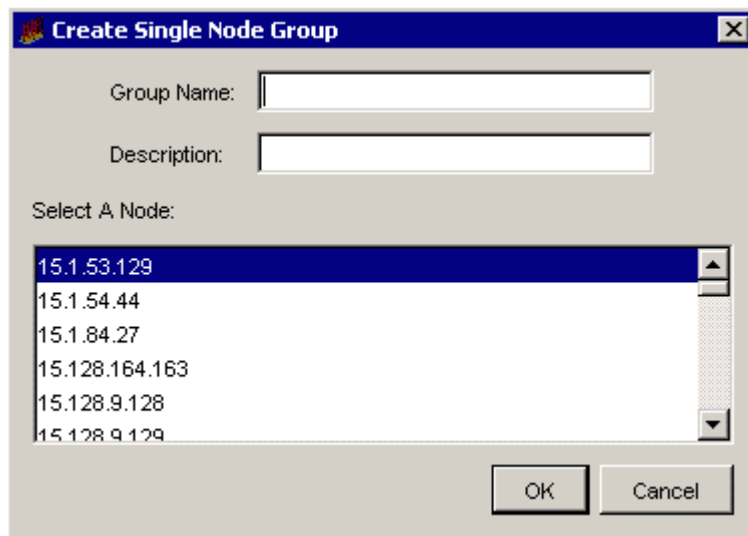
- 2 Type the name of the node that you are looking for in the Node Name box.
- 3 If you want to only locate complete words in the name instead of partial words, select the Match Whole Word Only check box.
- 4 If you want to locate names that match the case you specified in the Node Name box instead of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

Creating Nodes

To create a node, follow these steps:

- 1 From the Node Management window, click **Create**.

The Create Single Node Group dialog box opens.



2 Type the name for the node in the Node Name box. This name can be a maximum of 64 characters.

3 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 282 for details.

4 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 279 for details.

5 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.

6 Select the **SNMP Node** box if the node supports SNMP v1 data.

7 Select the **SNMP V2 Node** box if the node supports SNMP v2 data.

- 8 Click **OK**.

Editing Nodes

To edit an existing node, follow these steps:

- 1 From the Node Management window, 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 3](#).

If you want to find a node in the list of existing nodes, see [Finding a Node](#) on page 272 for details.

- 2 Click **Edit**.

The Edit Node dialog box opens.

- 3 Modify any value except the value in the Node Name box. See [Creating Nodes](#) on page 272 for details.

- 4 Click **OK**.

Deleting Nodes

To delete one or more existing nodes, follow these steps:

- 1 From the Node Management window, select one or more nodes that you want to delete by highlighting the corresponding rows. See [Selecting Items](#) on page 170 for details.

If you want to find a node in the list of existing nodes, see [Finding a Node](#) on page 272 for details.

- 2 Click **Delete**.

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

- 3 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.

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. The following table 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

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

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, follow these steps:

- 1 In the navigation tree provided in the Administration Snapin, click **Polling Policy**.
- 2 From the File menu, select **Import Nodes**.
The Import from dialog box opens.
- 3 Navigate to the file that contains the nodes you want to import.
- 4 Double-click the file name, or select the name and click **Open**.

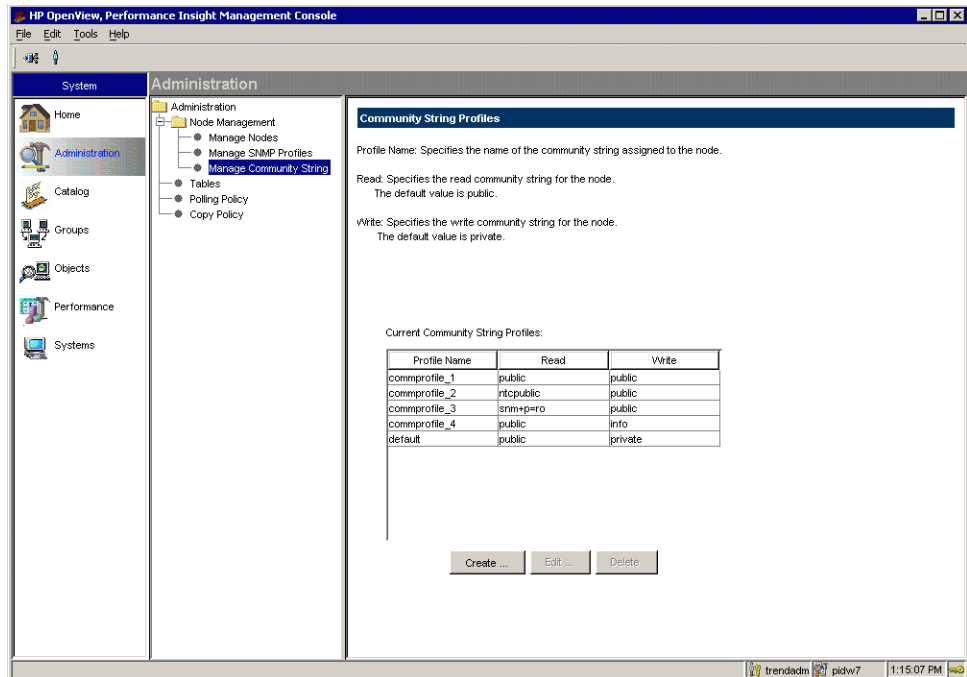
Alternatively, you can use the `node_manager` command to import or export nodes. For more information about the `node_manager` command, see the *HP Performance Insight Reference Guide*.

Managing Community String Profiles

To view a list of the existing community string profiles with the name of the profile and the corresponding read and write permissions, follow these steps:

- 1 In the navigation tree provided in the Administration Snapin, expand the Node Management folder.
- 2 Click **Manage Community String**.

The Community String Profiles window opens.



You can do the following:

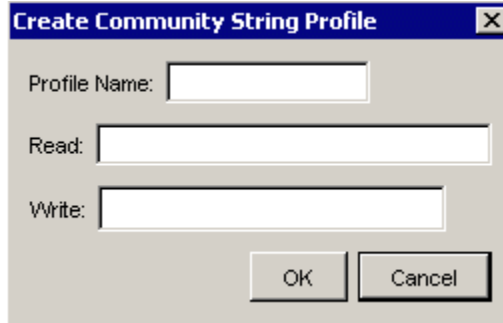
- Create a community string profile.
- Edit a community string profile.
- Delete a community string profile.

Creating a Community String Profile

To create a community string profile, follow these steps:

- 1 From the Community String Profiles window, click **Create**.

The Create Community String Profile dialog box opens



The image shows a dialog box titled "Create Community String Profile". It has a standard Windows-style title bar with a close button (X) in the top right corner. The dialog contains three text input fields: "Profile Name:", "Read:", and "Write:". Below the input fields are two buttons: "OK" and "Cancel".

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

Editing a Community String Profile

To edit a community string profile, follow these steps:

- 1 From the Community String Profiles window, select the community string profile you want to modify.
- 2 Click **Edit**. The Edit Community String Profile dialog box opens.



- 3 Modify the permission values in the Read and Write fields.



You cannot change the name of the profile.

- 4 Click **OK**.

Deleting a Community String Profile

To delete a community string profile, follow these steps:

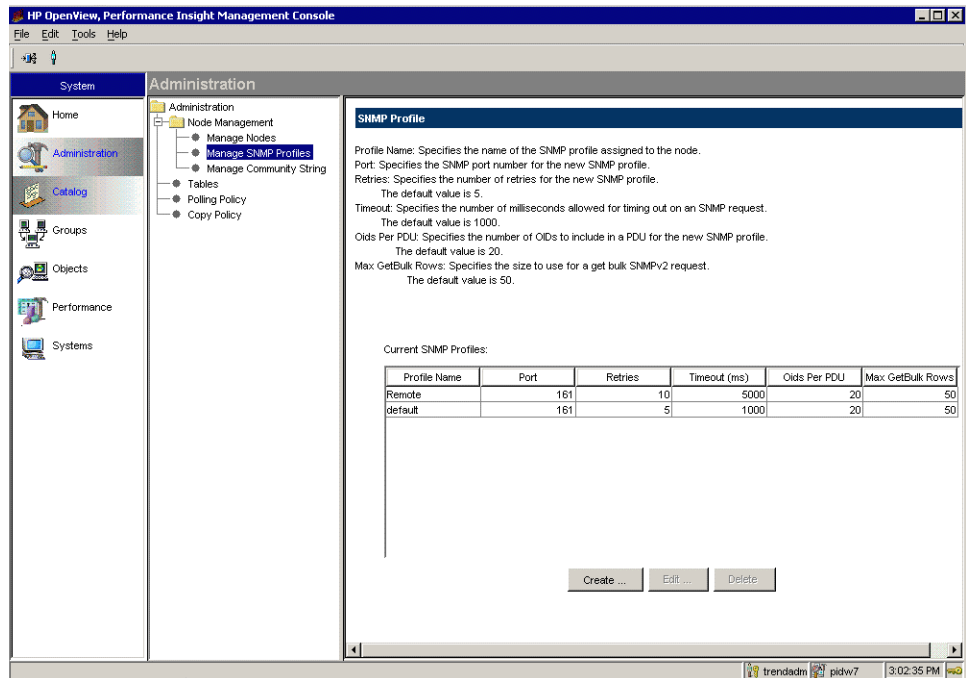
- 1 From the Community String Profiles window, select the community string profile you want to delete.
- 2 Click **Delete**.
A deletion confirmation window displays.
- 3 Click **Yes**.

Managing SNMP Profiles

To view the existing SNMP profiles with names and their corresponding settings, follow these steps:

- 1 In the navigation tree provided in the Administration Snapin, expand the Node Management folder.
- 2 Click **Manage SNMP Profiles**.

The SNMP Profile window opens..



You can do the following:

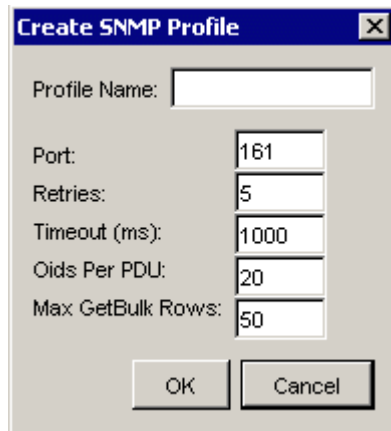
- Create a community string profile.
- Edit a community string profile.
- Delete a community string profile.

Creating an SNMP Profile

To create an SNMP profile, follow these steps:

- 1 From the SNMP Profile window, click **Create**.

The Create SNMP Profile dialog box opens.



The screenshot shows a dialog box titled "Create SNMP Profile" with a close button (X) in the top right corner. The dialog contains the following fields and values:

Profile Name:	<input type="text"/>
Port:	<input type="text" value="161"/>
Retries:	<input type="text" value="5"/>
Timeout (ms):	<input type="text" value="1000"/>
Oids Per PDU:	<input type="text" value="20"/>
Max GetBulk Rows:	<input type="text" value="50"/>

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

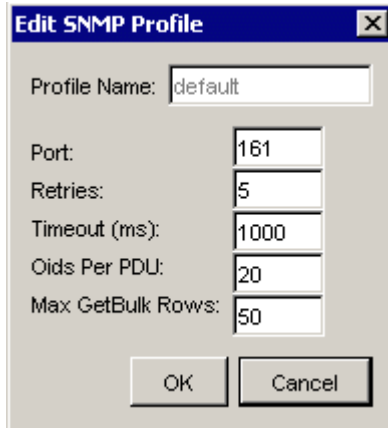
- 2 Type the name for the profile in the Profile Name box. This name can be a maximum of 64 characters.
- 3 Type the value for the port number in the Port box.
- 4 Type the number of retries allowed in the Retries box.
- 5 Type the number of milliseconds to wait for a timeout in the Timeout (ms) box.
- 6 Type the maximum number of OIDs that the agent will support for each PDU in the Oids Per PDU box.
- 7 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.
- 8 Click **OK**.

Editing an SNMP Profile

To edit an SNMP profile, follow these steps:

- 1 From the SNMP Profile window, select the SNMP profile you want to modify.
- 2 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:	default
Port:	161
Retries:	5
Timeout (ms):	1000
Oids Per PDU:	20
Max GetBulk Rows:	50

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

- 3 Modify the values for the settings that you want to change.
You cannot change the name of the profile.
- 4 Click **OK**.

Deleting an SNMP Profile

To delete an SNMP profile:

- 1 From the SNMP Profile window, select the SNMP profile you want to delete.
- 2 Click **Delete**.
A deletion confirmation window displays.
- 3 Click **Yes**.

Managing Datapipe Installations


To view the existing datapipes with a corresponding list of installations for each datapipe, follow these steps:

- 1 In the navigation tree provided in the Administration Snapin, click **Polling Policy**.

The Polling Policy window opens.

- 2 From the Edit menu, select **Datapipe Installations**.

The Edit Datapipe Installations dialog box opens.



You can do the following:

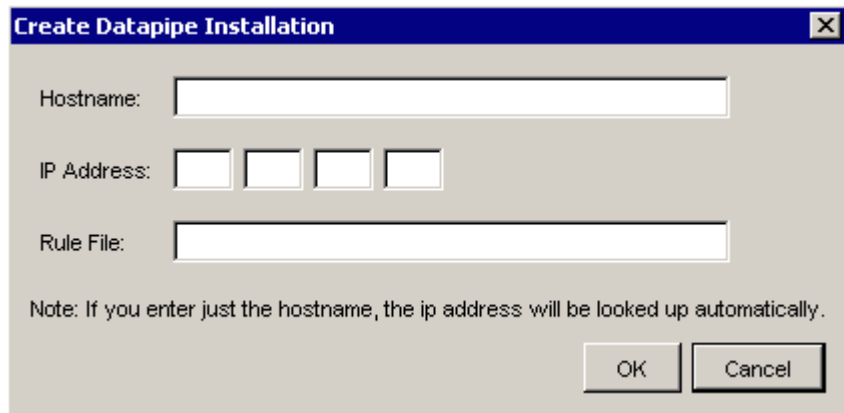
- Create a community string profile.
- Edit a community string profile.
- Delete a community string profile.

Creating a Datapipe Installation

To create a datapipe installation:

- 1 From the Edit Datapipe Installations dialog box click **Create**.

The Create Datapipe Installation dialog box opens.



The screenshot shows a dialog box titled "Create Datapipe Installation". It features three input fields: "Hostname" (a single text box), "IP Address" (four separate boxes for each octet), and "Rule File" (a single text box). Below the fields is a note: "Note: If you enter just the hostname, the ip address will be looked up automatically." At the bottom right are "OK" and "Cancel" buttons.

- 2 Type the hostname for the collection station in the Hostname box.

If you type 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 type a hostname, it is the name of the collection station.

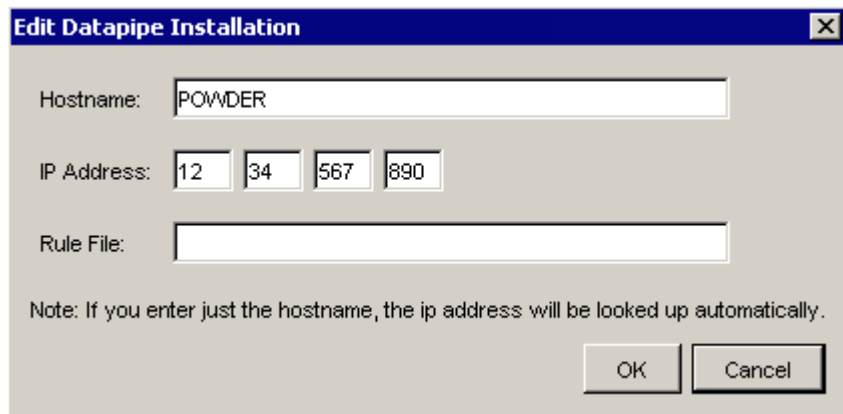
If the system cannot locate the IP address for the specified hostname, a window displays to confirm that it could not find the address.

- 3 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 type only the IP address, then the address is the name of the collection station.
- 4 Type the name of the file that contains the instructions for the collection station in the Rule File box. This box can remain blank.
- 5 Click **OK**.

Editing a Datapipe Installation

To edit a datapipe installation:

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



Edit Datapipe Installation [X]

Hostname: POWDER

IP Address: 12 34 567 890

Rule File:

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

OK Cancel

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

Deleting a Datapipe Installation


To delete a datapipe installation:

- 1 From the Edit Datapipe Installations dialog box, select the datapipe to delete from the Select Datapipe list.
- 2 Select the corresponding installation from the Current Datapipe Installations list.
- 3 Click **Delete**.
A deletion confirmation window displays.
- 4 Click **Yes**.

Modifying Property Table Values

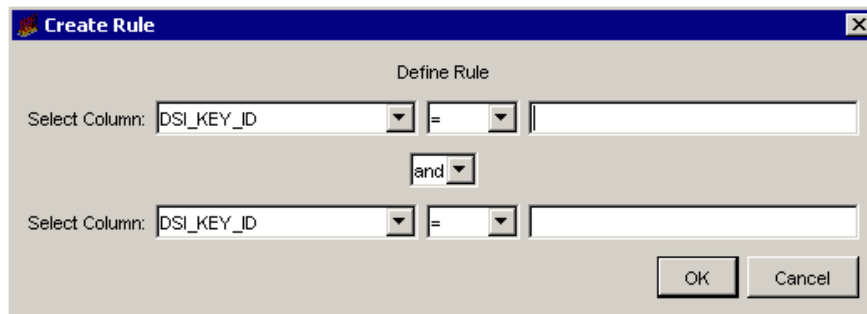
To modify any value in a property table:

- 1 Click the **Administration** icon in the navigation pane.
The Administration Snapin opens.
- 2 In the navigation tree provided in the Administration Snapin, click **Polling Policy**.
- 3 From the Edit menu, select **Descriptions**.
The Edit Descriptions dialog box opens.



- 4 Select the property table that contains the values you want to modify from the Select a Property Group list.

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




- 6 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.

- 7 Click **View Resulting Instances** to see all the columns for the selected rows in the property table.
- 8 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.
- 9 When you are finished, click **Close** in the Edit Descriptions dialog box.

9 Internal Monitoring of PI using Performance Graphs

PI Performance Graphs is an application that is available from the PI Management Console. It uses a set of graphs to show the performance of PI polling. It allows you to determine if a PI system is functioning normally or is overloaded.

You can view the performance graphs to see the amount of data collected from nodes and also the duration of each poll cycle.

There are three types of graphs that show the performance of polling in a PI system:

- **Process Duration:** To view the duration graphs for PI polling components
- **Polling Health:** To view the successfulness of the polling
- **Row Counts:** To view graphs of row counts by data table by polling cycle

Starting HP Performance Insight Performance Graphs

To start the PI Performance Graphs application:

- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 for details.

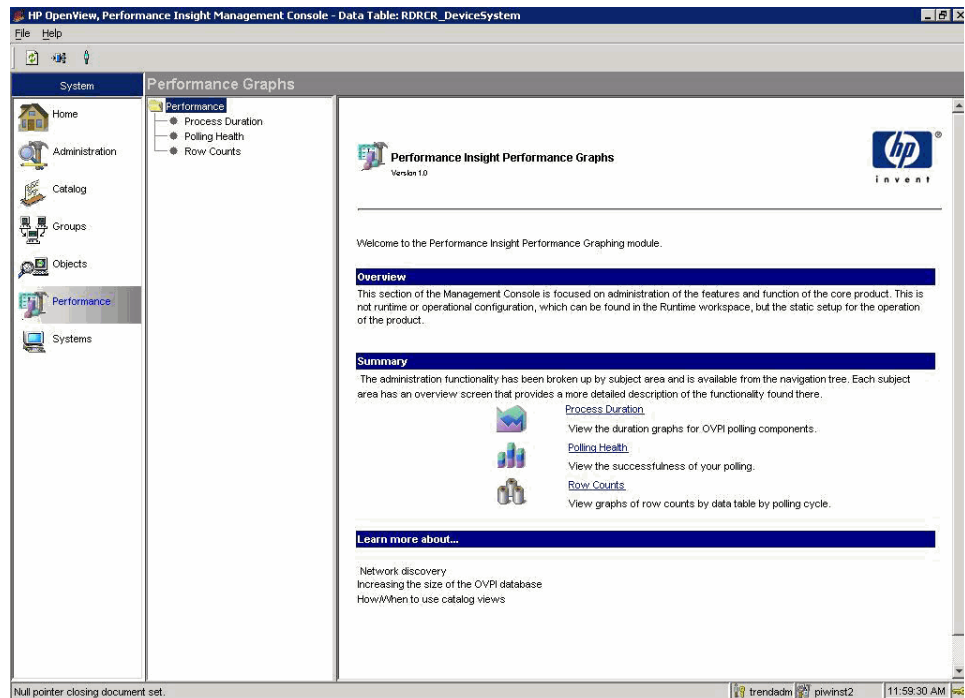


You must log on to the Management Console using an account that has Administrator privileges.

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

The OVPI Performance Graphs window opens in the Management Console. The left side of the window displays a navigation tree, which lists the three options for retrieving performance graphs.

The following figure shows the OVPI Performance Graphs window:



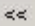
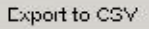
Viewing HP Performance Insight Performance Graphs

On the OVPI Performance Graphs window, click one of the following on the navigation tree:

- [Process Duration](#)
- [Polling Health](#)
- [Row Counts](#)

The data metrics required in any of the Performance Graphs are as follows:

- 1 **Collector:** The program that collects the data from sources outside of PI. You can select any one of the following collectors:
 - **MW:** You can use the MW option to poll nodes for SNMP data on a PI system. See the *HP Performance Insight Reference Guide* for more information.
 - **EE:** You can use the EE option to import data from a flat file into a datapipe on a PI system. It uses the instructions from a TEEL file associated with it to import the data. See the *HP Performance Insight TEEL Reference Guide* for more information.
 - **SR:** You can use the SR option to collect data from the HP Performance Agent or the HP Operations agent on a PI system. The command used is `pa_collect`. See the *HP Performance Insight Reference Guide* for more information.
- 2 **Polling Interval:** The time interval in minutes for the polling.
- 3 **Data Table:** The table into which the polled data from the nodes are stored.

The Start Date for any Performance Graph is the filter for start of data retrieval. It is midnight (00:00:00) of any given day. Similarly, the End Date is the filter for end of data retrieval. It is 11:59:59 for any given day. You can view graphs for the last seven days (including the current day) by clicking the day navigation bar  below the graph. You can click  to export data from a Performance Graph to a Comma Separated Values (.csv) file.

Process Duration

Process or polling duration data illustrates total polling time for a particular data table. These graphs should be uniform. Use the Process duration graphs to view the amount of time spent collecting data for a specific data table.

Using a combination of time range, polling interval, collector type, and data table, you can retrieve a dataset of process duration times. By creating a line graph using these parameters, you can see how close they are coming to the edges of their collection windows as well as any major variations to the processing streams, which should, overall, be fairly flat.



A graph with large swings could denote large number of systems being down or database issues.

Create Process Duration Graphs

To create a Process Duration Graph, perform the following steps:


- 1 In the OVPI Performance Graphs window, on the Performance navigation tree, click **Process Duration**. The Performance - Polling Duration Data window opens.
- 2 From the Collector drop-down list, select MW, SR or EE Collector.
- 3 From the Select Polling Interval drop-down list, select the required interval.
- 4 From the Select Data Table drop-down list, select the data table for which process duration is to be calculated. Choosing All Policies for the data table will show a true process time for the mw_collect, pa_collect, or ee_collect process.
- 5 Click **Graph It**.

A graph is drawn between Collection times and Duration in seconds.

The red line represents the poll interval selected. If the graph crosses this line, your PI system may be missing polls as the previous poll is running longer than its expected duration. Performance tuning may be necessary if this occurs.

You may only graph data for collectors that have been used in the past seven days.

The following figure shows a sample PI Process Duration Graph:



Polling Health

Polling Health data graphs the total number of nodes that responded to your collection of data from SNMP and HP Performance Agent agents. A large number of failures in a given poll cycle can denote network outages or problems.

The polling health graphs draw a picture of how well your system is polling for the policies that you have set up.

Create Polling Health Graphs

To create a Polling Health Graph, perform the following steps:

- 1 In the OVPI Performance Graphs window, on the Performance navigation tree, click **Polling Health**. The Performance - Polling Health Data window opens.
- 2 From the Collector drop-down list, select MW, SR or EE Collector.
- 3 From the Select Polling Interval drop-down list, select the required interval.
- 4 From the Select Data Table drop-down list, select the data table for which Polling Health is to be calculated.
- 5 Click **Graph It**.


A graph is drawn between time and number of nodes.

The bar graph is a stack of successful (green) and unsuccessful (red) polls for each node in the polling group over the course of a single day. The successful nodes are shown in green and unsuccessful nodes in red.

You can use the buttons at the bottom of the graph to move forward and backward in one day increments to view the latest seven days worth of data.

You may only graph data for collectors that have been used in the past seven days.

The following figure shows a sample PI Polling Health Graph:



Row Counts

The row counts graphs show the volume of the data by polling period for a particular data table. These are useful in showing where your peak collection times are and in determining database size information.

This can be helpful to visualize your busiest times or to help troubleshoot when report graphs are devoid of data.

Create Row Counts Graphs

To create a Row Counts Graph, perform the following steps:

- 1 In the OVPI Performance Graphs window, on the Performance navigation tree, click **Row Counts**. The Performance - Row Count Data window opens.
- 2 From the Collector drop-down list, select MW, SR or EE Collector.
- 3 From the Select Polling Interval drop-down list, select the required interval.
- 4 From the Select Data Table drop-down list, select the data table for which Row Counts is to be calculated.
- 5 Click **Graph It**.

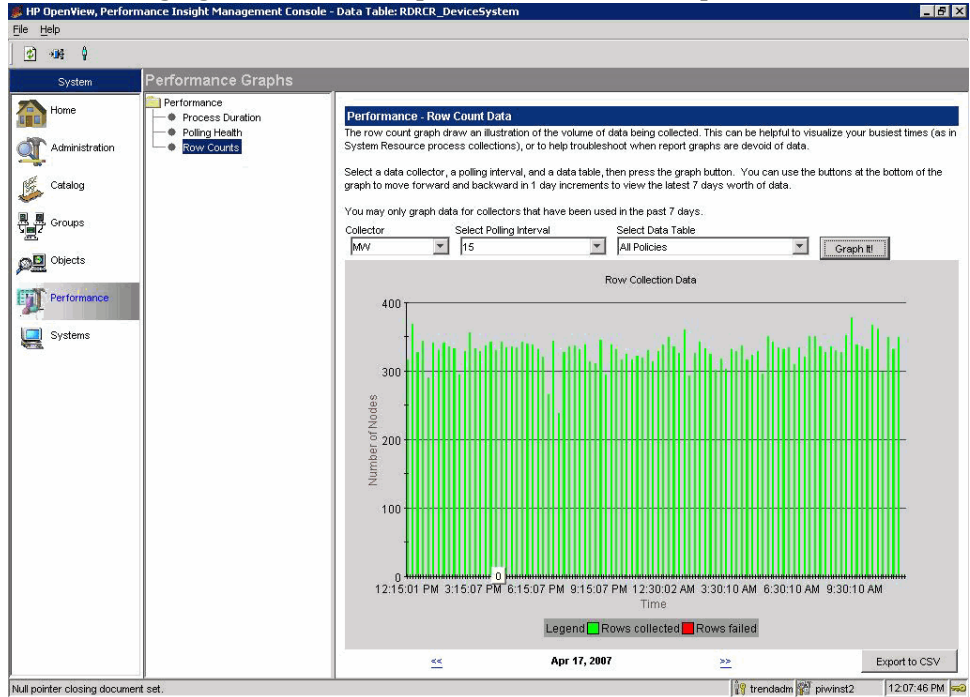
A graph is drawn between time and number of nodes.

You can use the buttons at the bottom of the graph to move forward and backward in one day increments to view the latest seven days worth of data.



The values displayed as Rows Collected will not always reflect the number of rows saved in the database. This is because in certain instances the collectors will retrieve data records that represent duplicate information from the performance agents running on polled nodes. HP Performance Insight is designed to discard duplicate records before inserting them into the database.

The following figure shows a sample PI Row Counts Graph:



10 Managing Objects

Object Manager, an application that is available from the Management Console, lets you locate and select managed objects, access object-specific property management forms and report pack configuration forms, and generate and view reports that are relevant to a selected object. Additionally, it lets you import node objects into the database and create export files containing nodes.

An *object* is any entity managed (analyzed and reported on) by PI. Objects are nodes, CPUs in nodes, memory and buffer space in nodes, an interface, a Frame Relay PVC, an ATM PVC, a slot, or a backplane. Different object types are managed by different report packs.

A *form* gives you a way to view and modify information about the managed objects in PI's property tables. 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.

Understanding the Object Model

PI's default object model looks only at the `ksi_managed_node` property table. When you install report packs, they extend the size and scope of this model.

The *model* is the major category you want to group something under. When you install report packs, default, customer and location models are created. Other technologies may also add their own model.


A category describes specifics in the model and is built from property tables and the columns within those tables. Categories are devices, interfaces, protocols, customers (which is both a model and a category), location, servers, and file systems. Categories are tied to a specific property column. Following are some examples:

Devices	<code>KSI_MANAGED_NODE.DSI_TARGET_NAME</code> <code>K_NODE.NODE_NAME</code> <code>KIRDEVPORTS.DSI_TARGET_NAME</code>
Interfaces	<code>KIRDEVPORTS.INTERFACENAME</code> <code>KV_FRSWITCH_PORT.PORT^a</code>

a. This resolves to `KIRDEVPORTS.INTERFACENAME`.

Currently, PI's object model only goes to the category level. The figure below shows an object model that is expanded beyond the category level. The bold text shows what PI supports. The category object (`foo.com`), child category

(Interface Type, and child category objects (Ethernet and ATM) are not supported by PI; however, it is possible to add them to the model. See [Adding New Object Models](#) on page 312 for details.



Locating Managed Objects

To view managed objects:

- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 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 525.

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

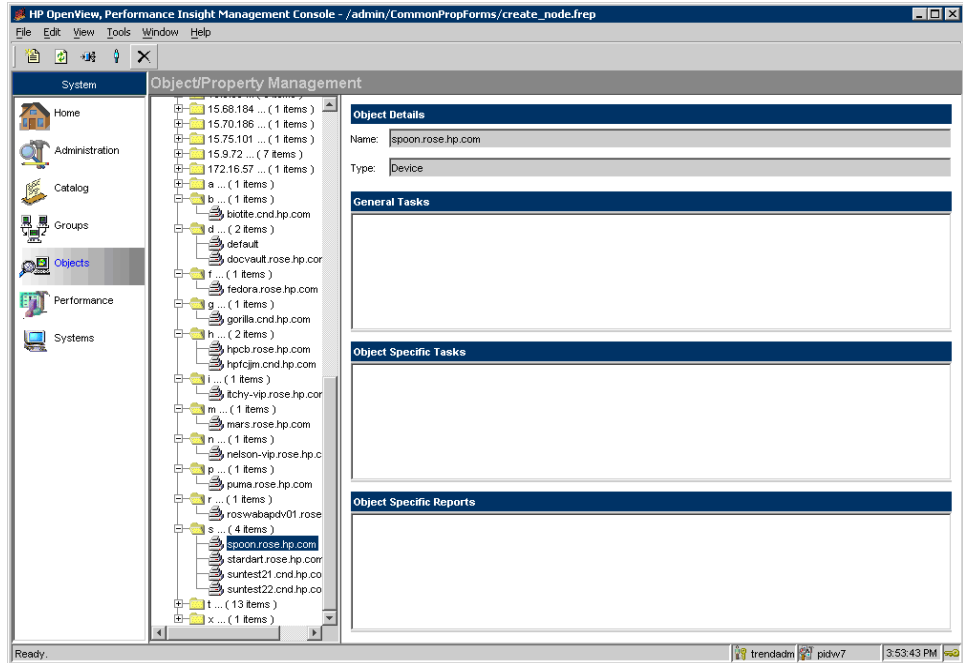
The Object/Property Management window opens in the Management Console. The left side of the window displays the Objects tree, which lists all the objects for the type of view that is selected.

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 309.

If you know the exact name of the object you want, you can use the find feature. See [Finding a Specific Object](#) on page 308.

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 319.

- The reports that apply to the object you selected appear in the Objects Specific Reports section.



If you created your own form or report and it does not display in Object Manager, it could be because you did not create the required property parameters for the form or report. You must create a property parameter and associate it with the table used to create the form or the tables used to create the report's elements. A form requires one property parameter. A report requires a property parameter for each of its elements. See the *HP Performance Insight Guide to Building and Viewing Reports* for more information about creating a property parameters for forms and reports.

3 Select a managed object in the Objects tree or use the find feature. (See [Finding a Specific Object](#) on page 308.)

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, PI generates and displays that report for the managed object you selected. You can view and print the report.
- If you select a task, PI 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 318.

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.

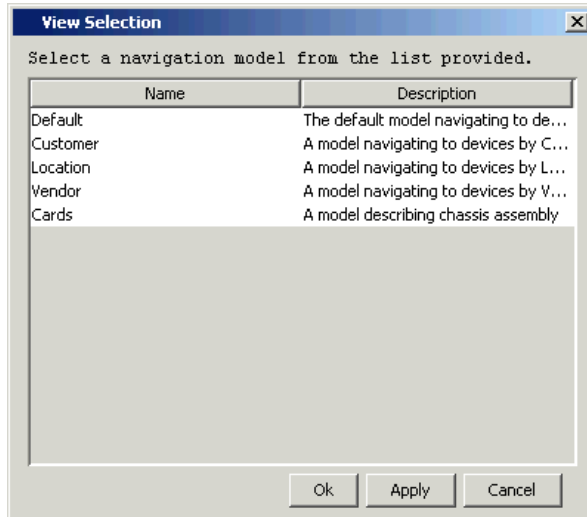
Changing the Objects View

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, select **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, select **Find**.

The Find dialog box opens.



- 2 Type the object name in the text area and click **Find**.

If PI 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.

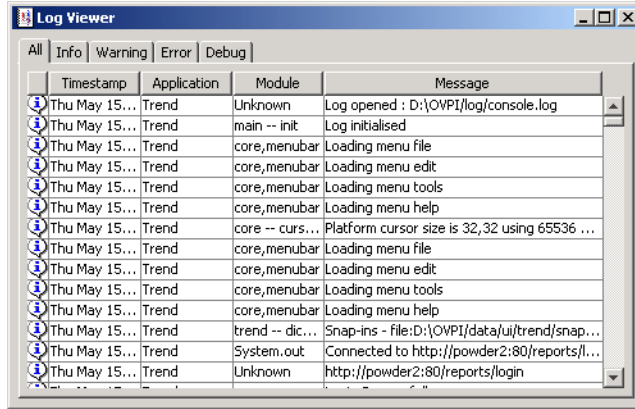
- 3 Do one of the following:
 - If PI did not find the object name you typed, you can try typing another name.
 - If the correct object name is highlighted in the tree, click **Close** in the Find dialog box.
- 4 When the desired object name is highlighted, go to [step 4](#) on page 306.

Viewing Log Information

The View menu includes two options that enable you to view log file and debugging information.

To view log file information, select **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, select **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, PI 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, PI 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 PI 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. PI will display more than the first character of each object name. (Tree Nodes Per Level must be a relatively low value; otherwise PI 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).

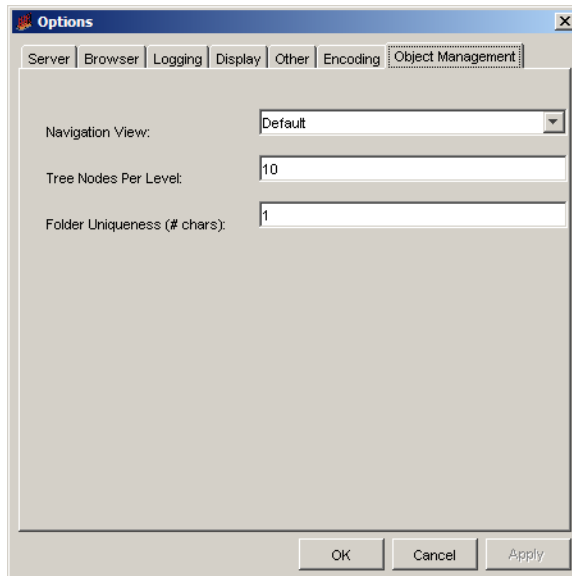
To modify the appearance of the Objects tree:

- 1 From the Edit menu, select **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, type 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, PI looks at the Folder Uniqueness value to determine how to reduce the number of displayed folders and leaves.
- 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.

PI implements this value only if the number of objects in the database exceeds the value in the Tree Nodes Per Level field.
- 6 Click **Apply** (or **OK** if you finish making all of your changes on the Options window).
- 7 From the View menu, select **Refresh**.

PI updates the view in the Objects tree.

Adding New Object Models

If you want to manage something by a particular category, you can create it. Suppose you want to register a new object model for the System Resources report pack (Oracle database?) in which you are only interested in CPU information for your managed servers.

To develop a proper object model, you need to be able to describe the following:

- The model
- The categories in the model
- The property columns that describe those categories
- The parent/child relationships to build the tree
- Forms, reports, or both that take advantage of the model

The master model is Servers, and the categories below the model will be Server and CPU. A hierarchical tree might look something like the following:

Servers

 Server

 CPU

To create the object model:

- 1 Create a file using a text editor and name it something like `create_sys_resources_ObjectModel.sql`.
- 2 In the SQL file, create the model container Servers by typing the following:

```
dsi_relation_model_p (in_action => 1,  
in_name => 'Servers',  
in_image_name => 'device',  
in_descr => 'Model for servers managed by OVPI');
```

Because this is just a container with no relations currently, there is no visual representation of the Servers container.

- 3 Create the category containers for Server and CPU by typing the following in your SQL file:

```
dsi_object_category_p(in_action=> 1,
```



```

in_name=> 'Server',
in_image_name=> 'device',
in_descr=>'Server Category');
dsi_object_category_p(in_action=> 1,
in_name=> 'CPU',
in_image_name=> 'device',
in_descr=>'CPU Child Category');

```

Again, at this point, these two categories are just also empty containers. To describe what belongs in these containers, you must associate them with property columns in the database.

- 4 Create the mappings between categories and property columns by typing the following:

```

dsi_category_instance_p(in_action=> 1,
in_object_category=> 'Server',
in_property_table=> 'K_SYSTEM_CPU',
in_property_column=> 'DSI_TARGET_NAME');

```

These statements place all values of DSI_TARGET_NAME in the table K_SYSTEM_CPU and place them into the Server category.

```

dsi_category_instance_p(in_action=> 1,
in_object_category=> 'CPU',
in_property_table=> 'K_SYSTEM_CPU',
in_property_column=> 'CPUID');

```

These statements place all values of CPUID in the table K_SYSTEM_CPU and place them into the CPU category. At this point, no relationships have been built yet. To build the tree, you need to create the hierarchical relationships.



Regardless of category name, all parent-child relationships are built on the same property table. The first level of any object model has a parent-child relationship to itself.

- 5 To create the hierarchical relationships, type the following in the SQL file:

```

dsi_model_instance_p(in_action=> 1,

```

```

in_model=>'Servers',
in_parent_object_category=>'Server',
in_parent_property_table=>'K_SYSTEM_CPU',
in_parent_property_column=>'DSI_TARGET_NAME',
in_child_object_category=>'Server',
in_child_property_table=>'K_SYSTEM_CPU',
in_child_property_column=>'DSI_TARGET_NAME');

```

If you looked at the Object Manager after performing this step and had data in K_SYSTEM_CPU, you would see the following:

```

--Servers
  -Server
    -ServerA
    -Server B
    -Server C

```

- 6 Create the CPU category by typing the following in the SQL file:

```

dsi_model_instance_p(in_action=> 1,
in_model=>'Servers',
in_parent_object_category=>'Server',
in_parent_property_table=>'K_SYSTEM_CPU',
in_parent_property_column=>'DSI_TARGET_NAME',
in_child_object_category=>'CPU',
in_child_property_table=>'K_SYSTEM_CPU',
in_child_property_column=>'CPUID');

```



The parent is the column you registered first, and the child is a child of that parent from the same property table but from a different column. Children should always be some kind of subset of the parent object. A server will have CPUs. But a CPU will not have many servers.

If you looked at the Object Manager after performing this step, you would see the following:

```
--Servers
  -Server
    -ServerA
      -CPU
        1
        2
        3
    -Server B
      -CPU
        1
  -Server C
```

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 275 describes the parameters for the import file.

Importing the File

To import a file containing nodes to be managed:

- 1 From the File menu, select **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, select **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.

To create your own forms for managed objects, see the *HP Performance Insight Guide to Building and Viewing Reports*.



If you created your own form and it does not display in Object Manager, it could be because you did not create a property parameter for the form. The form must have a property parameter that is associated with the table used to create it. See the *HP Performance Insight Guide to Building and Viewing Reports* for more information about creating a property parameter for a form.

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.

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->25	2	1	90.00	2.00	2.00
5	rfe877->25	2	1	90.00	2.00	2.00
6	ethernet-csma/cd	1	1	30.00	1.00	2.00
7	iso88023-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	iso88026-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**.

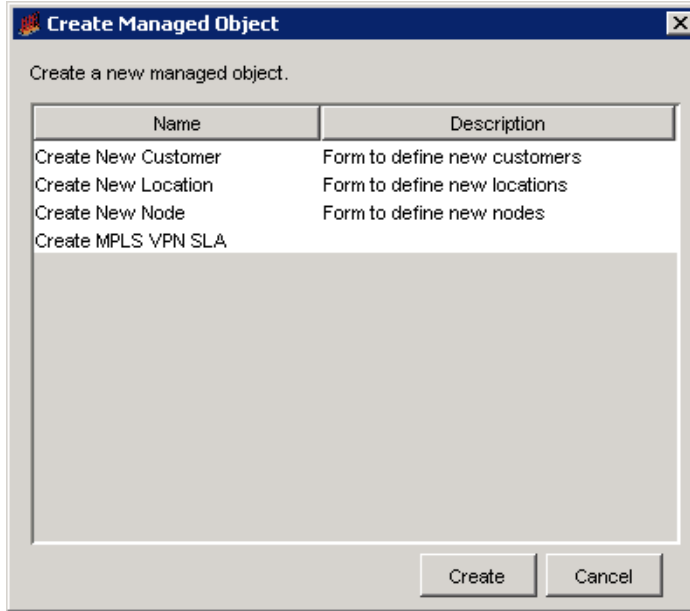
PI changes the property table values associated with the changes you made.

Creating New Managed Objects

To use a form to add a managed object to the database:

- 1 From the File menu, select **New**.

The Create Managed Object dialog box opens.



- 2 Select the form you want to use, and then click **Create**.

The form opens. 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**.

PI adds the objects to the database.

Deleting Nodes

You can delete a node or object from the PI database when such a node or object is not required to be polled any more.



This feature is fully supported only by the Interface Reporting Report Pack.

Deleting a node from the PI database may be required in the following situations:

- A managed object such as a system or interface is no longer in the network or you do not wish to collect data for it.
- There is no need to collect data from some objects. You can mark the objects so that polling will be stopped. Data associated with these objects will be aged out automatically. When all the data is removed, the object will also be removed.
- The managed object is still on the network, but only very specific collections need to occur and it should not appear in other management solutions. For example, it should be managed for net flow, but not interface reporting.

You can delete only Devices or Interface Types or Interfaces using the Delete Managed Object wizard. This wizard will not delete any other type of objects.

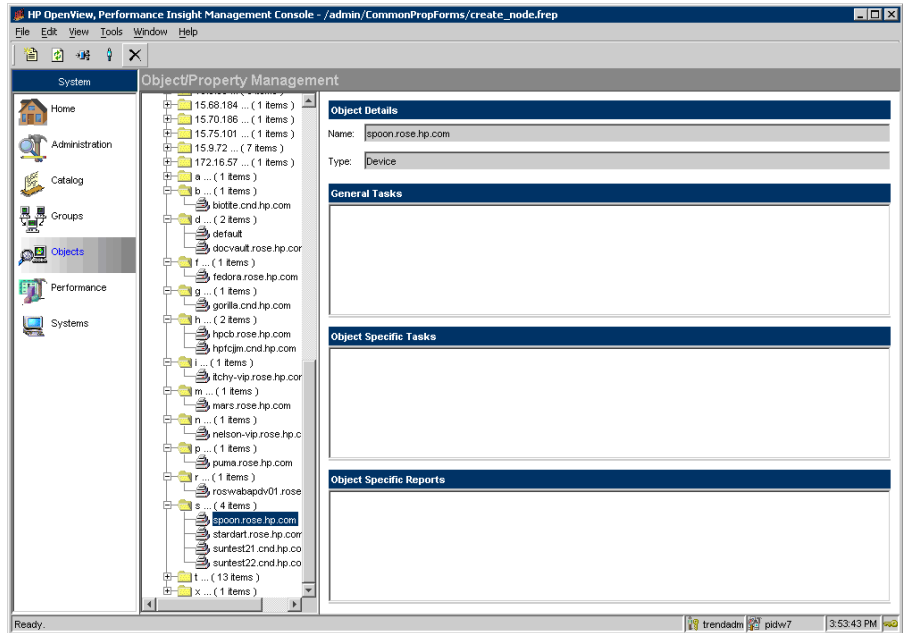
Perform the following tasks to delete a node or group of nodes.


Task 1: [Select a node or group of nodes and start the delete process.](#)

- 1 Click the **Objects** snapin on the Management Console.

The Object/Property Management window opens in the Management Console.

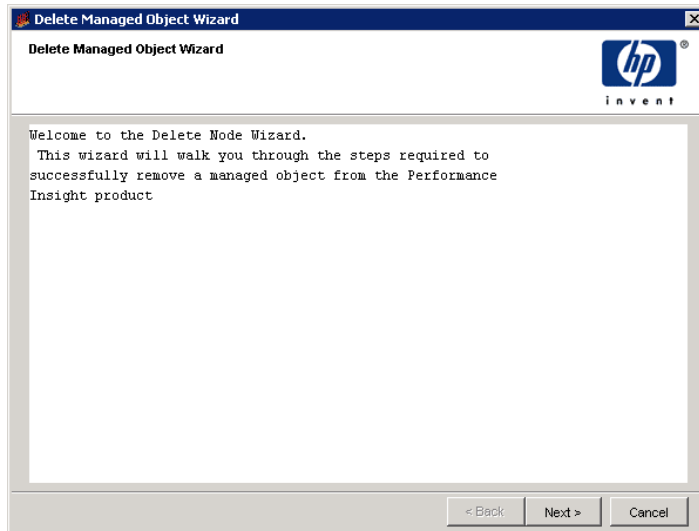
The following figure shows the Object/Property Management window:



- 2 Select the node or group of nodes that you want to delete.
- 3 On the toolbar, click the **Delete** button .

The Preparation for Objects Deletion dialog box opens.

- 4 Click **Yes** to start the object or node deletion process. The Delete Managed Object wizard opens.

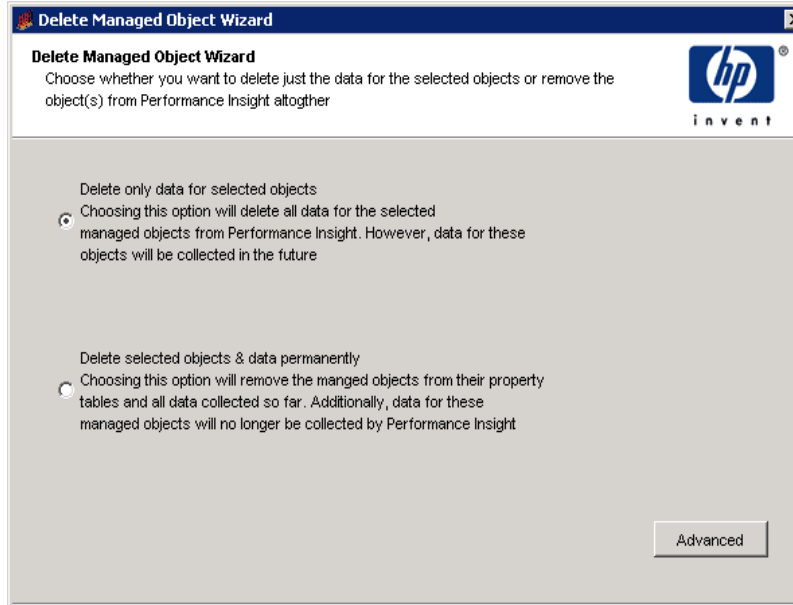


The Delete Managed Object wizard can be invoked only to delete Devices or Interface Types or Interfaces. This wizard will not delete any other type of objects.

Task 2: Select type of deletion.

- 1 On the Delete Managed object wizard, click **Next** to continue. A wizard opens that allows you to select between two options.
- 2 Select one of the following.
 - **Delete only data for selected objects:** If you select this option, it will delete all data for the selected node from the PI database. The node will not be deleted and data can be collected for this node or group of nodes in the future.
 - **Delete selected objects and data permanently:** If you select this option, it will remove the selected node or group of nodes from their property tables. All the data collected so far will be deleted. Data cannot be collected for this node or group of nodes in the future.


The following figure shows the two delete options.



Task 3: Control the type of deletion.

Delete only data option:

- 1 Click **Advanced** (after choosing Delete only data for selected objects). The Advanced Delete wizard that allows you to control the type of deletion.



- 2 To stop data collection for any category of the node, select the **Stop Collect** box.
- 3 By default, the **Data Only** box will be selected. You can deselect the **Data Only** box if you do not want the data to be deleted. By default, the **Object & Data** box will be disabled and it will show a .

- 4 Click **OK** to continue.



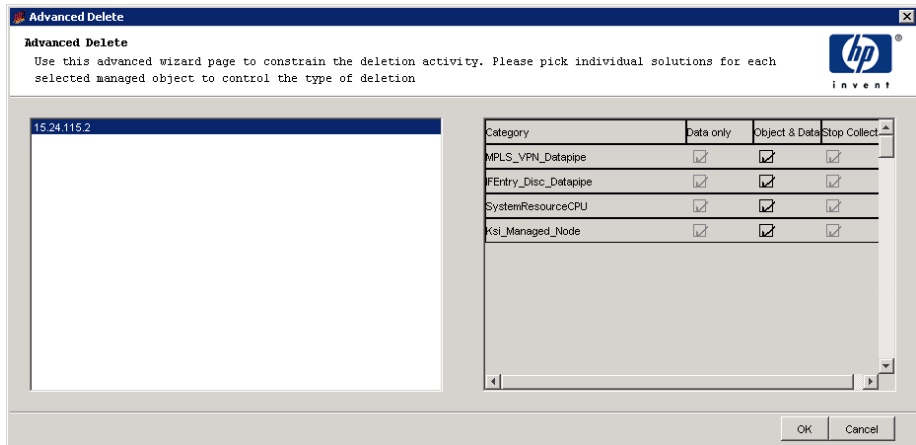
If you are selecting **Stop Collect**, make sure that you do the following:

- **Stop Collect** means new data will not come for the selected Report Pack in the PI database. It does not mean Network Polling for the device or interface will be stopped. Poll from network will still happen but data will not be uploaded into the PI Report Pack schema. This is not an option to reduce network load.
- The Stop Collect option is supported for the following categories only:
 - ksi_managed_node
 - CommonProperty Tables
 - Interface Reporting

For other categories, you can select **Stop Collect** but no GUI support is available to enable collection for those categories.

Permanent node delete option.

- 1 Click **Advanced** (after choosing Delete selected objects and data permanently in step 6). The Advanced Delete wizard allows you to control the type of deletion.



- 2 By default, the **Data Only** and **Stop Collect** boxes will be checked and disabled. Also, the **Object & Data** box will be selected by default. You can deselect it if you do not want to delete the data for a specific report pack.

- 3 Click **OK** to continue.




If you are selecting/deselecting report packs from this GUI, make sure that in the final selection of report packs you include either `ksi_managed_node` or Common Property Tables or Interface Reporting along with your other selections.

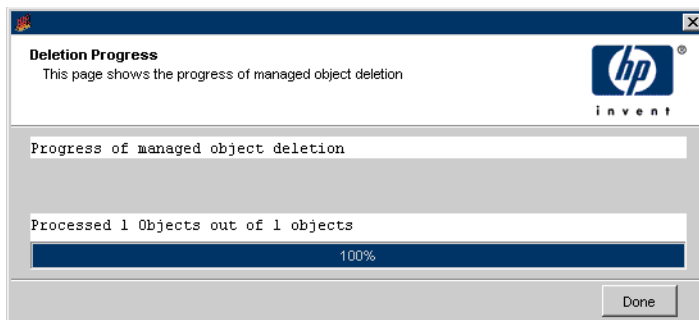
If you do not select any of the previously mentioned report packs, the wizard will not include such objects in the PI database. They will be permanently disabled.

Task 4: Complete the deletion process.

- 1 On the Delete Managed Object wizard, click **Next**.
- 2 The Delete Managed Object - Completion page opens. This page provides the information related to the tables which are affected. It also shows all the options you selected for the deletion: Data delete or Object Delete or Stop Collect.
- 3 Click **Finish** to complete the deletion process. The following Confirmation dialog box opens.



- 4 Click **Yes** if you are sure of the deletion operation. The Deletion Progress page opens which shows the progress of node deletion process.



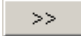
You can click on **Done** to close the Deletion Progress page, but the process runs in the background. You can view the same Deletion Progress page by choosing **Deletion Progress** from File menu.

To edit the Exclusion List of nodes

The Exclusion List keeps a record of all the deleted objects and also those objects that were marked to stop collect data. This list has the following properties:

- This list is similar to keymap tables and mbr tables associated with the property table but created when required by the Delete Managed Object service.
- It is associated with the property table and is deleted when `datapipe_manager` deletes the property table.

Perform these steps to edit the Exclusion List.

- 1 On the Objects snapin, from File menu, select **Edit Exclusion List**. The Edit Exclusion List page opens.
- 2 Select any category to view the excluded objects on the left-hand side pane.
- 3 Select the excluded object that you want to include for collecting data again.
- 4 Click  to move the selected object from the Excluded objects list to the Included objects list.
- 5 Click **Edit** and then click **Done** to confirm the inclusion.

11 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 PI; 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.

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 PI 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.



PI 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 group 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 conveniently 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 118 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 525.

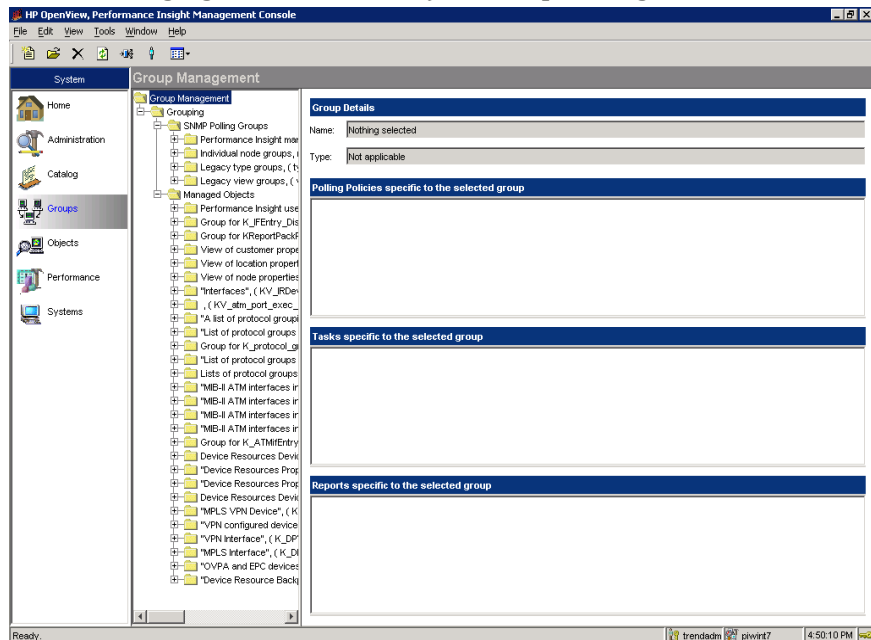
- 2 Click the **Groups** snapin in the navigation pane.

The Object Group Management window appears in the Management Console.

The Object Group Management window has a navigation and a viewing pane.

The navigation pane 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 336).

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, PI generates and displays that report for the group you selected. You can view and print the report.
- If you select a task, PI 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 318.

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.

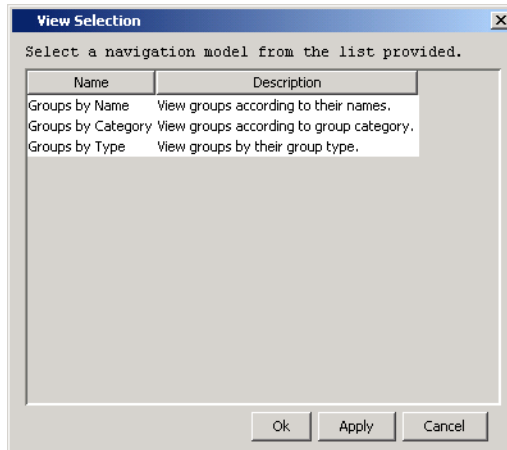
Changing the Groups View

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 353.)

To change the groups view:

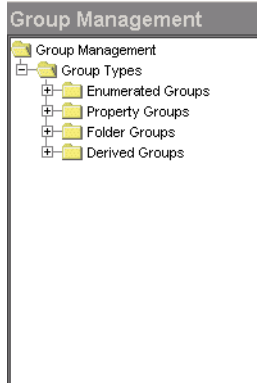
- 1 From the View menu, select **Change View**.

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, select **Find**.

The Find window opens.




- 2 Type the group name in the text box and click **Find**.

If PI 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 PI 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 334.

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.



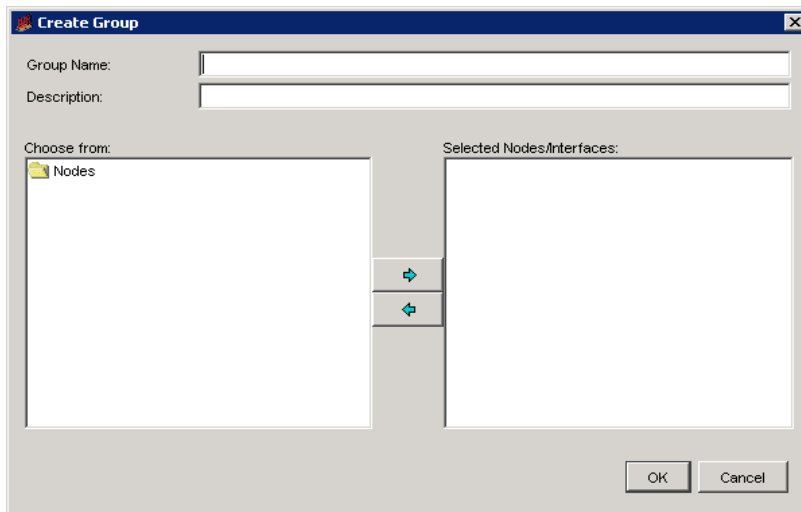

Creating an Enumerated List Group

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 select **Create Enumerated** from the menu.
 - From the File menu, point to New, and select **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 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.

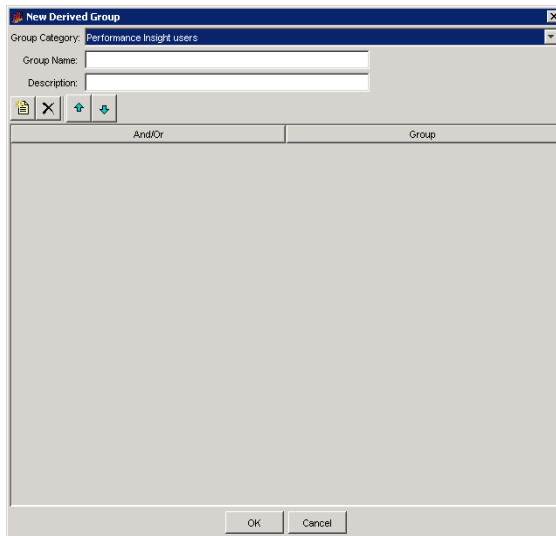
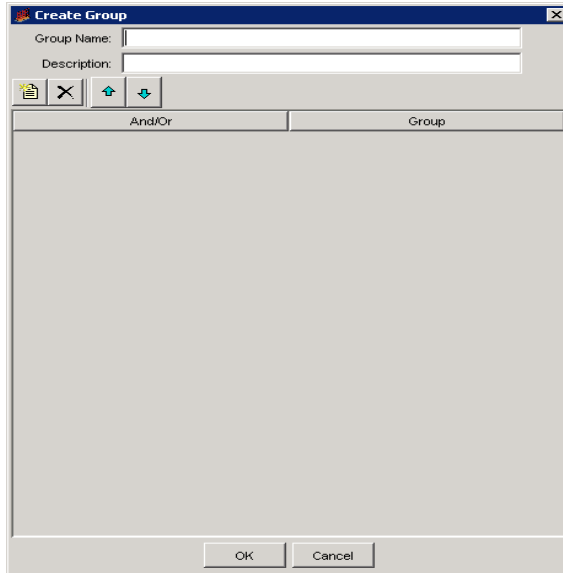
Creating a Derived Group

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 select **Create Derived** from the menu.
 - From the File menu, point to New, and select **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 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

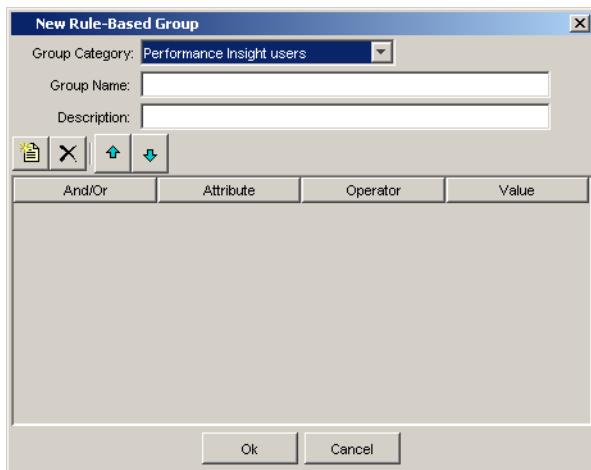
A *property group* comprises a group of managed objects that are based on certain properties in their property tables.

▶ 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 select **Create Property** from the menu.
 - From the File menu, point to New, and select **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 type 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 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 click .
- 8 Click **OK** when you are finished.

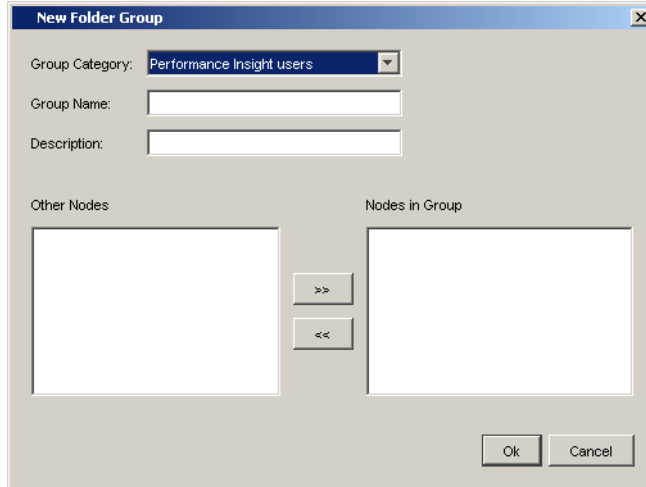
Creating a Folder Group

A *folder group* can contain an enumerated list of managed objects and other folder groups.


To create a folder group:

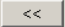
- 1 Do either of the following:
 - If you know the group category that you want to use to create the group, right-click the group category's folder in the navigation pane, and select **Create Folder** from the menu.
 - From the File menu, point to New, and select **Folder Group**.

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 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 you must edit the group to create subfolder groups by moving them into another folder group. See [Editing a Group](#) on page 350.

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, select **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 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, select **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, select **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 356.




Viewing a Group

To view the contents of a group, do one of the following:

- Right-click a group name, and select **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 353.)



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.

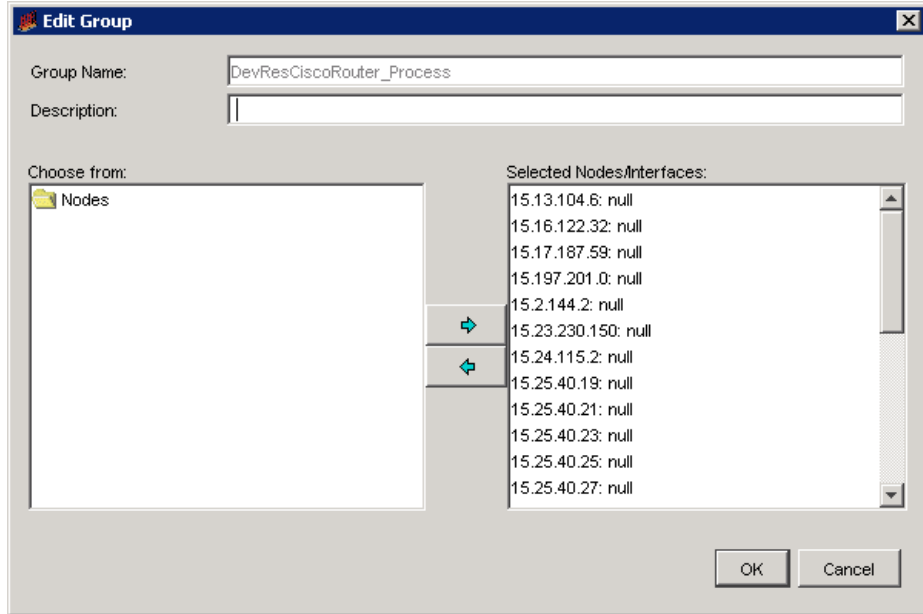


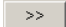
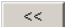
Group Manager does not allow you to edit All type groups.

To edit a group:

- 1 Right-click the group name, and select **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 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 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:

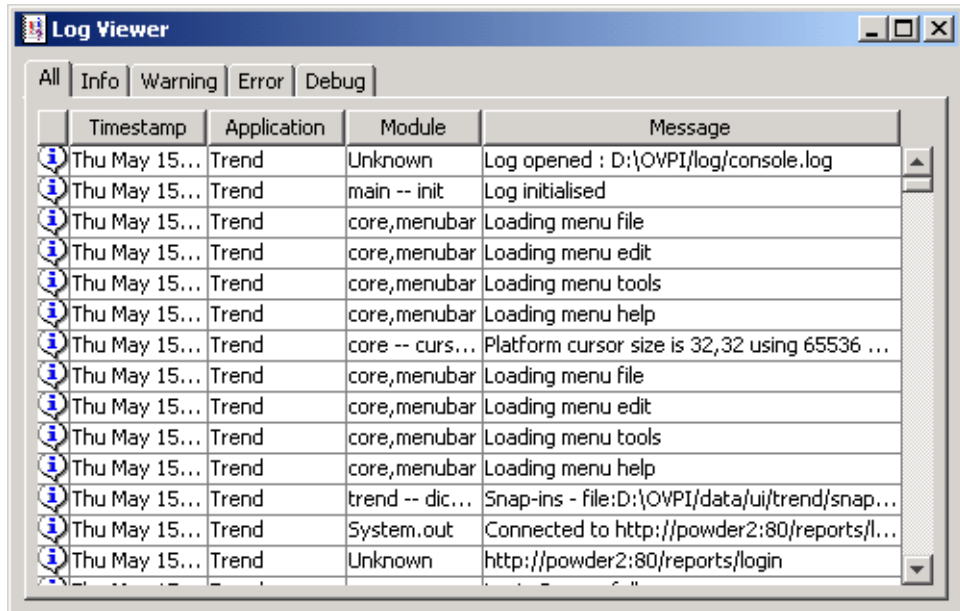
- 1 Right-click the group name, and select **Delete Group** from the menu.
A deletion confirmation window displays.
- 2 Click **Yes**.

Viewing Log Information

Using the options on the View menu, you can view log file and debugging information.

To view log file information, select **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, select **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, select **Options**.
The Options window opens.
- 2 Select the Group Management tab.

The Group Management page opens.

The screenshot shows a dialog box titled "Options" with a close button (X) in the top right corner. The dialog has several tabs: "Server", "Browser", "Logging", "Display", "Other", "Encoding", and "Group Management". The "Server" tab is selected. Inside the dialog, there are three input fields: "Host" with the value "piwint7.india.hp.com", "Port" with the value "8088", and "Use SSL" with an unchecked checkbox. Below these fields is a button labeled "Reset to System Defaults". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Apply". A note at the bottom of the dialog reads: "NOTE: When you make changes to the server information, you must logout and log back in for the changes to take effect."

- 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 PI 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

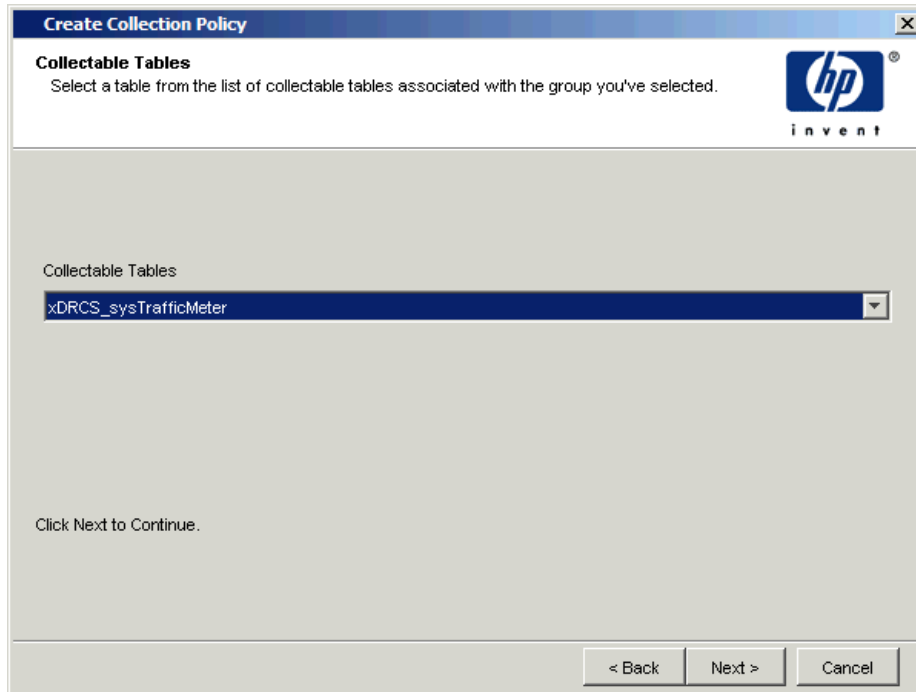
You can create a group polling policy using the Create Collection Policy wizard. You can also edit or delete a polling policy associated with any group by using the Polling Policy tab that is available when you view a group.

Creating a Group Polling Policy

To create a polling policy for a group:

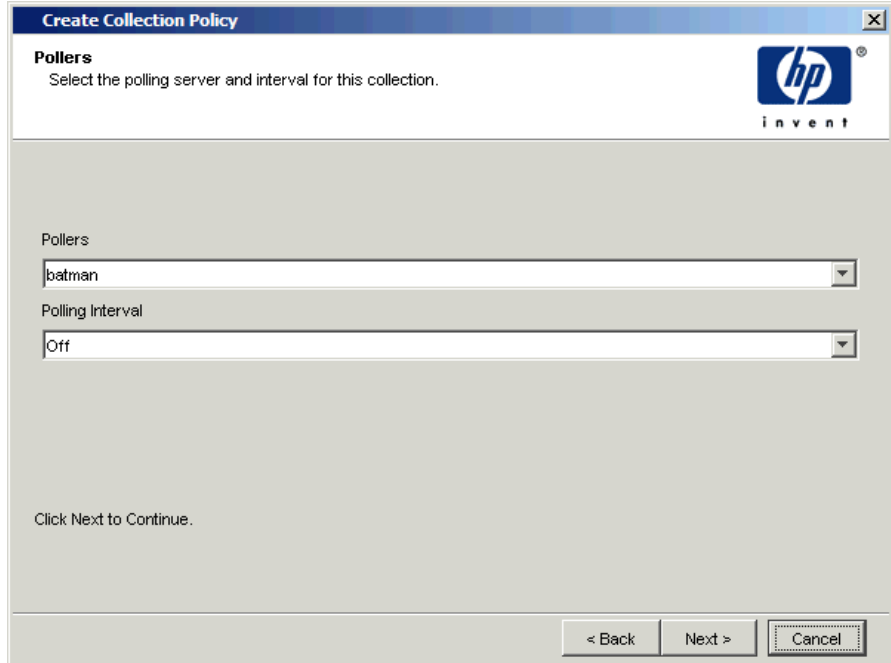
- 1 Right-click the group for which you want to create a polling policy, and select **New Collection** from the shortcut menu.

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, select a table, and 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



Create Collection Policy

Pollers
Select the polling server and interval for this collection.

hp
invent

Pollers
batman

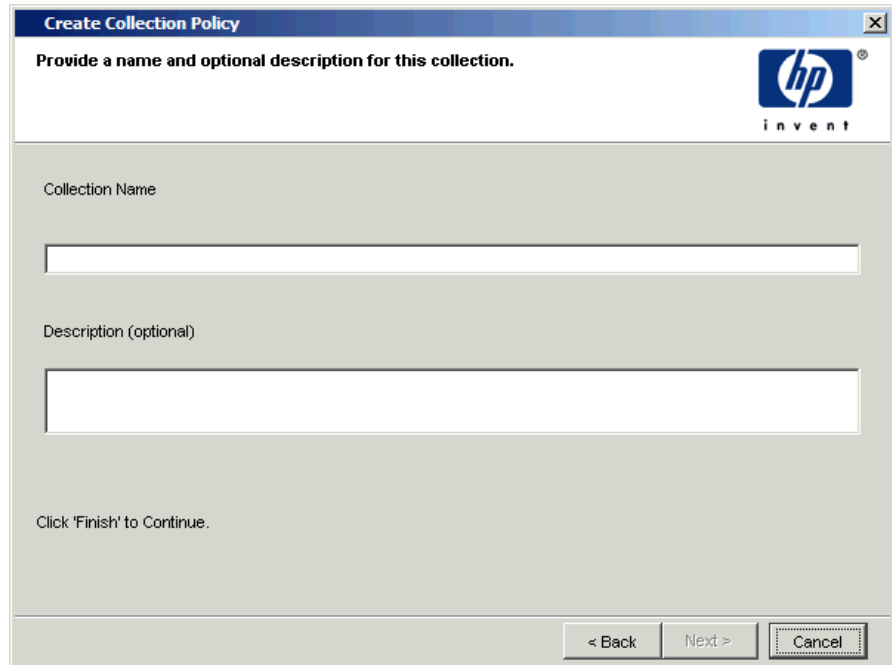
Polling Interval
Off

Click Next to Continue.

< Back Next > Cancel

- 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



Create Collection Policy

Provide a name and optional description for this collection.

hp
invent

Collection Name

Description (optional)

Click 'Finish' to Continue.

< Back Next > Cancel

- 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 select **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 257.

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 select **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 select **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.

See the *HP Performance Insight Guide to Building and Viewing Reports* for more information about selecting deferred nodes for reports.

12 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.
- **Remote Pollers.** A remote poller is a system that collects data from network nodes and stores the data in a PI database. The remote poller collects data from the nodes in the network and stores the data on a server running the PI database. A Remote Poller is set up to be a dedicated poller. It has no database or clients. You can add, delete, and modify the properties of these Remote Pollers

Starting the System and Network Administration Application

To start the System and Network Administration application:

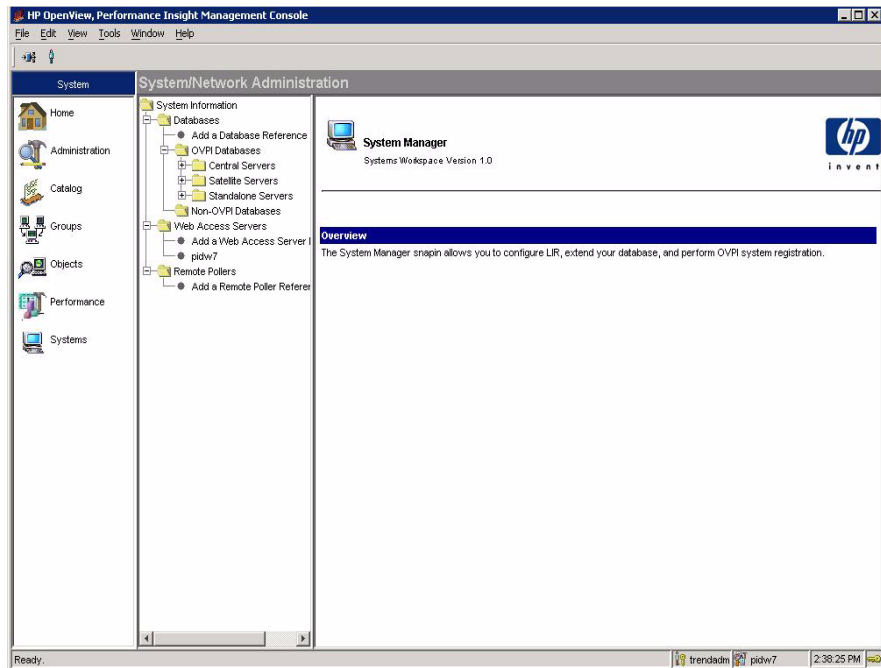
- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 for details.




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 525.

- 2 Click the **Systems** snapin in the navigation pane.

The System/Network Administration window appears on the Management Console. The left side of the window displays the Databases, Web Access Servers and the Remote Pollers.



When the  icon is blue, this indicates the system that you are currently logged on to.

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 (-).

Managing Your Web Access Server Systems

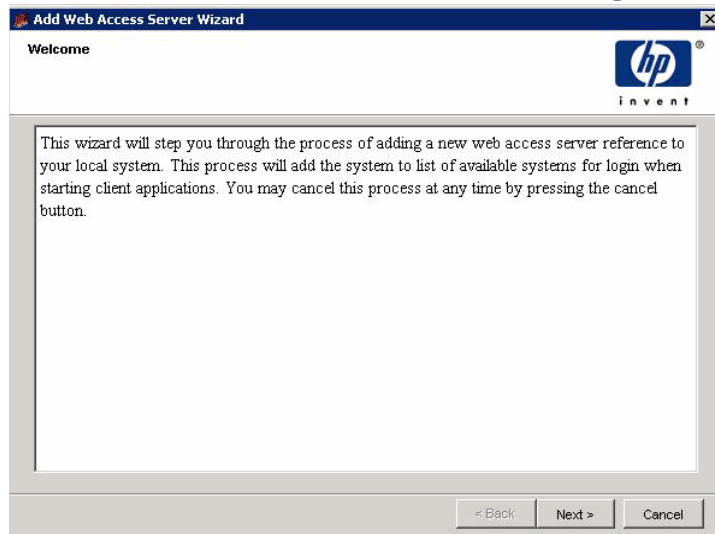
Using the System/Network Administration component, you can add Web Access Server systems, edit their properties, remove them from the list of available systems, and synchronize the local system with the currently selected Web Access Server.

Adding a Web Access Server

To add a Web Access Server to the list of systems that you can access:

Task 1: Start the Add Web Access Server wizard.

- 1 Expand the Web Access Server list by clicking on the (+) symbol. The option to **Add a Web Access Server Reference** appears.
- 2 Click the **Add a Web Access Server Reference** option.



- 3 Click **Next**.

Task 2: Supply the information needed to connect to the Web Access Server.

Add Web Access Server Wizard

Connection Information
Enter the hostname and http port of the OVPI management server you wish to connect to.
Check the Use TLS/SSL box if you wish to connect via SSL.

HP
invent

Hostname
[Empty text box]

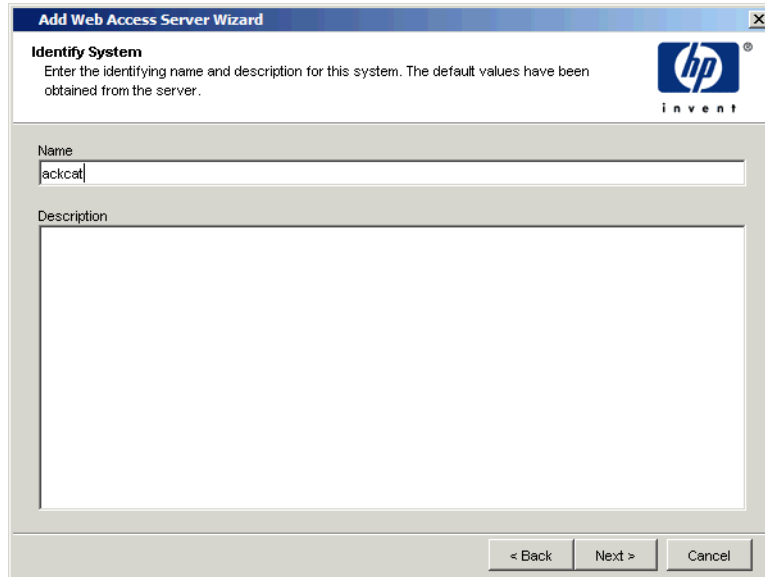
HTTP Port
[80]

Use TLS/SSL

< Back Next > Cancel

- 1 In the Hostname box, type the hostname of the Web Access Server.
- 2 In the HTTP Port box, type the port number for the Web Access Server if its port is not 80, which is the default.
- 3 Select the **Use TSL/SSL** box if the Web Access Server is using SSL and you want secure communications with it.
- 4 Click **Next**.

Task 3: Type a name and description for the Web Access Server.



5 Type a name for the server in the Name box (for example, neast2).




The name you specify cannot already exist in PI's system configuration file (`systems.xml`).

After you add the database, it is registered in the system configuration file with this name.

6 Type a description for the server in the Description box. A description is optional.

7 Click **Next**.

Task 4: Review the information you entered



- 1 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.

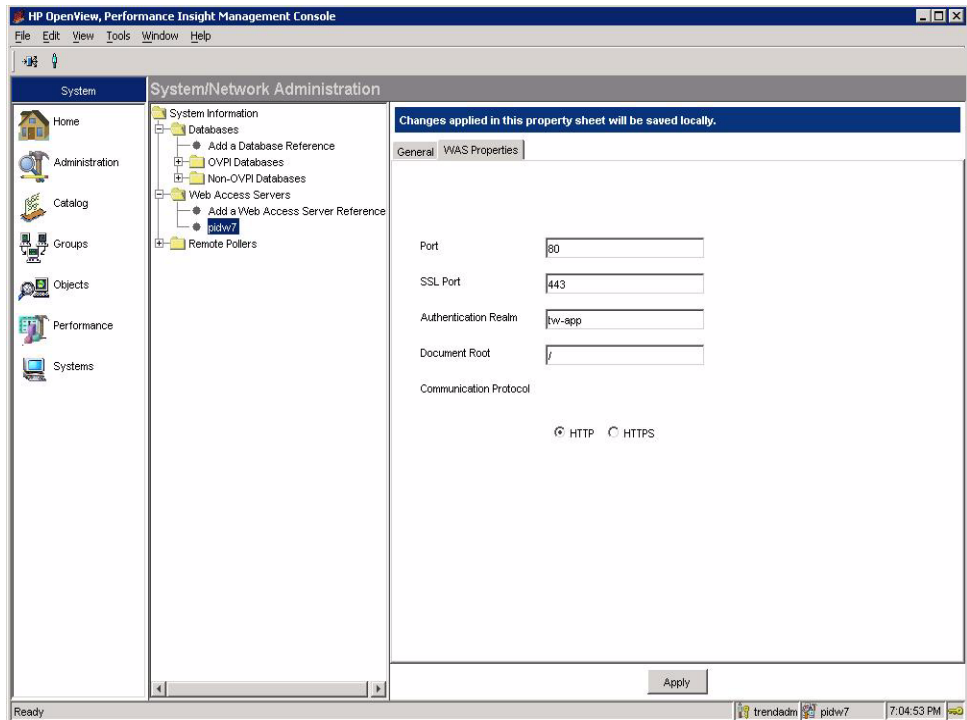
- 2 Click **Finish** to add the Web Access Server.

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 the server name. The viewing pane displays two tabs, **General** and **WAS Properties**.



This dialog box contains the following tabs:

- **General.** It contains the following information:
 - **Name.** Type the name of the Web Access Server.
 - **Hostname.** Type the hostname of the Web Access Server.
 - **IP Address.** Type the IP address of the Web Access Server.
 - **Description.** Type a description.
- **WAS Properties.** It contains the following information:
 - **Port.** Type the port number of the system.
 - **SSL Port.** Type the SSL Port number of the system.

- Authentication Realm. Type the realm. For more information about the realm, see [step b](#) on page 566.
- Communication Protocol. Click **HTTP** or **HTTPS**.

You can change the information and click **Apply** to save them.

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.

- 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 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_dpipes.
 - 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 373 for details.

Synchronizing Web Access Servers

To synchronize a system with the current Web Access Server:

- 1 Expand the folder for your Web Access Server in the navigation pane, and click the server name. The viewing pane displays two tabs, **General** and **WAS Properties**
- 2 Click the **General** tab.

If you select your current system, you will not be able to select any synchronization options because you cannot synchronize a system with itself. Select another system.

- 3 Click **Synchronize** from the menu, and select 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 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. 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 system's 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.

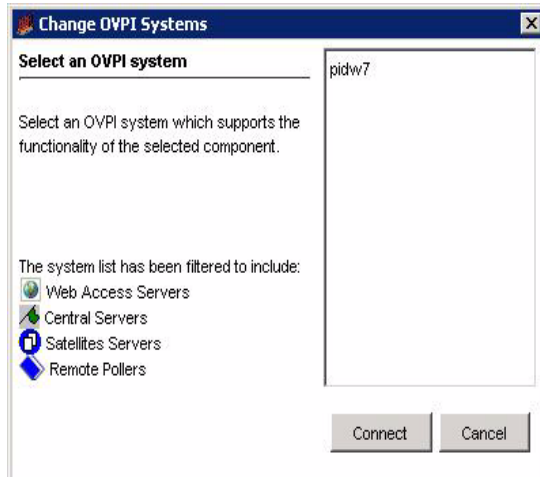
For example, if a system you compare the Web Access Server to has a database that you want added to the Web Access Server, the synchronization process will add the database to the Web Access Server.

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 select **Change Server** from the File menu. The Change OVPI System 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 navigation pane of the System/Network Administration window.
- 2 Select **Delete System Reference**. 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 Databases folder under the System/Network Administration window.
- 2 Click **Add a Database Reference**. The **Add Database Wizard** opens.
- 3 Click **Next** to continue. You can select the Type of Database here.
- 4 Do the following:
 - a Select one of the following options:
 - **OVPI Databases**. This option lets you add a database that contains the OVPI database schema.
 - **Non-OVPI Databases**. This option lets you add a generic database. A generic database does not contain the PI database schema. Supported generic database types are Sybase, Oracle, SQLServer, and Solid.



You must assign a non-PI database to a user before the user can access the database. Users cannot build reports using data from a non-PI database unless their user accounts have been set up to access that particular database.

To assign the non-PI database to a user, you can make the database the default for the user's account or create a second user account for the user, which will access the non-PI database. See [Modifying a User Account](#) on page 531 or [Creating a User Account](#) on page 527 for details.

- b Click **Next**.
 - If you selected a PI database, the Connection Information window opens. Go to [step 5](#).
 - If you selected a non-PI 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 Identify System 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 Identify System page opens.
- 7 Do the following:
 - a Type a name for the database in the Name box (for example, `neast2`).



The name you specify cannot already exist in PI's system configuration file (`systems.xml`).

After you add the database, it is registered in the system configuration file with this name.

- b Type a description for the server in the Description box. A description is optional.
 - c Click **Next**. The Summary window opens.
- 8 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 databases for that server.

- ▶ To build and view reports for non-PI database systems, see *HP Performance Insight Guide to Building and Viewing Reports*.

Viewing Database Properties

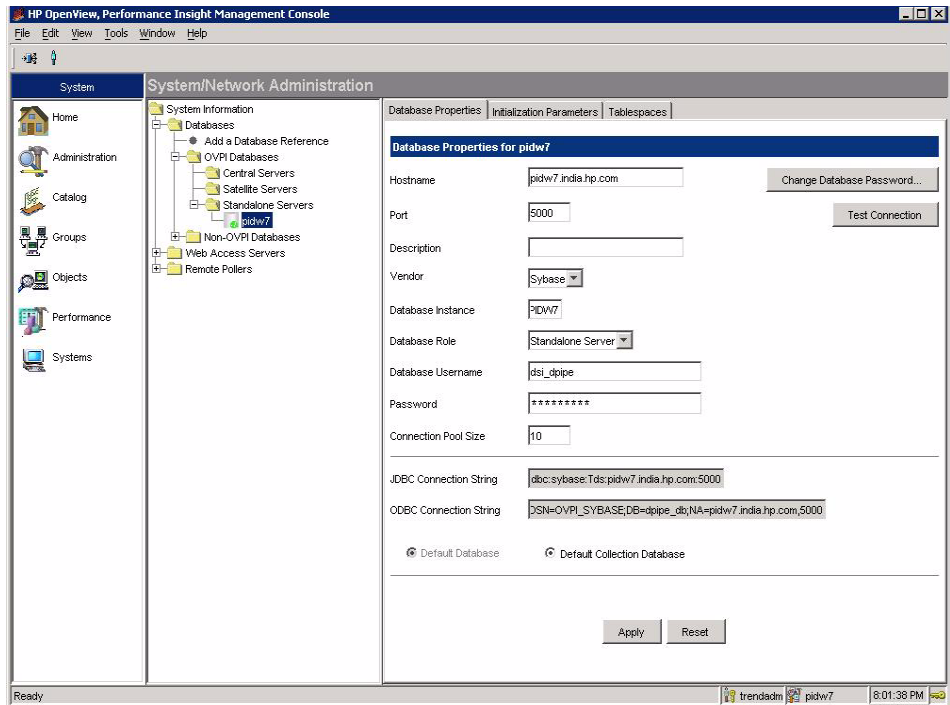
You can view the properties for your database and modify the database password, if needed.

- ▶ 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 On the System Information navigation tree, expand the Databases folder.
- 2 Click the database you want to view. A window containing three tabs opens.

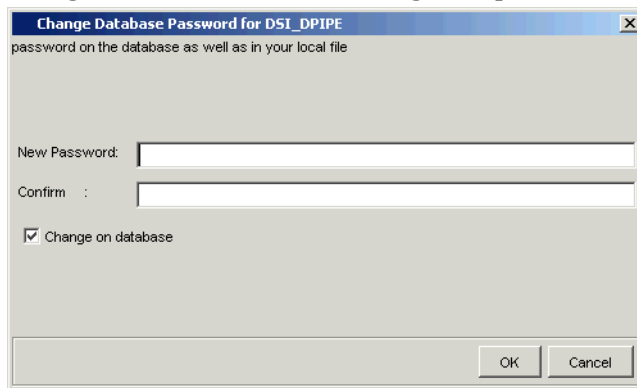
- Click the **Database Properties** tab. The database's properties display in the viewing pane.



The properties shown are as follows:

- HostName.
- Port.
- Description.
- Vendor. This is the database type (Sybase or Oracle).
- Database Instance. By default, these names are DSQUERY (Sybase) or \$ORACLE_SID (Oracle).
- Database Role. Select your Database role from the drop-down menu: Standalone, Central or Satellite Server.
- Database Username.
- Password.

- **Connection Pool Size.** This is the number of connections that you have made available to PI.
 - **JDBC Connection String.** This is the string that PI's Java components use to connect to the database.
 - **ODBC Connection String.** This is string that PI's C-based components use to connect to the database.
 - **Test Connection button.** Click **Test Connection** to test the connection 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 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.

Changing Database Initialization Parameters

To change database initialization parameters:

- 5 On the System Information navigation tree, expand the Databases folder.
- 6 Click the Database you want to view and click the tab **Initialization Parameters**.

The database initialization parameters display in the viewing pane.

Modifi...	Requires R...	Parameter	Value	Description
<input checked="" type="checkbox"/>	<input type="checkbox"/>	SQL Perfmom Integration	0	SQL Perfmom 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, it may be necessary for you to extend the database.

Database Component Sizing Guidelines

Table 8 Database Component Sizing Guidelines

Database Component	Guideline
Temporary database (tempdb)	50% of your PI database (minimum), up to 1 GB
Transaction log	50% of your PI 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 PI 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

- ▶ If the system on which you installed PI is using Oracle Client to access a remote Oracle database, you cannot extend tablespaces.

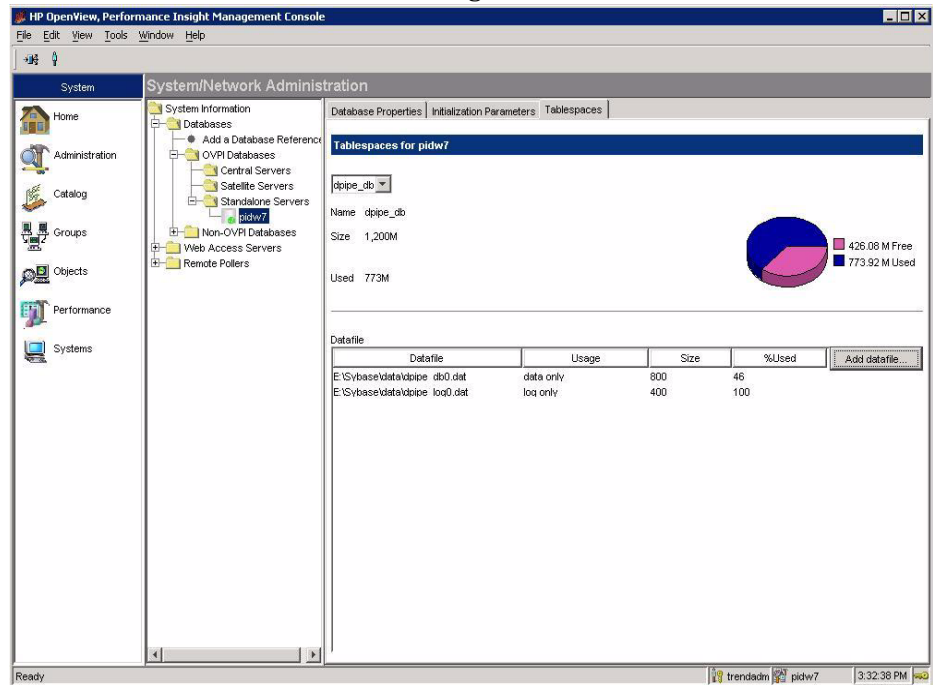
To expand the database:

- 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 PI 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`.


- Click on a tablespace. The Tablespaces window opens in the viewing pane. It will look similar to the following:



The Tablespaces 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 tempdb 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 PI 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.

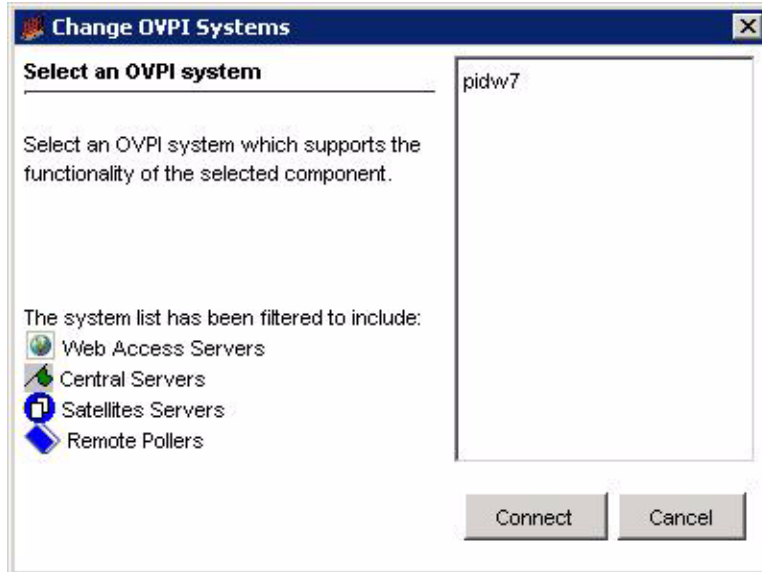
- d Click **Extend**.

The new data file will appear in the window.

Changing the Database Server

To change the database server you are using:

- 1 Click  on the toolbar or select **Change Server** from the File menu. The Change OVPI Systems dialog box opens.



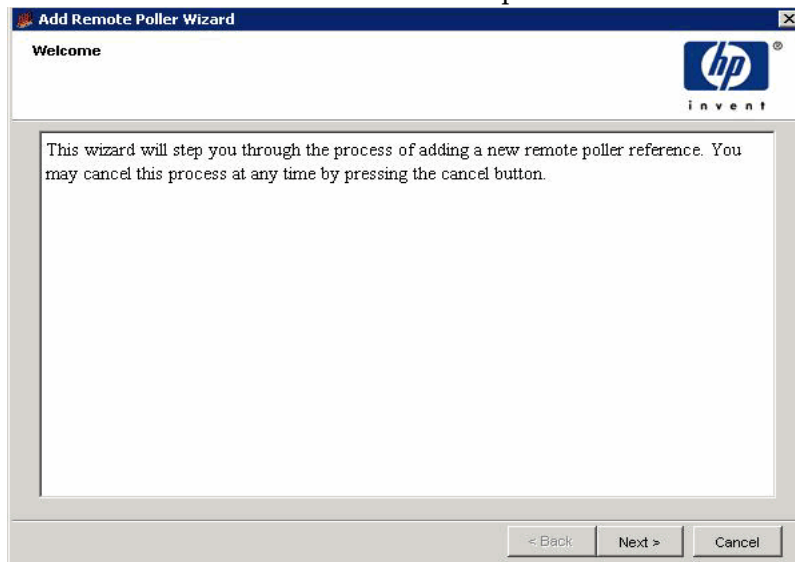
- 2 Select the database you want, and click **Connect**. After the connection is made, the database icon next to this system should be highlighted green.

Add a Remote Poller

Using the System/Network Administration component, you can add a Remote Poller to the list of systems that you can access:

Task 1: Start the Add Remote Poller wizard.

- 1 Expand the Remote Pollers list by clicking the (+) symbol. The option **Add a Remote Poller Reference** appears.
- 2 Click the **Add a Remote Poller Reference** option.

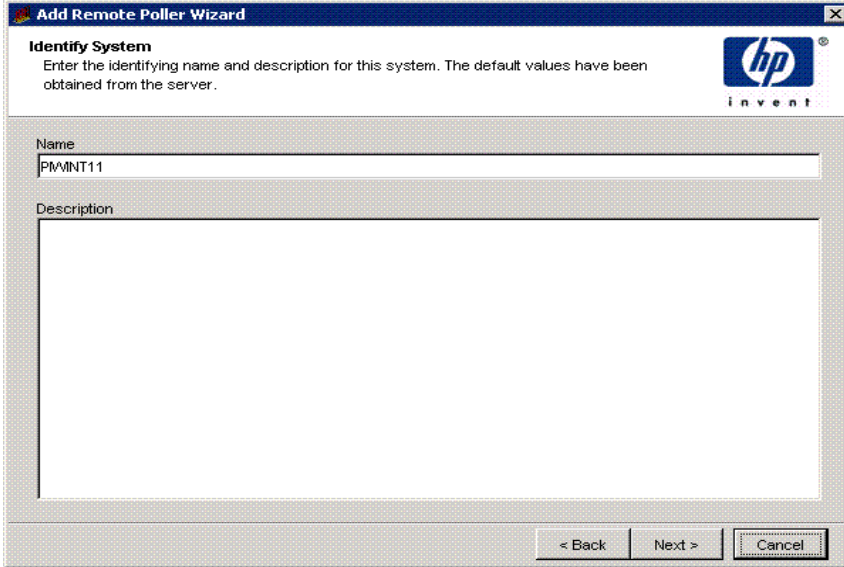


- 3 Click **Next**. The Connection Information window opens.

Task 2: Supply the information needed to connect to the Remote Poller.

- 1 In the Hostname box, type the hostname of the Remote Poller.
- 2 In the HTTP Port box, type the port number for the Remote poller.
 - ▶ You cannot connect to a Remote Poller using a TSL/SSL connection.
- 3 Click **Next**.

Task 3: Type a name and description for the Remote Poller.



The screenshot shows a Windows-style dialog box titled "Add Remote Poller Wizard". The main heading is "Identify System". Below the heading, there is a short instruction: "Enter the identifying name and description for this system. The default values have been obtained from the server." In the top right corner, there is the HP logo and the word "invent" below it. The dialog contains two input fields: a "Name" field with the text "PMVNT11" and a larger "Description" field which is currently empty. At the bottom of the dialog, there are three buttons: "< Back", "Next >", and "Cancel".

4 Type a name for the Remote Poller in the Name box.




The name you specify cannot already exist in PI's system configuration file (`systems.xml`).

5 Type a description for the Remote Poller in the Description box. A description is optional.

6 Click **Next**.

Task 4: Review the summary of the information you have entered for the Remote Poller.



1 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.

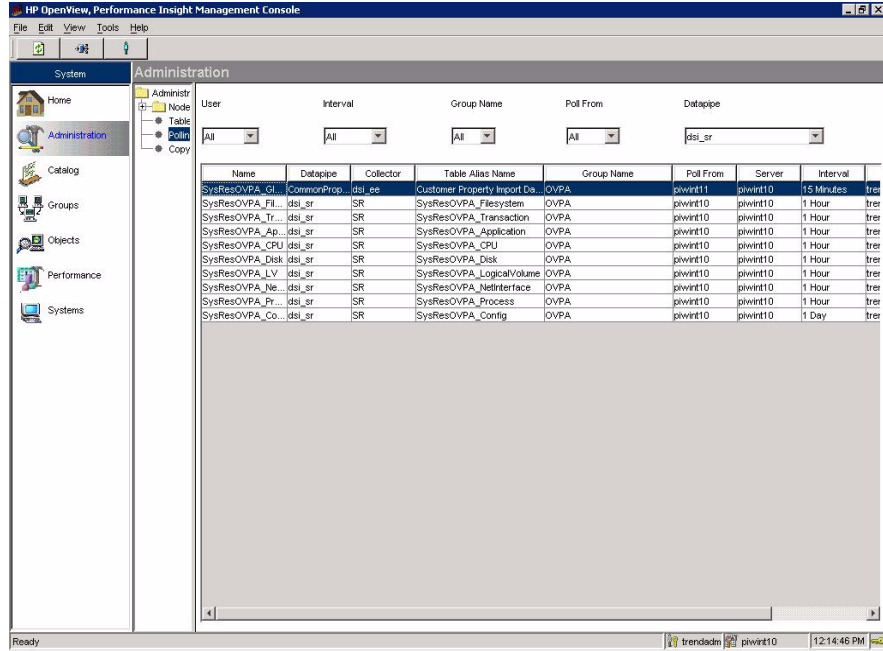
2 Click **Finish** to add the Remote Poller.

The wizard closes, and the Remote Poller is added to the list.

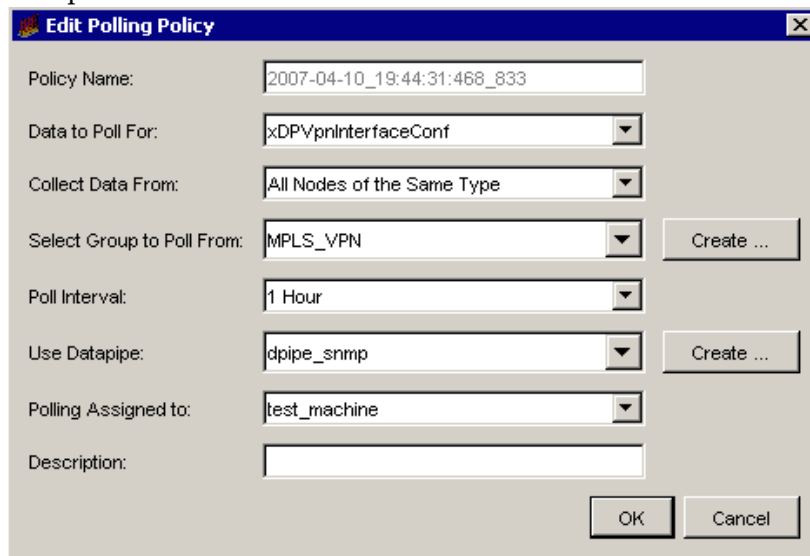
Task 5: [Edit Polling Policy](#)

1 Click the **Administration** snapin in the navigation pane and click **Polling Policy**.

2 Select a Polling Policy from the list of Polling Policies.



3 From the Edit menu select Polling Policy...The Edit Polling Policy window opens.



- 4 Click **Create** located beside Use Datapipe. The Create Datapipe Installation window opens.

Create Datapipe Installation [X]

Hostname:

IP Address:

Rule File:

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

OK Cancel

- 5 Type the Hostname of the Remote Poller and click **OK**. This will take you back to the Edit Polling Policy window.
- 6 Click **OK** on the Edit Polling Policy window.

Repeat steps 2 through 6 for each of the Polling Policies.

13 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 PI Table Manager application.

For more information about tables, see the *HP 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:

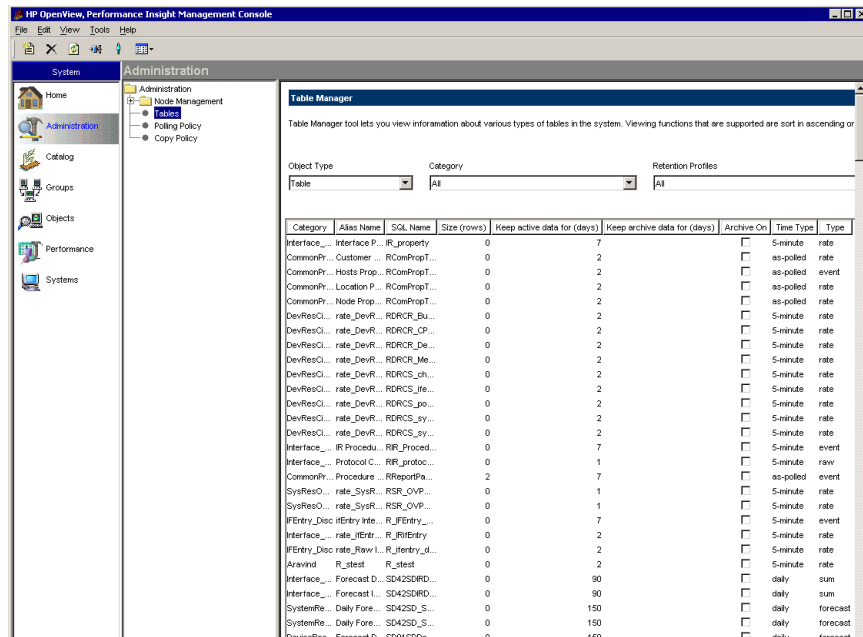
- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 for details.



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 525.

- 2 Click the **Administration** icon in the navigation pane. The Administration snapin opens.

- 3 In the navigation tree provided in the Administration snapin, click **Tables**. The Table Manager window opens.



The Table Manager 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. The pane shows a tree. The tree contains folders. The folders contain items, which may include other items or additional folders.

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 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 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 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 and click.

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 select **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, follow these steps:

- 1 From the Edit menu, select **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 Alias Name...** to locate the alias name.

Viewing Data Tables by Type or Category

The viewing pane contains various filters. You can use the following filters to view only a subset of the available data tables:

- Object Type
- **Category**
- **Retention Profiles**
- **Table Types**
- **Time Types**

For example, the following shows the results that are displayed in the viewing pane when you select router from the Category folder.

Table Manager tool lets you view information about various types of tables in the system. Viewing functions that are supported are sort in ascending or descending order, filter through category to time type, refresh, find, view column details etc.

Object Type Category Retention Profiles Table Types Time Ty

Table CommonPropertyTables All Rate All

Category	Alias Name	SQL Name	Size (rows)	Keep active data for (days)	Keep archive data for (days)	Archive On	Time Type
CommonPropertyTables	Customer P...	RComPro...	0	2		<input type="checkbox"/>	as-poll
CommonPropertyTables	Location Pr...	RComPro...	0	2		<input type="checkbox"/>	as-poll
CommonPropertyTables	Node Prop...	RComPro...	0	2		<input type="checkbox"/>	as-poll



The columns that you see may be different from the columns in this example. You can select the columns you want to view when the data table is displayed. See [Changing the Table Column Display](#) on page 399 for details.

The following columns are displayed:

Category This is the category for the table. The category is a standard code assigned to a group of tables by the PI 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.

By Default, no category is selected. To view the details from a particular category, select the category from 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.

Size (rows) This is the number of rows in the table.

Keep active 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:

- permanently—never deletes the data
- *n*—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.

Keep archive data for (days)	<p>This is the number of archive retention days for the data in the table. Data is discarded when it is age + 1 days old. The available archive retention options are as follows:</p> <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 archive retention period for the time type. This value is set using the Archive Retention box in the Set Table Details dialog box.
Archive On	This is a flag to indicate whether an archive table will be created or not.
Time Type	This specifies the time type of the data. The available time types are as-pollled, sub-minute, 5-minute, hourly, daily, weekly, monthly, quarterly, and yearly.
Type	This is the type of data stored in this table. The available types are raw, rate, sum, forecast, baseline, event, and forecast. For a description of the types, see Step 3 on page 442.

Viewing Property Tables

To view existing property tables in the database, select **property** from the Table Types list. The following example shows the results in the viewing pane when you select property from the Table Types list.

Table Manager tool lets you view information about various types of tables in the system. Viewing functions that are supported are sort in ascending or descending order, filter through category to time type, refresh, find, view column details etc.

Table Manager tool interface showing filters and a table of property tables.

Object Type: Table | Category: CommonPropertyTables | Retention Profiles: All | Table Types: property | Time Types: All

Alias Name	SQL Name	Type	Is Standard	Size (rows)
ComPropTab Procedure Log Property Table	KReportPackProcLog	property	<input type="checkbox"/>	2
prop_SysResAppOVPA_app_history	K_APPOVPA_APP_HISTORY	property	<input type="checkbox"/>	2
prop_AppOVPA_app	K_AppOVPA_application	property	<input type="checkbox"/>	6
SR OVPA and EPC Application by Application Property Table	K_AppOVPA_byApp	property	<input type="checkbox"/>	6
SR OVPA and EPC Application Property Table	K_Application_app	property	<input type="checkbox"/>	6
ComPropTab Customer Property Table	K_Customer	property	<input type="checkbox"/>	2
prop_Disk_Dsk	K_Disk_Dsk	property	<input type="checkbox"/>	3
ComPropTab Location Property Table	K_Location	property	<input type="checkbox"/>	2
prop_LogicalVolume_LV	K_LogicalVolume_LV	property	<input type="checkbox"/>	1
prop_NetInterface_netif	K_NetInterface_Netif	property	<input type="checkbox"/>	4
ComPropTab Node Property Table	K_Node	property	<input type="checkbox"/>	6
prop_OVPA_CPU	K_OVPA_CPU	property	<input type="checkbox"/>	2
prop_SysResOVPA_CPU_history	K_OVPA_CPU_HISTORY	property	<input type="checkbox"/>	2
prop_SysResOVPA_Disk_history	K_OVPA_DISK_HISTORY	property	<input type="checkbox"/>	2
prop_SysResOVPA_Disk	K_OVPA_Disk	property	<input type="checkbox"/>	3
prop_SysResOVPA_filesystem_history	K_OVPA_FILESYSTEM_HISTORY	property	<input type="checkbox"/>	2
prop_SysResOVPA_global_history	K_OVPA_GLOBAL_HISTORY	property	<input type="checkbox"/>	2
prop_SysResOVPA_LV_history	K_OVPA_LV_HISTORY	property	<input type="checkbox"/>	1
prop_SysResOVPA_LogicalVolume	K_OVPA_LogicalVolume	property	<input type="checkbox"/>	1
prop_SysResOVPA_NetInterface_history	K_OVPA_NETIF_HISTORY	property	<input type="checkbox"/>	2
prop_SysResOVPA_NetInterface	K_OVPA_NetInterface	property	<input type="checkbox"/>	4

The following columns are displayed:

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.
Type	This is the type of data stored in the table. In this case, the value is always property .
Is Standard	If selected, identifies a well-known table to the PI database.
Size (Rows)	This is the number of rows in the table.


Changing the Table Column Display

You can change whether a table column displays and change the order in which table columns are displayed.

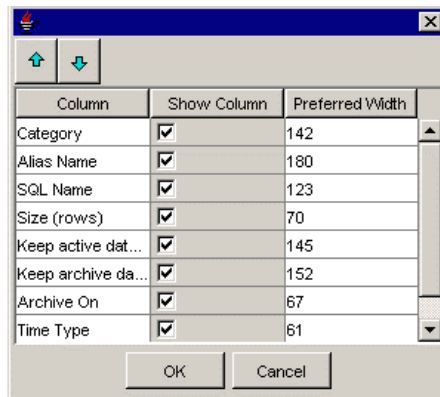
To change how table columns are displayed:



- 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 the table in Table Manager's viewing pane.

The Data Table Details dialog box opens. The **Table Info** tab opens by default.



The Data Table Details dialog box consists of three tabs: **Table Info**, **Retention Profile**, and **Index Profile**.

The Data Table Details dialog box lets you do the following:

- View information about the table's columns.
- Add a description.
- View information about the table's columns.
- View information about the corresponding indexes.
- View and manage retention profiles.

- View and manage index profiles.

Viewing Data Table Column Information

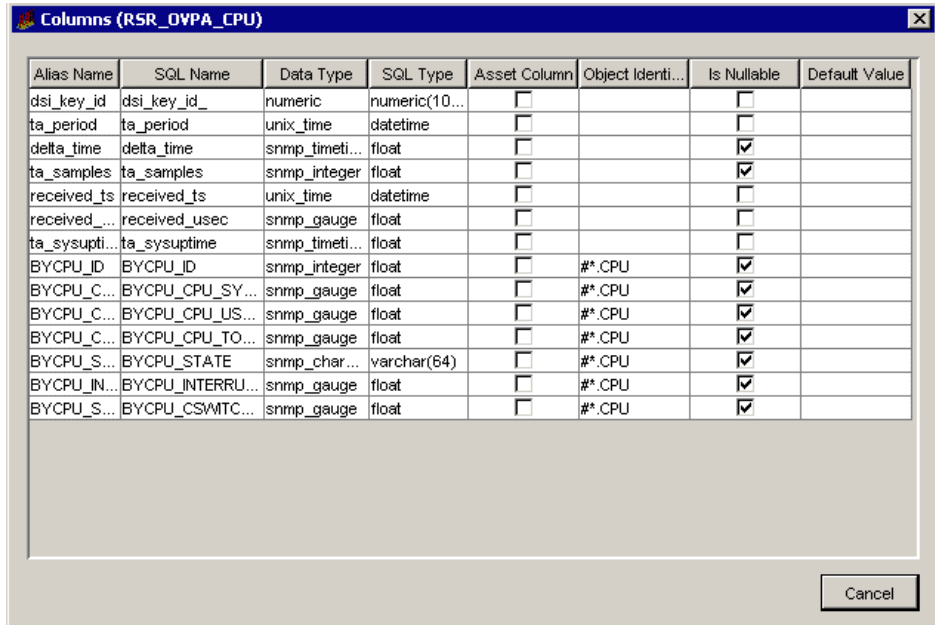
The Table Info box in the Data Table Details dialog box shows the table information for the selected table. For data tables, the Table Info box also shows the name of the property table that this data or internal table uses.

Adding a Data Table Description

The Description box in the Data Table Details dialog box lets you add a description of the table. Type 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. The Columns window opens.



The screenshot shows a dialog box titled "Columns (RSR_OVPA_CPU)" with a table of column information. The table has eight columns: Alias Name, SQL Name, Data Type, SQL Type, Asset Column, Object Identifi..., Is Nullable, and Default Value. The rows list various columns such as dsi_key_id, ta_period, delta_time, ta_samples, received_ts, received_..., ta_sysupti..., BYCPU_ID, BYCPU_C..., BYCPU_C..., BYCPU_C..., BYCPU_S..., BYCPU_IN..., and BYCPU_S... with their respective data types and properties.

Alias Name	SQL Name	Data Type	SQL Type	Asset Column	Object Identifi...	Is Nullable	Default Value
dsi_key_id	dsi_key_id_	numeric	numeric(10...	<input type="checkbox"/>		<input type="checkbox"/>	
ta_period	ta_period	unix_time	datetime	<input type="checkbox"/>		<input type="checkbox"/>	
delta_time	delta_time	snmp_timeti...	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
ta_samples	ta_samples	snmp_integer	float	<input type="checkbox"/>		<input checked="" type="checkbox"/>	
received_ts	received_ts	unix_time	datetime	<input type="checkbox"/>		<input type="checkbox"/>	
received_...	received_usec	snmp_gauge	float	<input type="checkbox"/>		<input type="checkbox"/>	
ta_sysupti...	ta_sysuptime	snmp_timeti...	float	<input type="checkbox"/>		<input type="checkbox"/>	
BYCPU_ID	BYCPU_ID	snmp_integer	float	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	
BYCPU_C...	BYCPU_CPU_SY...	snmp_gauge	float	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	
BYCPU_C...	BYCPU_CPU_US...	snmp_gauge	float	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	
BYCPU_C...	BYCPU_CPU_TO...	snmp_gauge	float	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	
BYCPU_S...	BYCPU_STATE	snmp_char...	varchar(64)	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	
BYCPU_IN...	BYCPU_INTERRU...	snmp_gauge	float	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	
BYCPU_S...	BYCPU_CSWITC...	snmp_gauge	float	<input type="checkbox"/>	#.CPU	<input checked="" type="checkbox"/>	


Cancel

The Columns window shows the following information for each column in the table.

Alias Name	A logical, unique, descriptive name of the column in the table.
SQL Name	The SQL name of the column in the table.
Data Type	This is the PI data type for the column. For a description of these types, see Step 4 on page 443.
SQL Type	The database type of the data for the column. For a description of these types, see Step 4 on page 443.
Asset Column	If selected, specifies that the column contains an asset property for use in an asset summary.
Object Identifier	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, specifies that the column can be NULL.
Default Value	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. The Indexes window opens.



The Indexes window shows the following information about each index for the table.

Index Name	The name of the index for the table.
Unique	If selected, specifies that the index is unique.

Clustered	If selected, 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.

Retention Profiles

A data retention profile for a data table is created when the data table is created in the database. The `datapipe_manager` command uses the information specified in the TEEL file to create a retention profile. The `datapipe_manager` command deletes a retention profile if no data table has reference to the retention profile. For more information about retention profiles, see [Chapter 16, Archiving Data](#).

Index Profiles

The `indexmaint` command maintains an index of existing data tables on a PI system and performs the daily tasks associated with a data or a property table. Some of the tables change infrequently. The `indexmaint` command runs daily on such tables. This adversely affects the resource utilization and processing time of the PI system.

An index profile is a group of statistics or indexing related tasks performed by `indexmaint` at the specified interval. Each data table or property table has a default index profile associated with it. To maintain indexes in an optimal manner, you can create a customized index profile using the Index Profile wizard from the PI Management Console. For more information, see [Creating an Index Profile](#).

When `indexmaint` launches on a table, it performs the tasks associated with the index profile of the table. If no customized index profile is assigned to the table, `indexmaint` performs the tasks listed in the default index profile for that table.

The default index profile performs the following tasks for each type of table:

- Gathers statistics of all rate, hourly, and daily tables everyday.

- Gathers statistics of all weekly tables every Monday.
- Gathers statistics of all monthly, quarterly, and yearly tables every second day of the month.

If `indexmaint` launches without a table name, it performs the tasks listed in the default index profile of all the data and property tables.

► For detailed documentation about the `indexmaint` command, see the *HP Performance Insight Reference Guide*.

Managing Index Profiles

The Index Profile wizard lets you to perform the following tasks:

Viewing and Assigning an Index Profile

The Index Profile tab in the Data Table Details dialog box shows the current index profile settings for the selected data table.

Statistics	Frequency	Launch Day	Index Task	Frequency	Launch Day
GatherStatisticsT...	Daily	All Days	VoidTask	Daily	All Days

To view and assign an existing index profile, follow these steps:

- 1 Click **View Index Profiles**.

The Index Policy Profile dialog box opens. The dialog box shows a list of existing index profiles and the following information:

Name	Name of the index profile.
Description	Description of the index profile.
Task List	The list of tasks for the selected index profile.

- 2 To assign an index profile to the table, select an index profile from the list, and then click **Assign Index Profile**.
- 3 When you are finished, click **Close**.

Creating an Index Profile

The Index Profile wizard lets you create a new index profile. An index profile can have one or more statistics or indexing related task.

To create an index profile, follow these steps:

- 1 Double-click a table in the Table Manager's viewing pane.
The Data Table Details dialog box opens.
- 2 Click **Index Profile** tab, and click **Create**.

The Index Policy Profile dialog box opens.

Index Policy Profile

Welcome to Index Policy Profile Editor

Name

Description

Statistics Related Tasks Indexing Related Tasks

Statistics Tasks: Index Tasks:

Daily Weekly Monthly Daily Weekly Monthly

Frequency Frequency

Add

Task List

Statistics	Frequency	Launch Day	Index Task	Frequency	Launch Day
------------	-----------	------------	------------	-----------	------------

Delete

Cancel Finish

- 3 In the Name box, type a name for the index profile.

The name cannot exceed 64 characters. It can contain the following characters: alphabetic, numeric, an underscore (_), or a hyphen (-).

- 4 In the Description box, type a description for the index profile.

A description is optional. If you specify a description, it cannot exceed 255 characters.

- 5 Create Statistics and Indexing Tasks.

Statistics Task:

- a Select the **Statistics Related Tasks** check box if you want to analyze the table.
- b From the Statistics Tasks list, select one of the following options:
 - **GatherStatisticsTask**: task to gather statistics is created.

- **AnalyzeTableTask**: task to analyze the table is created.
- c Select one of the following options:
 - **Daily**. The Frequency list is unavailable and data for the table is collected daily. If you want to create an indexing task go to [step a](#); otherwise proceed to [step 6](#).
 - **Weekly**. Data for the table is collected weekly. The Frequency list is populated with the days of the week. Proceed to [step d](#).
 - **Monthly**. Data for the table is collected monthly. The Frequency list is populated with values from **1** through **31** and **last day**. Proceed to [step d](#).
- d From the Frequency list, select an appropriate value.

Indexing Task:

- a Select the **Indexing Related Tasks** check box if you want to rebuild or coalesce an index.
 - b From the Index Tasks list, select one of the following options:
 - **RebuildIndexTask**: task to rebuild the index is created.
 - **CoalesceIndexTask**: task to coalesce index is created.
 - c Select one of the following options:
 - **Daily**. The Frequency list is unavailable and the index task occurs daily. Proceed to [step 6](#).
 - **Weekly**. The index task occurs weekly. The Frequency list is populated with the days of the week. Proceed to [step d](#).
 - **Monthly**. The index task occurs monthly. The Frequency list is populated with values from **1** through **31** and **last day**. Proceed to [step d](#).
 - d From the Frequency list, select an appropriate value.
- 6 Click **Add**.
- If the index profile creation is successful, the task appears in the Task List section.
- 7 To create more tasks repeat [step 4](#) through [step 6](#).
- 8 Click **Finish**.

An Index profile with one or more tasks is created and assigned to the table. The `indexmaint` command uses the profile that you assign to schedule and run the tasks related to the table.

Editing an Index Profile

Editing an index profile involves adding new tasks or deleting existing tasks associated with the index profile.

To edit an existing index profile, follow these steps:

- 1 Double-click the table in the Table Manager's viewing pane.
The Data Table Details dialog box opens.
- 2 Click the **Index Profile** tab, and then click **Edit**.
The Index Policy Profile dialog box opens.
- 3 Do either or both of the following:
 - To add new tasks, perform [step 5](#) and [step 6](#) listed in the [Creating an Index Profile](#) section.
 - To delete existing tasks, select the task you want to delete from the Task List section, and then click **Delete**.
- 4 Click **Finish**.

You cannot edit the name of the index profile.

Viewing Property Table Details

To view detailed information about any property table, follow these steps:

- 1 In the Table Manager's viewing pane, select **Property** from the Table Types list.

The viewing pane displays all the property tables.

- 2 Double-click a property table in the viewing pane.

The Property Table Details dialog box opens.

Property Table Details

Table Info

Alias Name: SR OVPA and EPC Application by Application Property Table
SQL Name: K_AppOVPA_byApp
Group Member Table: K_AppOVPA_byApp_mbr
Keymap Table: K_AppOVPA_byApp_keymap

Child Property Tables:

KV_AppOVPA_byApp

Related Data Tables:

SH_AppOVPA_byApp
SD_AppOVPA_byApp

Index Profile

Name: DefaultpropertyProfile
Description: index profile for all property tables

Statistics	Frequency	Launch Day	Index Task	Frequency	Launch Day
GatherStatistic...	Daily	All Days	VoidTask	Daily	All Days

View Index Profiles **Create** Edit

Description: id EPC application metrics by application

Columns ... Indexes ... OK Cancel

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.
- View or managing an index profile.

Viewing Property Table Information

The Table Info box in the Property Table Details dialog box (see figure on [page 409](#)) contains the following information for property tables.

Alias Name	A logical, unique, descriptive name of the table in the database.
SQL Name	The SQL name of the table.
Group Member Table	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	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:

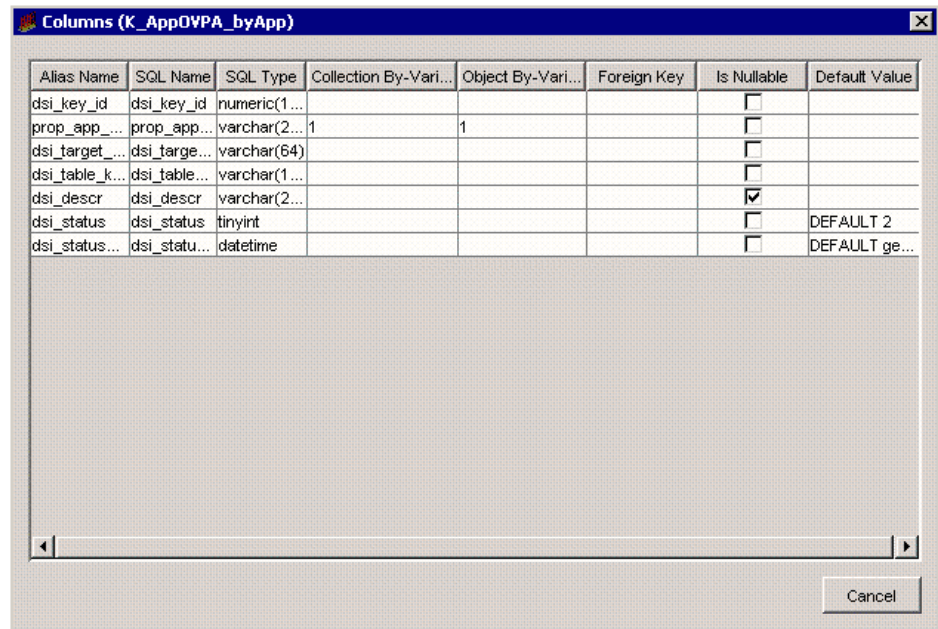
Child Property Tables	A list of property tables that reference this property table.
Related Data Tables	A list of data tables that use this property table.

Adding a Property Table Description

The Description box in the Property Table Details dialog box lets you add a description of the table. Type 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. The Columns window opens.



The screenshot shows a window titled "Columns (K_AppDVPA_byApp)" with a table of column information. The table has the following columns: Alias Name, SQL Name, SQL Type, Collection By-Vari..., Object By-Vari..., Foreign Key, Is Nullable, and Default Value. The data rows are as follows:

Alias Name	SQL Name	SQL Type	Collection By-Vari...	Object By-Vari...	Foreign Key	Is Nullable	Default Value
dsi_key_id	dsi_key_id	numeric(1...				<input type="checkbox"/>	
prop_app_...	prop_app_...	varchar(2...	1	1		<input type="checkbox"/>	
dsi_target_...	dsi_targe...	varchar(64)				<input type="checkbox"/>	
dsi_table_k...	dsi_table...	varchar(1...				<input type="checkbox"/>	
dsi_descr	dsi_descr	varchar(2...				<input checked="" type="checkbox"/>	
dsi_status	dsi_status	tinyint				<input type="checkbox"/>	DEFAULT 2
dsi_status...	dsi_statu...	datetime				<input type="checkbox"/>	DEFAULT ge...

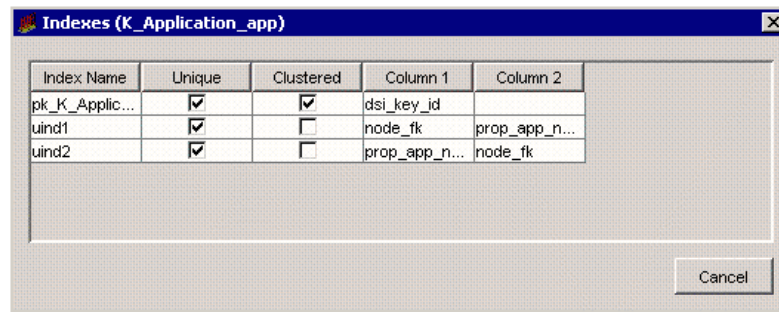
The Columns window shows the following information for each column in the table.

Alias Name	A logical, unique, descriptive name of the column in the table.
SQL Name	The SQL name of the column in the table.
SQL Type	The database type of the data for the column. For a description of these types, see Step 4 on page 443.
Collection By Variable	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	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** Identifies the parent table and the column name that contain the data if the column is a foreign key. The format is *parent_table.column_name*.
- Is Nullable** If selected, 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, the Is Nullable column cannot be selected.
- Default Value** 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. The Indexes window opens.



The Indexes window shows the following information about each index for the table.

- Index Name** The name of the index for the table.
- Unique** If selected, indicates that the index is unique.
- Clustered** If selected, 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, only one index will appear in the list. If the by-variables have different columns in the indexes, at least two indexes will appear in the list. Also, a column appears for every element in an index.

Viewing or Managing an Index Profile

The Index Profile box in the Property Table Details dialog box contains the following information about the index profile associated with the property table.

Name	Name of the index profile.
Description	Description of the index profile.
Task List	The list of tasks for the selected index profile.

For more information about index profiles, see [Index Profiles](#) on page 403.

Displaying Table Data

Table Manager's **Display Table Contents** option lets you examine the data in the PI 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.

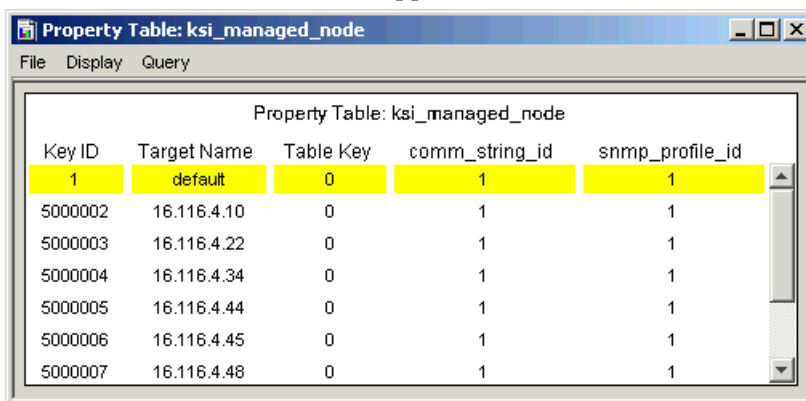
You can change the type of data displayed and its appearance, and save it to the PI report formats (.rep or .srep). You can view the saved reports using Web Access Server, Report Builder, or Report Viewer. Additionally, you can print these reports, and export the data to different formats (CSV and ASCII) 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 following menus: File, Display, and Query. Using the menu options, you can modify how the data will appear in the report.

Table 9 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 10 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 11 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.
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.

Table 11 Query Menu Options

Option	Description
Change Constraints	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.

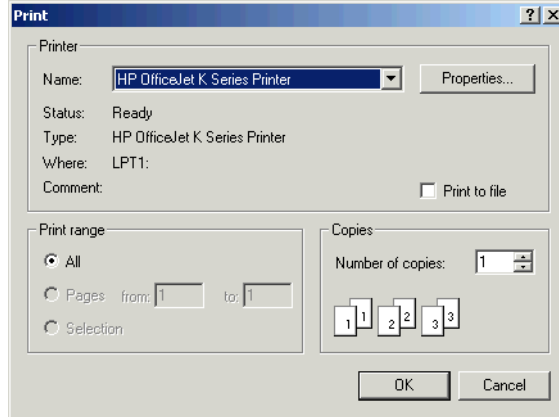
Printing Table Data

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 PI report formats:

- PI Dataset (.srep)


- PI Report Definition (.rep)

You can view the saved report using Web Access Server, Report Builder, or Report Viewer.

To save a table data:

- 1 Do either of the following:
 - To save a report with its current file name, select **Save** from the File menu or press **CTRL+S**.
 - To save a report with a different file name or format, select **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:
 - Type 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 ensures that the entire dataset you constrained will be used.
- 2 From the File menu, select the **Export** option.



- 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

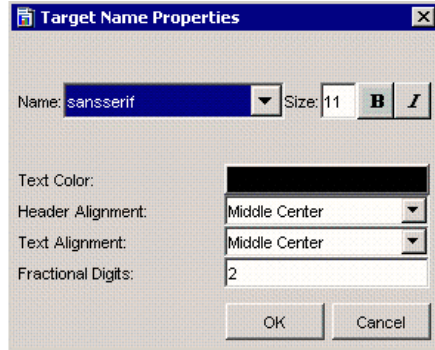
You can change how table data displays by using the options on the Display menu to change the format of the data's statistics, cells, columns, and rows.

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.

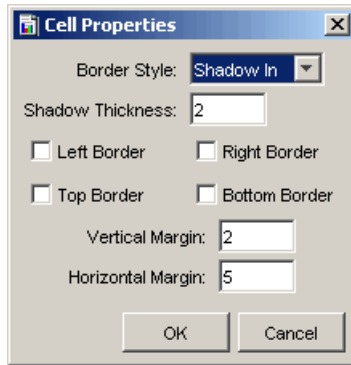


- 3 Do the following:
- a Select a font from the Name list to change the font used for the statistic values.
 - b Type a size for the font used for the statistic values in the Size box.
 - c Click **Bold** to bold the statistic values.
 - d Click **Italics** to italicize the statistic values.
 - e Click the Text Color box to display the Select a Color dialog box, which lets you select a color for the statistic values.
 - f To position the statistic's name in the heading of the table column, select an alignment option from the Header Alignment list.
 - g To position the statistic text that appears in the column below the statistic's heading, select an alignment option from the Text Alignment list.
 - h Type a digit to set the number of fractional digits for the numeric columns.
- 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**.



- 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:

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**.



- 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.

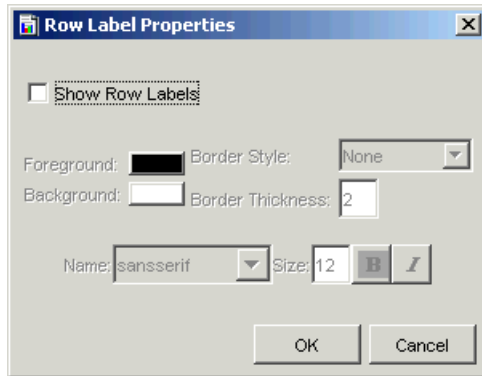
By default, this is set to **None**, which shows no border for the column heading boxes. If you select **Shadow Out**, for example, the column heading box looks like the following:


 - 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 **Bold** to bold the column heading text.
 - h Click **Italics** 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**.



The Row Label Properties dialog box is almost identical to the Column Properties dialog box.

- 2 See the steps for the procedure in [Changing Column Properties](#) on page 425 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: 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.

PI'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 . 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 .
- 5 To rearrange the order of the statistics in the list, do *one* of the following:
 - Click  to move a statistic up in the list.
 - Click  to move a statistic down in the list.
- 6 Click **OK**.

Changing the Time Period

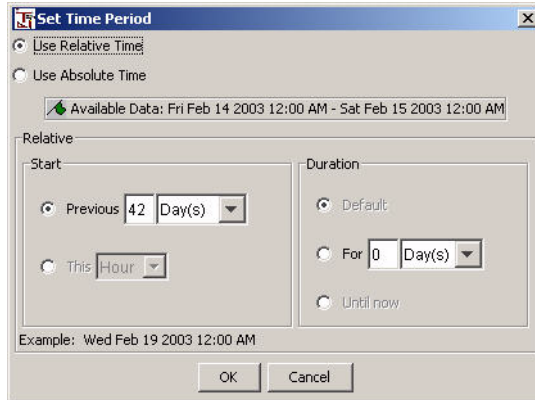
You can change the time period for data displayed in a report. You can select an absolute period (specific dates) or a relative time period (a period of time relative to the time you view the data).

Selecting an Absolute Time

An absolute time period identifies 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 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 **Browse** located next to the Start Time box or the End Time box. A window similar to this one opens:



- Click the down arrow to display a calendar, similar to the one below:



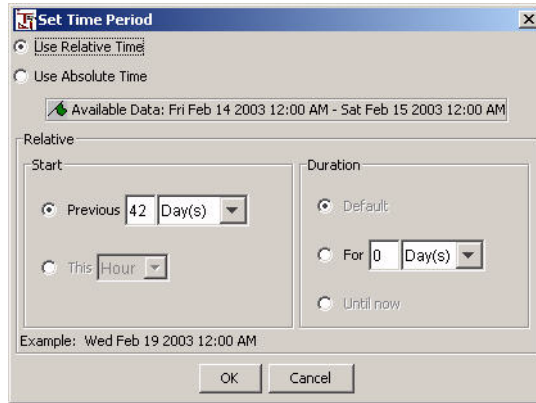
- Click the Month and Year arrows to select the month and year you want.
- Click the *Hour*, *Minutes*, *Seconds*, and *Time of Day* arrows to select the time.
- Click a date in the calendar.
After you click a date, the calendar window closes, and the selected date and time display in the Start Time (or End Time) window. You can click the down arrow to reopen the calendar if you want to make changes.
- 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 opens.



- 2 Click **Use Relative Time**.
- 3 In the Start box, do *one* of the following:
 - To select a current time period, click **This time period**. 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 **Previous n time period**. Type a number in the box, and 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. Type a number in the box, and 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**.

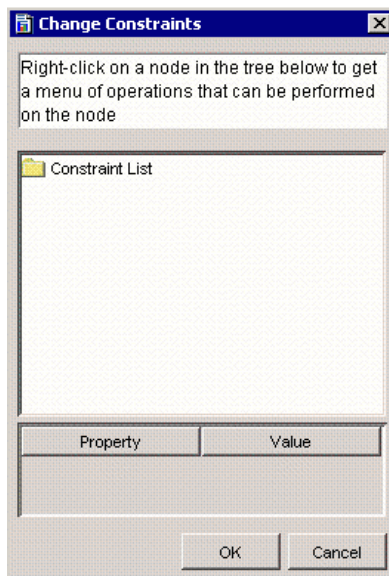
Adding Constraints

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.

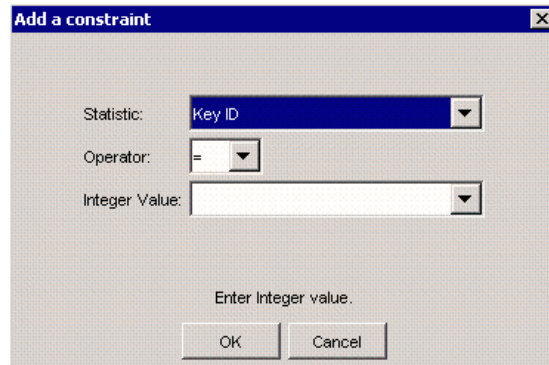
Add a Basic Constraint

To add a basic constraint:

- 1 From the Query menu, select **Change Constraints**.



- 2 Right-click **Constraint List**, and click **Add Constraint** on the shortcut menu.



- 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 select a parameter, if available. (See *HP Performance Insight Guide to Building and Viewing Reports* for more information about creating a parameter.)
- 4 Click **OK**.

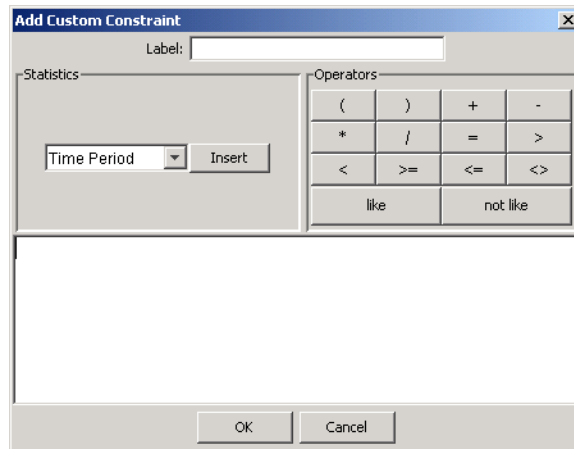
Adding a Custom Constraint

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 **Change Constraints**. The Change Constraints dialog box opens. See the figure on [page 435](#).
- 2 Right-click **Constraint List**, and click **Add Constraint** on the shortcut menu. The Add a Constraint dialog box opens. See the above figure.

- 3 Right-click **Constraint List**, and click **Add Custom Constraint** on the shortcut menu.



- 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 Type 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.

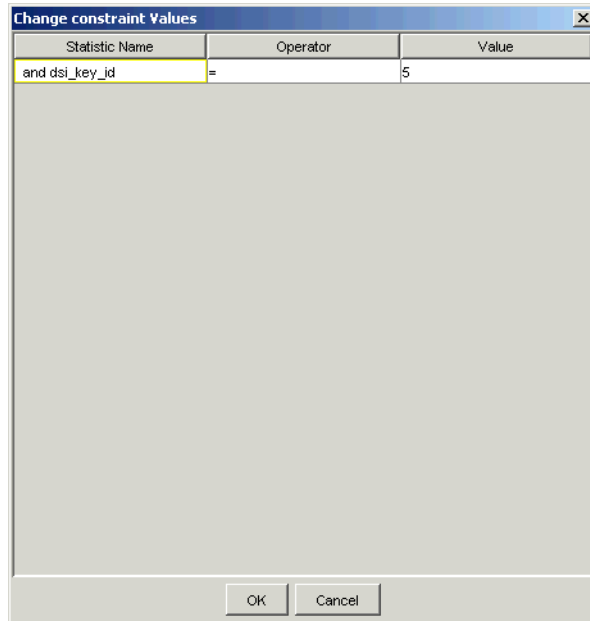
To add an AND or OR clause:

- 1 From the **Query** menu, select **Change Constraints**. The Change Constraints dialog box opens. See figure on [page 435](#).
- 2 Right-click **Constraint List**, and click **Add And** or **Add Or** on the shortcut menu.

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. The Change Constraint Values dialog box opens.



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, see the *HP 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 **Change Constraints**. The Change Constraints dialog box opens. See figure on [page 432](#).
- 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 have your changes take effect.

Deleting Constraints and Clauses

To delete a constraint or clause:

- 1 From the **Query** menu, select **Change Constraints**. The Change Constraints dialog box opens. See figure on [page 432](#).
- 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.

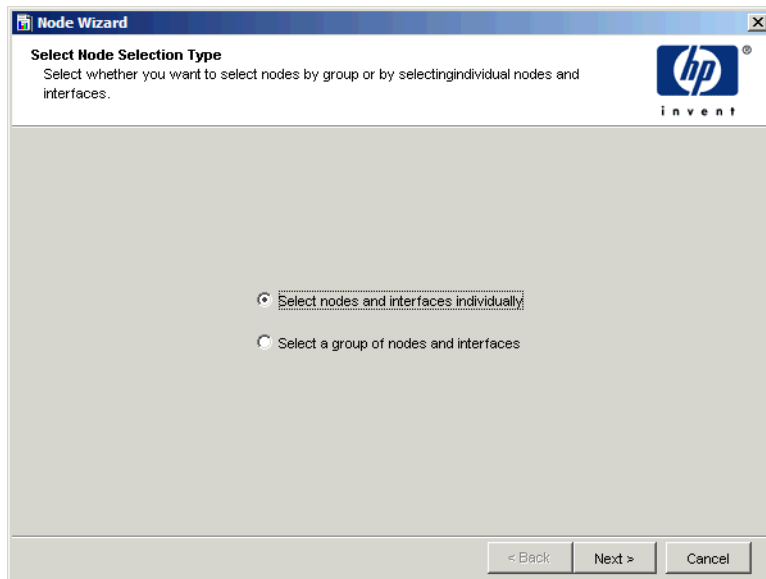
Selecting Nodes and Interfaces

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 [Step 3](#) on page 438.
- 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 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, see the *HP 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, PI 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**.



- 2 In the Max Rows box, type the number of rows you want, and click **OK**.

Creating Tables

Using Table Manager, you can create custom data tables and property tables. You can create collectible data tables. To create an SNMP collectable table, use the MIB Browser. See [Creating a Collection Table](#) on page 656 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 440. If you want to create a property table, see [Creating Property Tables](#) on page 453.

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. You must allow for 13 default columns when you select the objects for the data table. If you choose to define a property table, you must allow for seven default columns 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, see the *HP 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: Specify a collection name, collection category, and a collection description.

The Define Collection window opens.

Collection Definition Wizard

Define Collection

Enter a collection name that will be used to create the set of tables in the database. The name must be 30 characters or less. The collection category can be selected from the existing list or you can enter a new name to be used for grouping your collections. You may optionally enter a description up to 255 characters.

Collection Name:

Collection Category:

Collection Description:

< Back Next > Cancel

- 1 In the Collection Name box, type an alias name for the data table. This name is stored in the PI data repository.



The wizard will assign a name to the collection based on the alias name you provide.

The name must be unique and can be up to 255 characters. Valid entries include uppercase and lowercase alphabetic characters, all numeric digits, and the special characters hyphen (-) and underscore (_).

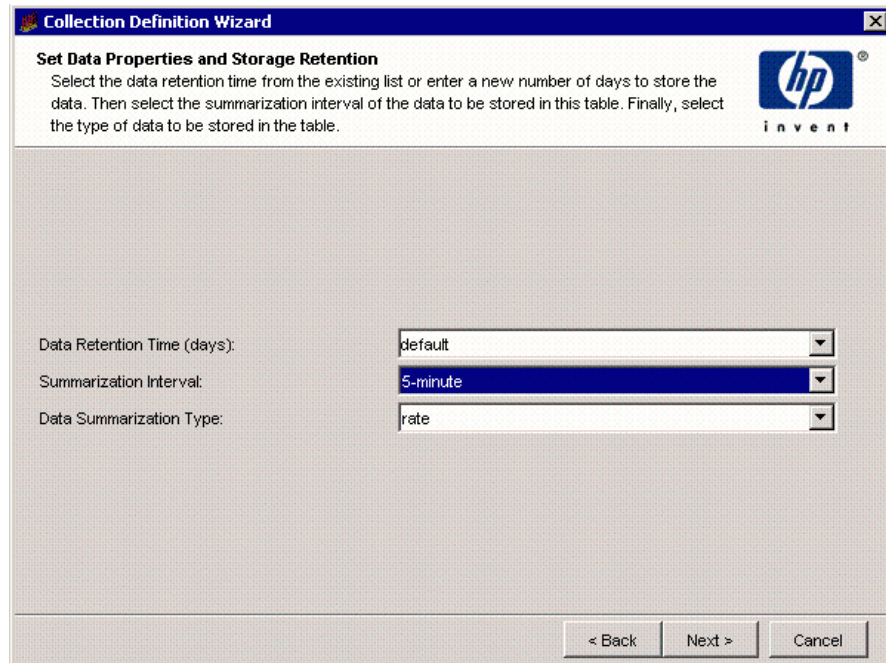
- 2 In the Collection Category box, type or select the category for the collection table.

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.

- 3 In the Collection 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: Set the data properties and storage retention for the collection table.

The Set Data Properties and Storage Retention window opens.



- 1 In the Data Retention Time box, type or select the number of days you want to keep the data in the table.
- 2 In the Summarization Interval box, select a time characteristic for the data in the table. For example, a 5 minute table contains data for a 5 minute interval. Data collected at 5 minute intervals is called rate data.
- 3 In the Data Summarization Type box, select a table type.

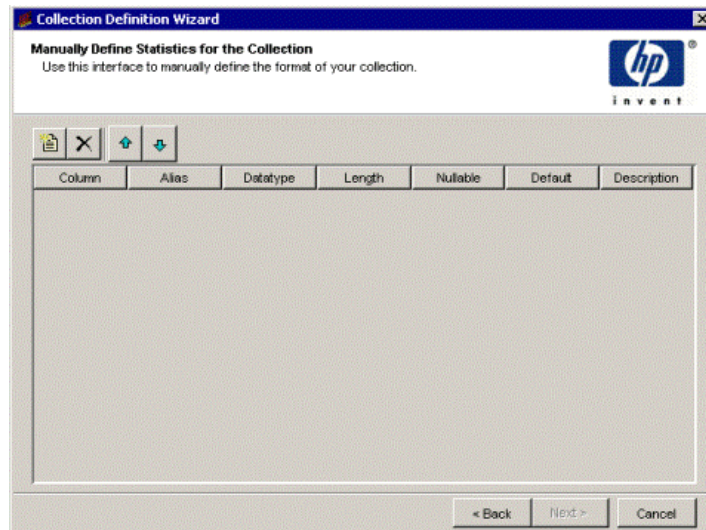
The table type specifies the class of data contained in the table. Following is the list of available table type options:


- 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.

4 Click **Next**.

Task 4: Define the statistics for the collection.

The Manually Define Statistics for the Collection page opens.



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:

PI 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

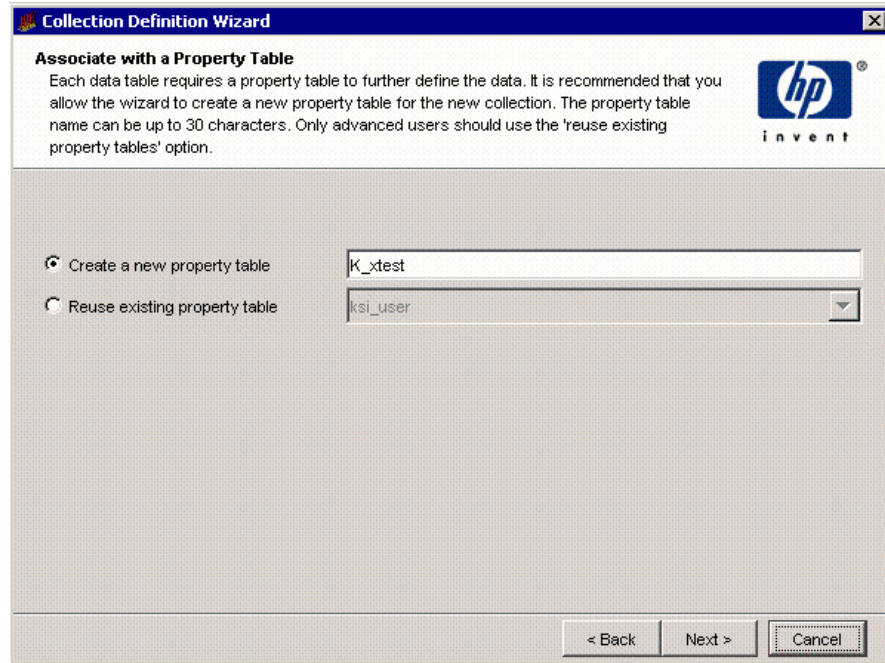
- 5 In the Length box, type a numeric value if needed.
Valid data types for length values are char_string, hex_string, numeric, octet_string, snmp_char_string, snmp_hex_string, or snmp_octet_string.
- 6 Click **Nullable** if the statistic can be NULL.
- 7 In the Default box, type the default value for 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.



Collection Definition Wizard

Associate with a Property Table

Each data table requires a property table to further define the data. It is recommended that you allow the wizard to create a new property table for the new collection. The property table name can be up to 30 characters. Only advanced users should use the 'reuse existing property tables' option.

Create a new property table

Reuse existing property table

K_xtest

ksi_user

< Back Next > Cancel

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.

Key	FK	Column	Alias	Datatype	Data Length
<input checked="" type="checkbox"/>	<input type="checkbox"/>	dsi_target_n...	dsi_target_n...	snmp_char_...	64
<input checked="" type="checkbox"/>	<input type="checkbox"/>	dsi_table_key	dsi_table_key	snmp_char_...	128
<input type="checkbox"/>	<input type="checkbox"/>	dsi_descr	dsi_descr	snmp_char_...	255

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 444.

- 5 In the Data Length box, type a numeric value if needed.

Valid data types for length values are `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, 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:

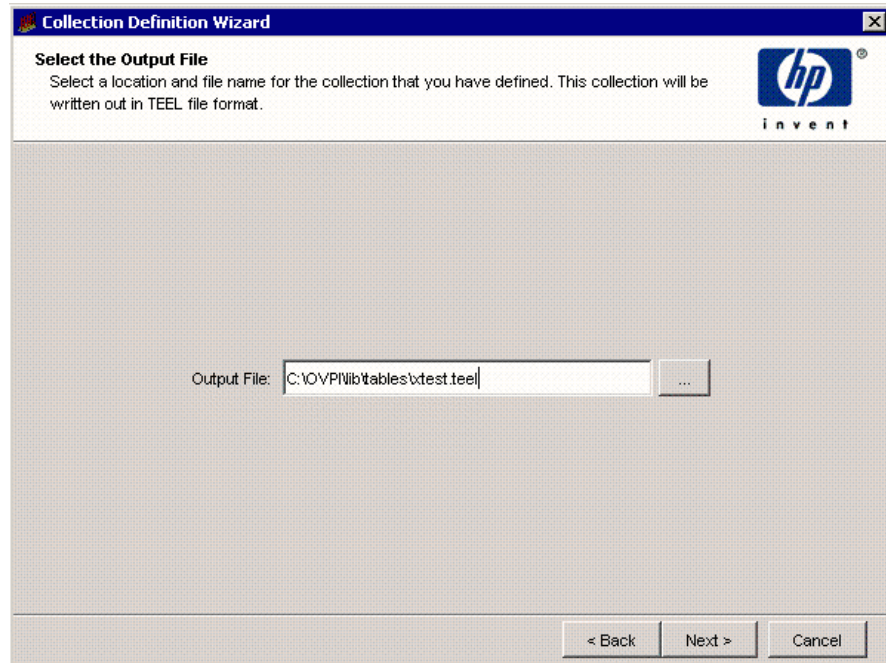
- To move the column up, click .
- To move the column down, click .
- To delete a column from the property table, click .




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.

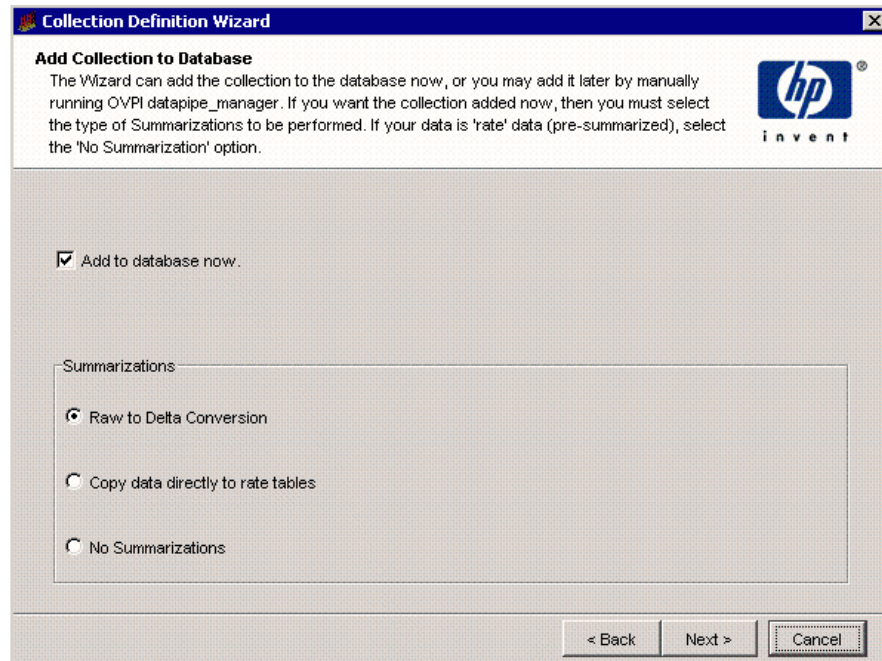
The Select the Output File page opens.



- 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.



- 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 type 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 [Step 7](#) on page 450.

If you need additional information about the options for the `datapipe_manager` command, see the *HP Performance Insight TEEL Reference Guide* and the *HP 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.

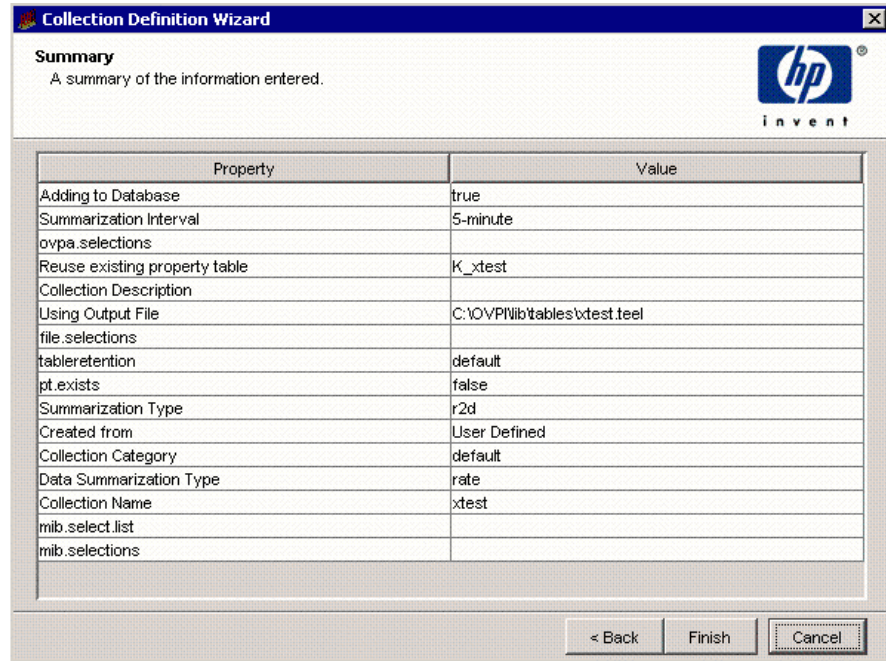
- **No Summarizations**

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 PI 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: Type a name for the property table.

The Create Property Table page opens.

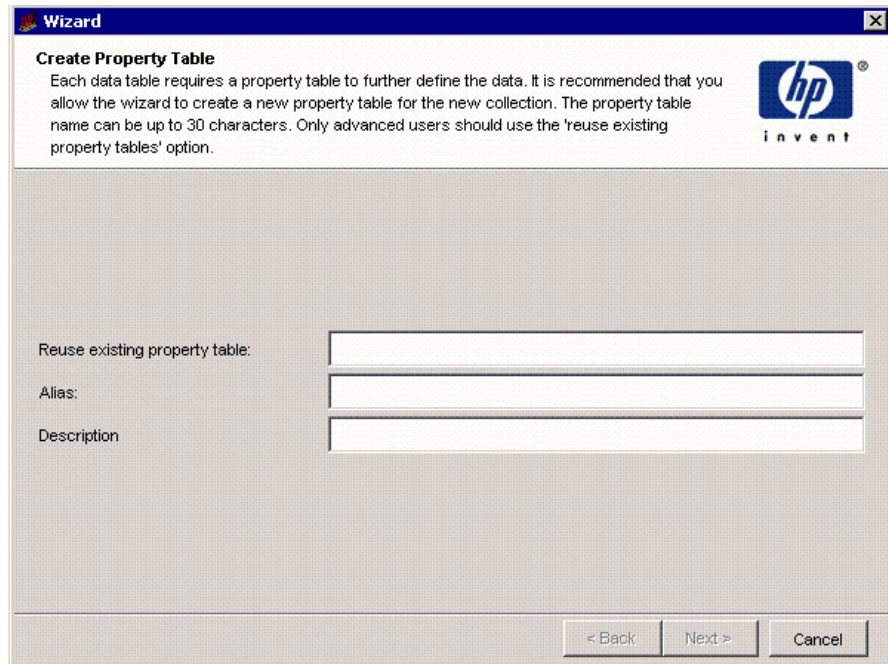


Figure 1 Create Property Table Page

- 1 In the Reuse existing 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 type 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.

To complete this task, see [Step 6](#) on page 447.

Task 4: Specify an output file for the collection.

The Select the Output File page opens.

To complete this task, see [Step 7](#) on page 450.

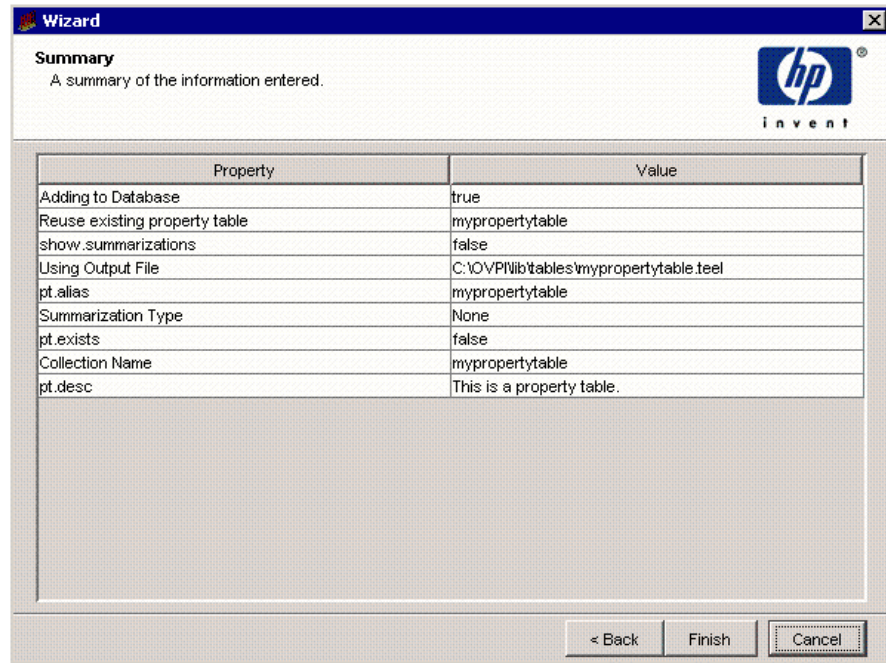
Task 5: Register the collection with the database.

The Add Collection to Database page opens ([page 451](#)).

To complete this task, see [Step 7](#) on page 450.

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 PI creates the TEEL file and stores it in the designated location.

Setting Table Defaults

Table Manager lets you set defaults for 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.




Care should be exercised before changing any of these values. When you modify the default values for data retention, the data in the installed PI reporting solutions will change. If you increase the default values, it will negatively affect overall system performance.

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 PI will keep the data for a particular table. It allows PI 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 PI 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. 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 PI 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: 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 PI optimize its searches of the database. See the *HP Performance Insight Reference Guide* for more information about tpmaint.




Care should be exercised before changing any of these values. When you modify the default values for data retention, the data in the installed PI reporting solutions will change. If you increase the default values, it will negatively affect overall system performance.

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**.

14 Using Catalog Manager

Using the Management Console's Catalog Manager application, you can create a *report link* in a deployed report that will open another report, a third-party application, or a Uniform Resource Locator (URL) from the deployed report.

A report link takes specific information from an element (graph, table, image, or label) in the source report, uses it as the input for a report parameter (in the case of a report to report link) or to an argument (in the case of a report to application or URL link), and displays the report, application, or URL with the specified information when the link is opened from the report. Thus, you can connect reports, URLs, and applications together to create navigation paths that build an intelligent and customer-specific flow of information.

In addition, you can manage report links and items that you deploy. Catalog Manager contains a Links and a Deployed Items catalog. These catalogs reside on the Web Access Server. The contents of the catalogs are as follows:

- Links. Contains all of the report links that you create in the deployed reports located in the Deployed Items catalog.
- Deployed Items. Catalog contains the following:
 - Deployed reports. You *deploy* reports to this catalog to make them available for users to view using PI's Web Access Server client application. You can deploy reports using Catalog Manager, Package Manager, and the Deployment wizard.
 - Deployed URLs. Deploying a URL in the Deployed Items catalog is useful when you will create a link from a report to a particular URL frequently, or when users need to be able to view a particular URL directly from the Web Access Server.

Additionally, you can perform management tasks specific to the items contained in each catalog.

Users can trigger report links when they view the reports using the Web Access Server, Report Viewer, and Report Builder.



Users cannot open report to application links using the Web Access Server. Instead, they must use Report Builder or Report Viewer to open these kinds of links. Additionally, the application must reside on the system from which Report Builder or Report Viewer are run to successfully invoke the application linked to the report.

Starting Catalog Manager

To start Catalog Manager:

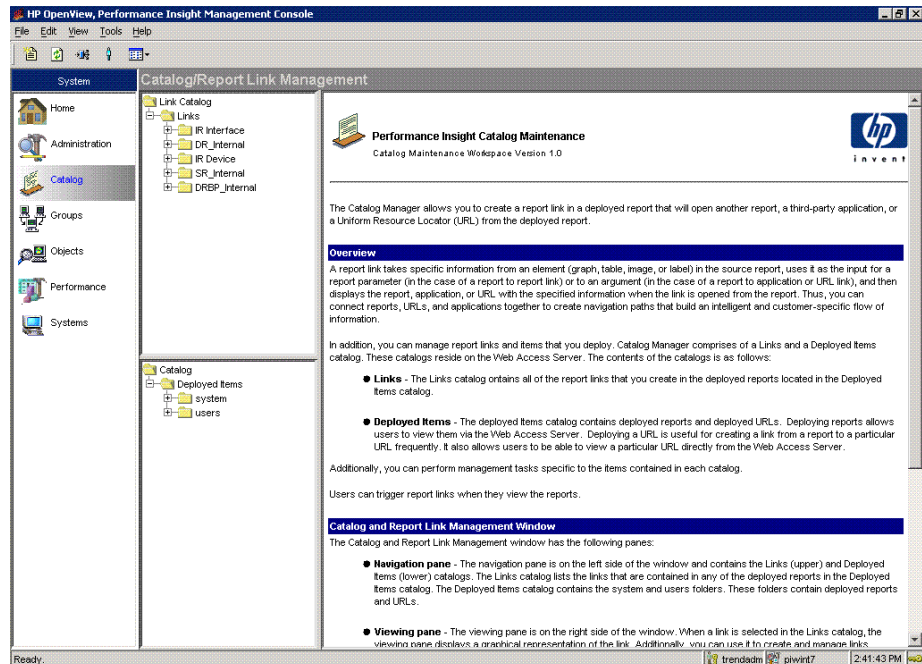
- 1 Start the Management Console. See [Starting the Management Console](#) on page 118 for details.



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 525.

- 2 Click the **Catalog** icon in the Management Console's navigation pane.

The following is an example of how the Catalog/Report Link Management window will appear.



Catalog/Report Link Management Window

The Catalog/Report Link Management window has the following panes:

- Navigation pane. The navigation pane is on the left side of the window and contains the Link and Deployed Items catalogs. The Link catalog contains the `Links` folder, which lists the links that are contained in any of the deployed reports in the Deployed Items catalog.

The Deployed Items catalog contains the `Deployed Items` folder. It contains the `system` and `users` folders. These folders contain deployed reports and URLs.

- Viewing pane. The viewing pane is on the right side of the window. When a link is selected in the Links catalog, it displays a graphical representation of the link. You can use it to create and manage links. See [Modifying Links in the Viewing Pane](#) on page 505 for more information.

Managing the Deployed Items Catalog

You can perform the following tasks on the Deployed Items catalog:

Deploying Reports

Deploying a report places the report (a `.rep` or `.srep` file) in either the `system` or `users` folder of the Deployed Items catalog.

You must deploy a report in order for users to view the report using the Web Access Server. You must deploy reports before you can create links in them. Catalog Manager's Report Deployment wizard guides you through each step necessary to deploy reports to the catalog.

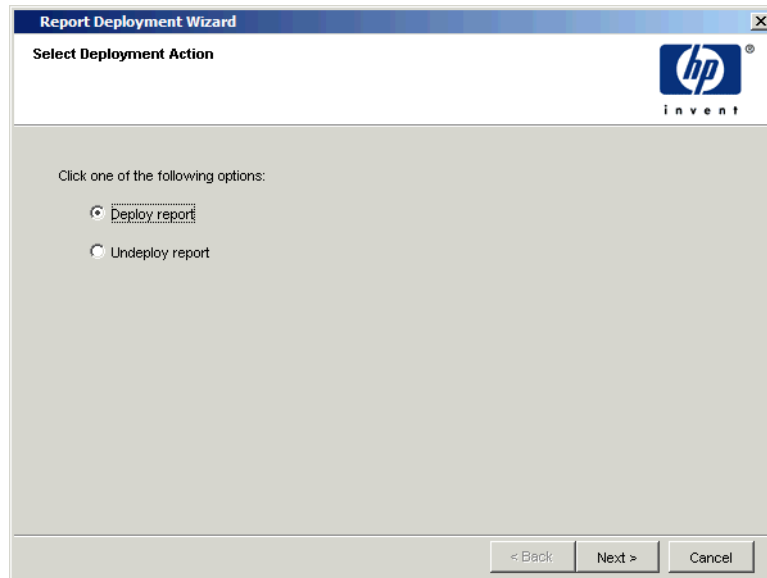
To deploy a report:

Task 1: Start the Deployment Wizard.

From Catalog Manager's File menu, select **Deploy**.

Task 2: Select the deploy action.

The Select Deployment Action window opens.



Select **Deploy report** and click **Next**.

Task 3: Specify the information needed to deploy the report to the Web Access Server.

The screenshot shows a dialog box titled "Report Deployment Wizard" with the subtitle "Enter the Application Server Information". The HP logo and the word "invent" are in the top right corner. The dialog contains two text input fields: "Web Access Server Host" with the value "ren.zko.hp.com" and "Web Access Server Port" with the value "443". Below these fields is a checked checkbox labeled "Communicate over SSL". At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

By default, the Web Access Server Information window displays the name and port number of the current Web Access Server.

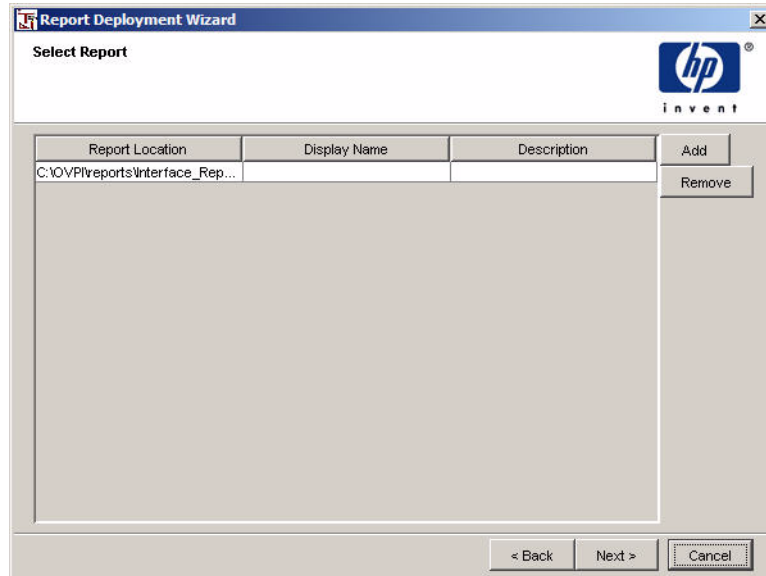
Do the following:

- 1 In the Web Access Server Host box, leave the default value supplied or type the name of the Web Access Server where you want to deploy the report.
- 2 In the Web Access Server Port box, leave the default value supplied or type the port number of the Web Access Server specified in the Web Access Server Host box.
- 3 Do *one* of the following:
 - Select the **Communicate over SSL** check box to enable SSL.
 - Click to clear the **Communicate over SSL** check box to disable SSL.

The Communicate over SSL option enables Secure Socket Layer (SSL) communication between the Web Access Server and any supported Web browser. When this option is not used, the communication between the two is over HTTP.

- 4 Click **Next** to continue.

Task 4: Select the report or reports you want to deploy.



- 1 Click **Add**.

The Open dialog box opens. Navigate to the folder that contains the report or reports you want to deploy.

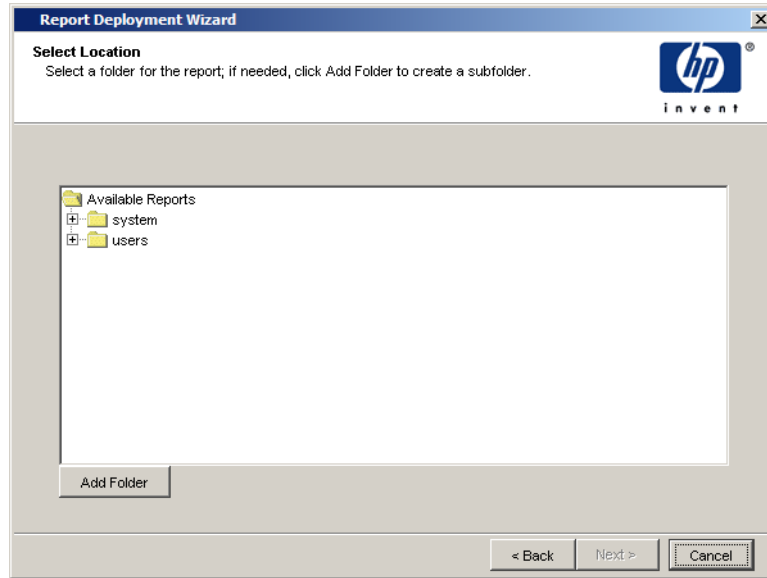
- 2 Do any of the following in the Open dialog box, then click **Open**:

- To select a single report, click the report file name.
- To select consecutive reports, click the first report, press and hold down the **SHIFT** key, and click the last report.
- To select reports that are not consecutive, press and hold down the **CTRL** key, and click each report.

- 3 If you need to remove a report that you added, select it, and click **Remove**.

4 Click **Next**.

Task 5: Select a folder in which to deploy a report or reports.



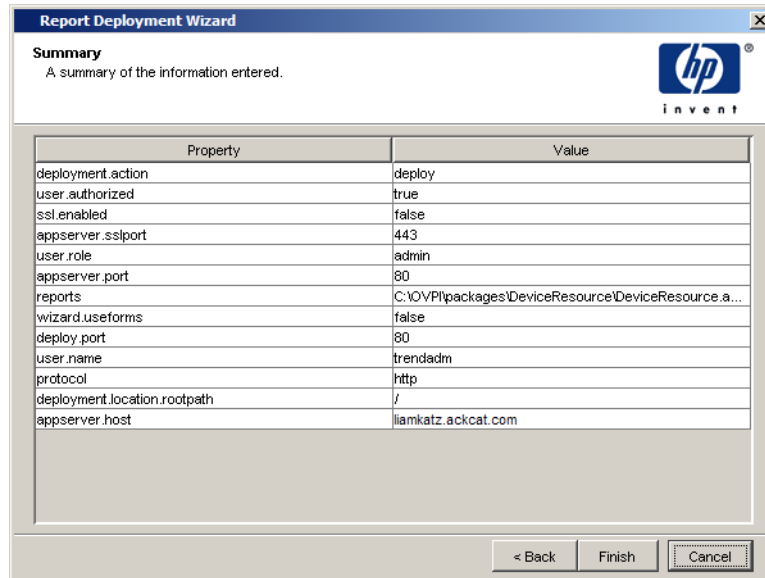
Do one of the following:

- Expand the tree, and select the folder where you want to store the report(s), and click **Next**.
- Expand the tree, and select the folder where you want to create a new folder, then click **Add Folder**.

A New Folder entry appears in the tree. Do the following:

- a Click the New Folder entry and type the name of the new subfolder.
- b Click **Next**.

Task 6: Deploy the reports.



- 1 Click **Finish**.

If the report is deployed successfully, a confirmation message displays.

- 2 Click **OK**.
- 3 To see that the report was deployed, select the folder to which you deployed the report, and click **Refresh** in the toolbar.

Undeploying Reports

You can undeploy a report by using either the Undeploy option, which is available from the shortcut menu in the Deployed Items catalog, or by using the Report Deployment wizard. Undeploying a report removes it from the Deployed Items catalog.

Undeploying a Report Using the Delete Menu Option

To undeploy a report:

- 1 In the Deployed Items catalog, navigate to the folder (`system` or `users`) that contains the report you want to undeploy.
- 2 Right-click the report.
- 3 Select **Delete** from the menu.

Undeploying a Report Using the Report Deployment Wizard

To undeploy a report:

Task 1: Start the Deployment Wizard.

From Catalog Manager's File menu, select **Deploy**.

Task 2: Select the deploy option.

The Select Deployment Action window opens. See [page 465](#).

Select **Undeploy report** and click **Next**.

Task 3: Specify the information needed to undeploy the report from the Web Access Server.

The Web Access Server Information window opens. See [page 466](#). By default, the window displays the name and port number of the current Web Access Server.

Do the following:

- 1 In the Web Access Server Host box, leave the default value supplied or type the name of the Web Access Server from where you want to undeploy the report.
- 2 In the Web Access Server Port box, leave the default value supplied or type the port number of the Web Access Server specified in the Web Access Server Host box.
- 3 Do *one* of the following:
 - Select the **Communicate over SSL** check box to enable SSL.
 - Click to clear the **Communicate over SSL** check box to disable SSL.

The Communicate over SSL option enables Secure Socket Layer (SSL) communication between the Web Access Server and any supported Web browser. When this option is not used, the communication between the two is over HTTP.

- 4 Click **Next** to continue.

Task 4: Select the report you want to undeploy.

The Select Location window opens. See [page 468](#).

- 1 Navigate to the folder that contains the report or reports you want to undeploy.
- 2 Select a report or reports by doing any of the following:
 - To select a single report, click the report file name.
 - To select consecutive reports, click the first report, press and hold down the **SHIFT** key, and click the last report.
 - To select reports that are not consecutive, press and hold down the **CTRL** key, and click each report.
- 3 Click **Next**.

Task 5: Undeploy the reports.

The Summary window opens. See [page 469](#).

- 1 Click **Finish**.
If the report is undeployed successfully, a confirmation message displays.
- 2 Click **OK**
- 3 To see that the report was undeployed, select the folder from which you undeployed the report, and click **Refresh** in the toolbar.

Deploying URLs

When you deploy a URL to the Deployed Items catalog, you store all the necessary information for the URL in the catalog. This is useful when you want to access a URL quickly from PI's Web Access Server client application.

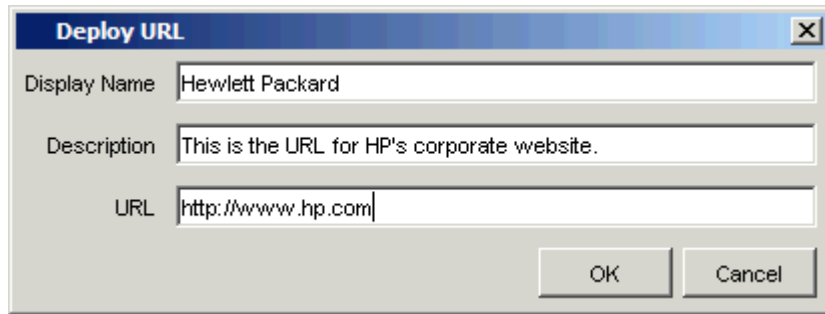
Once you deploy a URL, you can access it by clicking the **Catalog** link in the Web Access Server's links bar. For example, if you deployed a URL to the **System** folder, you would expand the System folder to display its deployed items. When you click the URL, its associated website displays in the Web Access Server's results window.

You do not have to deploy URLs to the URL catalog. This is optional.

To deploy a URL:

- 1 Right-click the folder in which you want to place the URL in the Deployed Items folder, and select **Deploy URL** from the menu.

The Deploy URL dialog box opens.



The screenshot shows a dialog box titled "Deploy URL". It has three text input fields. The first field is labeled "Display Name" and contains the text "Hewlett Packard". The second field is labeled "Description" and contains the text "This is the URL for HP's corporate website.". The third field is labeled "URL" and contains the text "http://www.hp.com". At the bottom right of the dialog are two buttons: "OK" and "Cancel".

- 2 In the Display Name box, type the name you want to use for the URL when the URL displays in the Deployed Items catalog.
- 3 In the Description box, type a description for the URL. A description is optional.

If you do not specify a description, you can do this later by right-clicking the URL and selecting **Properties** from the menu. The Properties dialog box opens and displays the properties of the URL (name, description, and URL address). See [Viewing and Changing URL Properties](#) on page 478 for more information.

- 4 In the URL box, type the address for the URL, supplying any parameters that you want to input to the website.

- 5 Click **OK**.

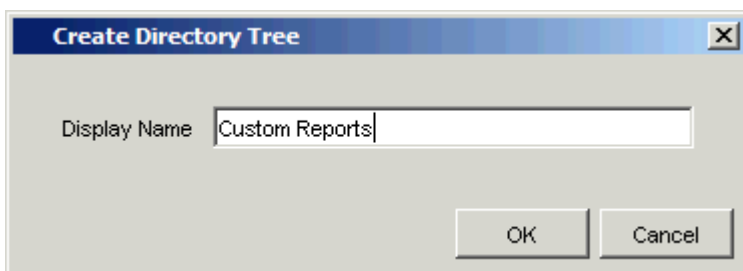
Creating New Directories in the Catalog

The Deployed Items catalog contains a folder called Deployed Items; this folder contains the `system` and `users` folders. You can create new directories within any directory in the `system` and `users` folders.

To create a new directory:

- 1 Right-click the folder in which you want to create a directory, and select **Create Directory** from the menu.

The Create Directory Tree dialog box opens.



- 2 In the Display Name box, type the name of the folder that you want to create, and click **OK**.

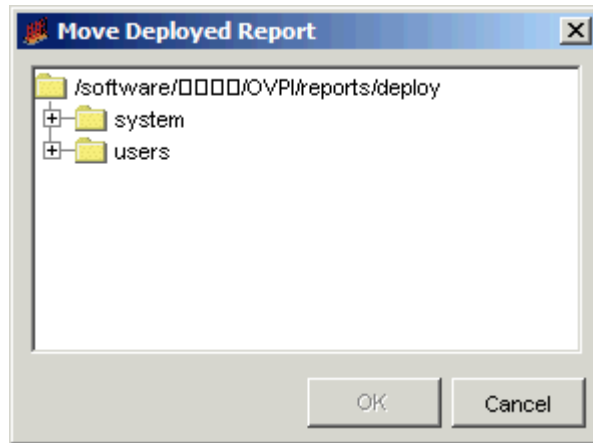
Moving Reports and Directories

You can move reports and directories.

To move a report:

- 1 In the Deployed Items catalog, locate and select the report that you want to move.
- 2 Right-click the report, and select **Move** from the menu.

The Move Deployed Report dialog box opens.

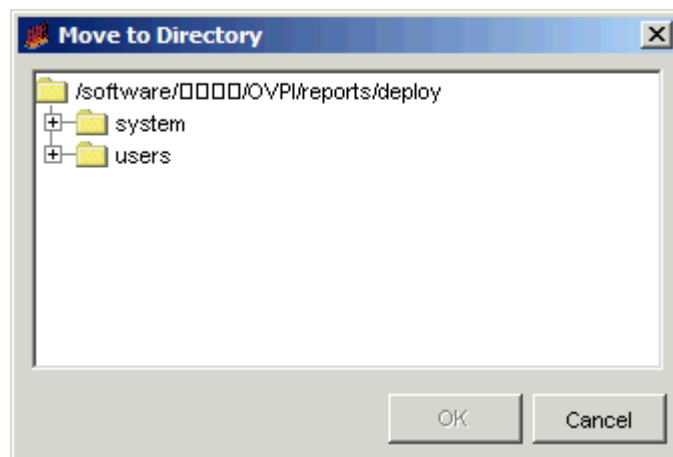


- 3 Select the folder to which you want to move the report, and click **OK**.

To move a directory:

- 1 In the Deployed Items catalog, locate and select the directory that you want to move.
- 2 Right-click the directory, and select **Move** from the menu.

The Move to Directory dialog box opens.



- 3 Select the directory to which you want to move to:

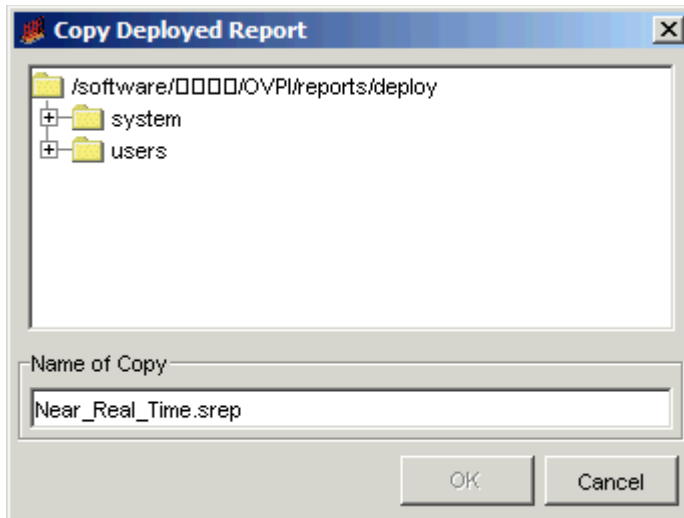
Copying Reports and Directories

You can copy reports and directories.

To copy a report:

- 1 In the Deployed Items catalog, locate and select the report that you want to copy.
- 2 Right-click the item, and select **Copy** from the menu.

The Copy Deployed Report or Copy to Directory dialog box opens.

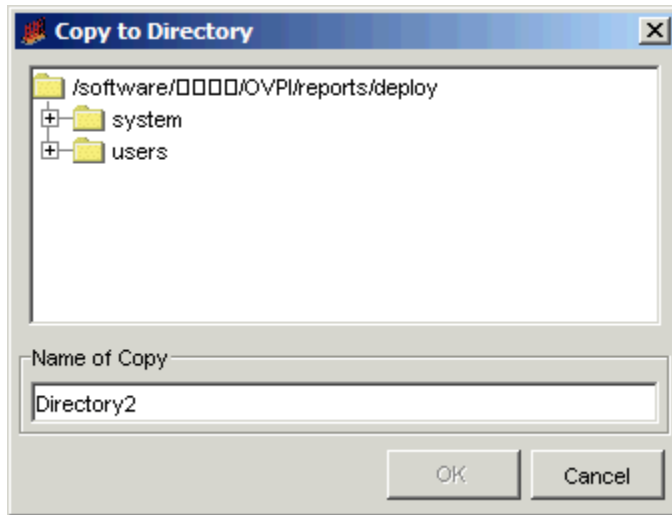


- 3 Do the following:
 - a Select the folder to which you want to copy the report.
 - b In the Name of Copy box, leave the name as is, or type a new name.
 - c Click **OK**.

To copy to a directory:

- 1 In the Deployed Items catalog, locate and select the directory to which you want to copy.
- 2 Right-click the item, and select **Copy** from the menu.

The Copy to Directory dialog box opens.



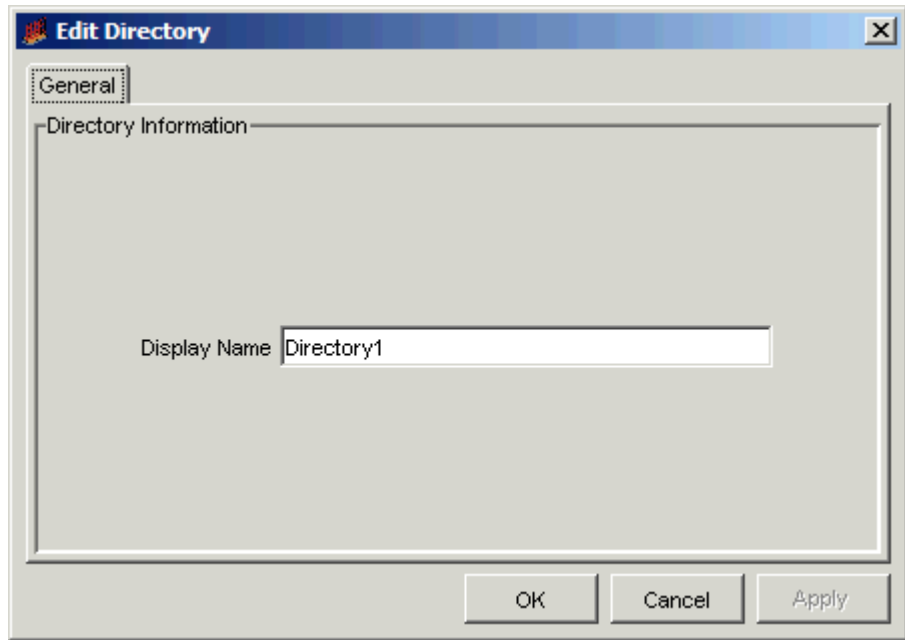
- 3 Do the following:
 - a Select the directory to which you want to copy the report.
 - b In the Name of Copy box, leave the name as is, or type a new name.
 - c Click **OK**.

Viewing and Changing Directory Properties

To view or change the properties of a directory:

- 1 Right-click the item, and select **Properties** from the menu.

The Edit Directory dialog box opens.



- 2 In the Display Name box, type a display name for the directory.
- 3 Click **OK**.

Viewing and Changing Deployed Item Properties

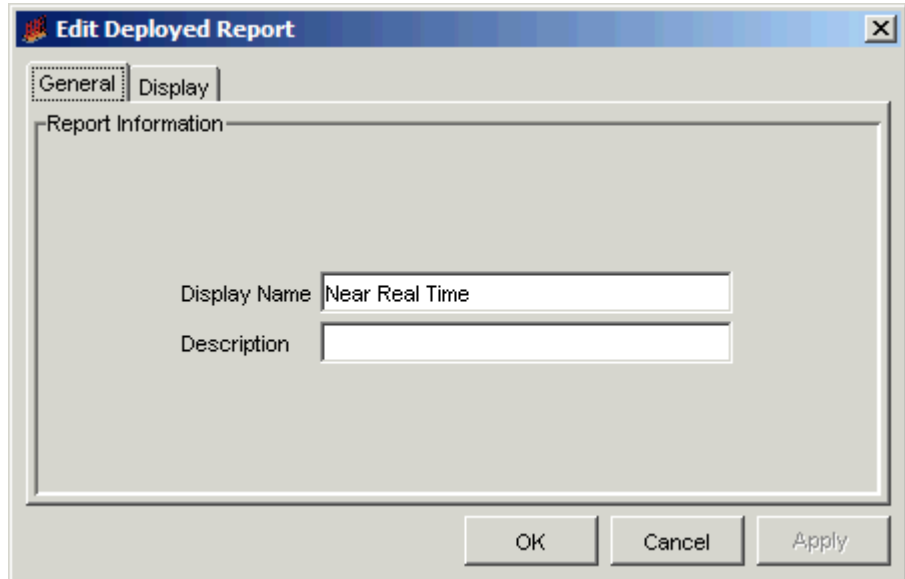
You can view and change the properties of a deployed report or URL.

Viewing and Changing Report Properties

To view or change report properties:

- 1 Right-click the item, and select **Properties** from the menu.

The Edit Deployed Report dialog box opens.



- 2 Click the General tab to change the report's name or description.

Do either of the following:

- In the Display Name box, type a display name for the report.
- In the Description box, type a description.

- 3 Click the Display tab to change how the report is displayed when it is viewed using the Web Access Server client application.

By default, when a user opens the report using the Web Access Server, the report is displayed in the current browser window. To make the report display in a new browser window, click **Open in New Browser**.

- 4 Click **OK**.

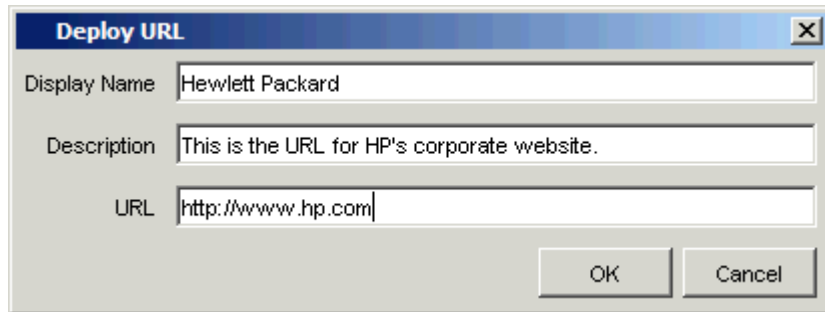
Alternatively, if you are only interested in viewing a report's name and description, hold the pointer over the report in the catalog.

Viewing and Changing URL Properties

You can view and change the properties of a URL, which are its display name (the name the appears for it in the Deployed Items catalog), description, and address.

To view and change a URL's properties:

- 1 Right-click the URL, and select **Properties** from the menu.



The URL Properties dialog box opens.

- 2 Do any of the following:
 - In the Display Name box, type a display name for the report.
 - In the Description box, type a description.
 - In the URL box, type a new URL and parameters.
- 3 Click **OK**.

Alternatively, if you are only interested in viewing a URL's name and description, hold the pointer over the URL in the catalog.

Deleting Reports and URLs

Deleting a report or a URL undeploys it from the Deployed Items catalog.

To delete a report or URL, right-click the item, and select **Delete** from the menu.

Refreshing the Catalog

Sometimes you may need to refresh a folder in the Deployed Items catalog. For example, if you and another person are using a shared administrator account, more than one person may be making changes at a given time. Thus, you may want to manually refresh the catalog to see the latest changes in the catalog.

To refresh a folder in the catalog, right-click the folder, and select **Refresh** from the menu.

Managing the Link Catalog

You manage the Link catalog by doing any of the following:

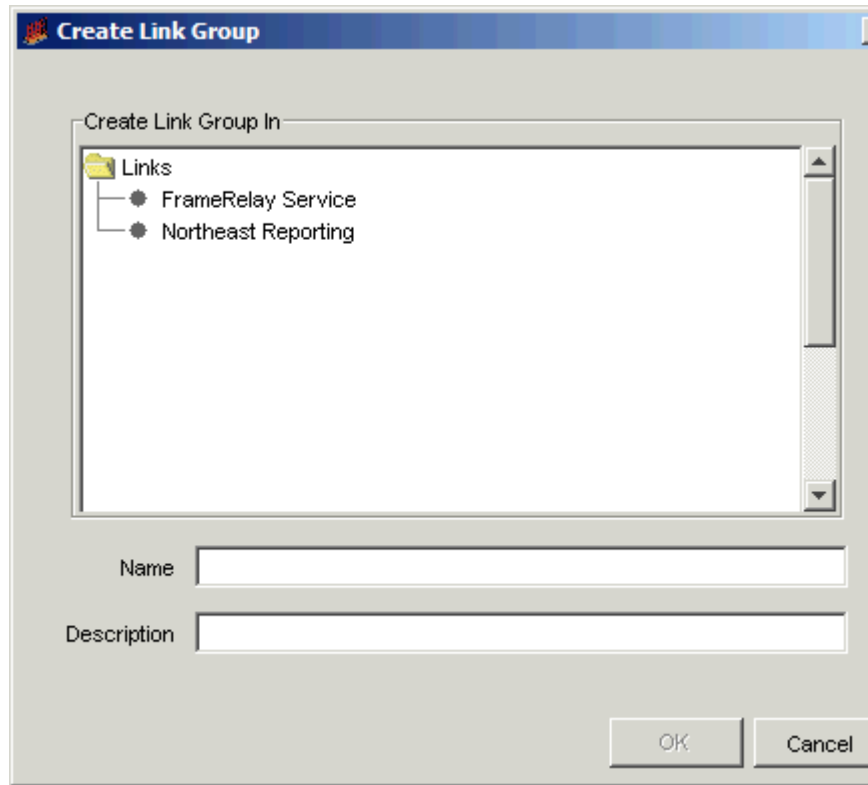
Creating Link Groups

By default, the Link Catalog has a folder called `Links`. You can create additional folders, called link groups, in the `Links` folder to organize your links by company, by group, by user, and so on.

To create a link group:

- 1 Do either of the following:
 - Right-click the `Links` folder in the Link Catalog, and select **New Link Group**.
 - From Catalog Manager's File menu, select **New**, and select **Link Group**.

The Create Link Group dialog box opens.



▶ If you selected Link Group from the File menu, this dialog box will also contain a Create Link Group In box.

- 2 If applicable, in the Create Link Group In box, select the folder in which you want to create the link group.
- 3 In the Name box, type a name for the link group folder. This folder will appear in the Links folder after it is created.
- 4 In the Description box, type a description of the link. This is optional.

To view these properties later, hold the pointer over the link group. Both the display name and the description will appear. To view these properties later, you can hold the pointer over the report in the catalog. The link's name and description are displayed.

To change the display name and description, right-click the link group folder, and select **Properties** from the menu. See [Viewing and Changing a Link's Properties](#) on page 508 for more information about changing a link's properties.

5 Click **OK**.

The link group displays in the `Links` folder.

Creating Report Links

A report link takes specific information from an element (graph, table, image, or label) in a source report, uses it as the input for a report parameter (in the case of a report to report link) or to an argument (in the case of a report to application or URL link), and displays the report, application, or URL with the specified information when a user opens the link in the report.

A report can be the destination for multiple report links. For example, you might view a report that displays customers and contains a report link that displays a device report, showing only the devices for a specific customer. You might view another report that displays locations. It contains a report link that displays the same device report, showing only the devices for a specific location. In this instance, both the customer and location reports link to the same destination report, which shows devices.

Likewise, one report can contain multiple links. One report can be linked to any number of reports, URLs, and applications.

Linking a Report to another Report

You can create a link in a deployed report (source report) to another deployed report (destination report). When a user clicks the link in the source report, the *link* opens the destination report using the value of a specific statistic associated with the selected element in the source report for the selected report parameter in the destination report. For example, you might want to create a report to report link to display more information about a device using an additional report.

Suppose you have both the Device Resources and Interface Reporting report packs installed. When these two report packs run together, Device Resources will indicate when a router CPU is spiking, and Interface Reporting will tell you whether interface utilization levels could be contributing to the problem.

Thus, you view the Device Resource Near Real Time Quickview report, which lists your devices' average utilization (buffer, CPU, and memory) for the last six hours, and you see that a particular device is overutilized. In this case, you would want to open the Interface Reporting Near Real Time report, which lists the individual interfaces up to the last poll cycle, to see if a particular interface is causing the problem. Additionally, you might always want to open this report regardless of whether any devices are experiencing problems in the future. Instead of having to open the report repeatedly, you can create a link from the Device Resource Near Real Time Quickview report to this report by creating a report to report link.

In this case, using the Report to Report Link wizard, you would want to link the Device statistic from the table that lists the devices in the Near Real Time Quickview report to the Device parameter in the Near Real Time report. The following example shows how the links would appear (underlined) in the report when viewing using the Web Access Server.

When you clicked on any link in the Device Resources report, the report would supply the information for the selected device to a selected parameter in the Near Real Time report. For example, the result of clicking on a link in the Device Resource Near Real Time Quickview would display the Interface Reporting Near Real Time report, which would show all the interfaces for the device you selected as follows:

Interface Reporting
Near Real Time

The Near Real Time report provides a detailed look at the performance of individual interfaces up to the last poll cycle. Select an interface from the list of active interfaces in the previous hour and the tabbed areas show exception counts, discards and errors as a percentage of utilization, and discards and errors as a percentage of available bandwidth. F/H indicates full or half duplex.

Interface Selection List
Active Interfaces with Statistics for the Previous Moving 6 Hours

Device	Interface	F/H	Speed	Avg Util.	Peak Util	% Errors	% Discards	Admin Status	Description
15.6.96.2	Tunnel5	F	In: 9 Kb/s Out: 9 Kb/s	In:2.6 Out:1.6	In:3.1 Out:2.3	In:0.0 Out:0.0	In:0.0 Out:0.0	Up	Tunnel5
15.6.96.2	Tunnel4	F	In: 9 Kb/s Out: 9 Kb/s	In:2.6 Out:1.3	In:3.1 Out:1.7	In:0.0 Out:0.0	In:0.0 Out:0.0	Up	Tunnel4
15.6.96.2	Tunnel7	F	In: 9 Kb/s Out: 9 Kb/s	In:1.3 Out:1.5	In:1.8 Out:2.2	In:0.0 Out:0.0	In:0.0 Out:0.0	Up	Tunnel7
15.6.96.2	Serial0	F	In: 1.5 Mb/s Out: 1.5 Mb/s	In:1.0 Out:0.6	In:2.6 Out:1.0	In:0.0 Out:0.0	In:0.0 Out:0.0	Up	Serial0
15.6.96.2	Tunnel6	F	In: 9 Kb/s Out: 9 Kb/s	In:1.3 Out:1.2	In:1.8 Out:1.6	In:0.0 Out:0.0	In:0.0 Out:0.0	Up	Tunnel6
15.6.96.2	Serial0.2	F	In: 1.5 Mb/s Out: 1.5 Mb/s	In:0.8 Out:0.5	In:2.3 Out:0.8	In: 0 Out: 0	In: 0 Out: 0	Up	Serial0.2
15.6.96.2	Serial0.1	F	In: 1.5 Mb/s Out: 1.5 Mb/s	In:0.3 Out:0.1	In:1.5 Out:0.4	In: 0 Out: 0	In: 0 Out: 0	Up	Serial0.1

Details Protocol Group Location Country
No Data



You can not add the reports saved in the Users folder to any custom view. The report to report link for the reports in the Users folder, therefore, does not work if you select a custom view.

The links will only work if you select the “default” view under Web Application Server → Administration → User Accounts → Users → User X → View.

In this instance, X is the name of the user.

To create a link from one report to another report:

Task 1: Start the Report to Report Link wizard.

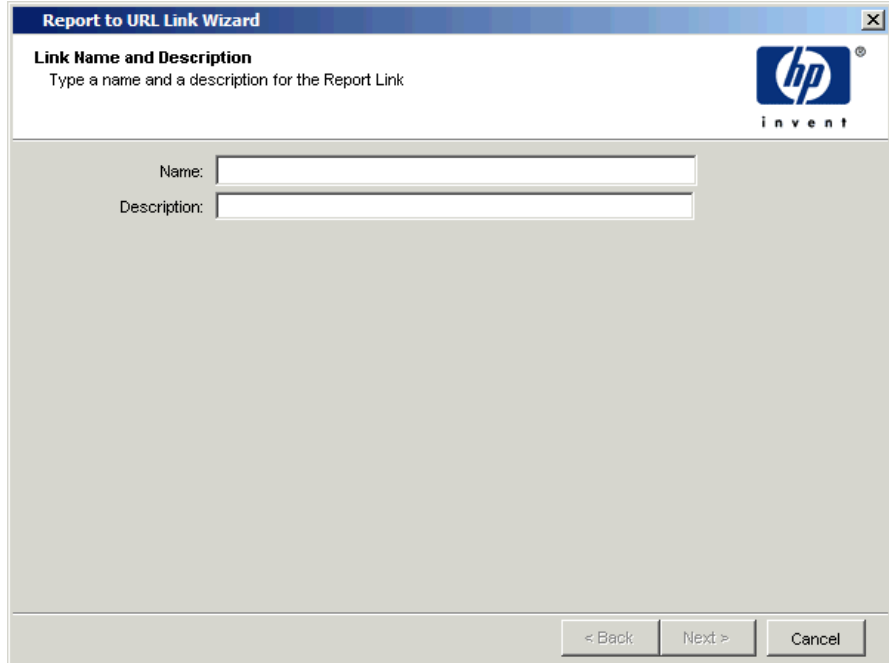
You can start the Report to Report Link wizard in the following ways:

- Right-click a link group in the Link catalog, select **New Report Link**, and select **New Report to Report Link** from the menu.

By default, Links is the only link group; however, you can create other link groups. For more information, see [Creating Link Groups](#) on page 480.

- From the Management Console’s File menu, select **New**, and select **Report to Report Link**.

The Link Name and Description window opens.



The screenshot shows a dialog box titled "Report to URL Link Wizard" with a close button in the top right corner. The main heading is "Link Name and Description" with the instruction "Type a name and a description for the Report Link". The HP logo and the word "invent" are in the top right. Below the heading are two text input fields: "Name:" and "Description:". At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

Task 2: Specify a name and a description for the link.

- 1 In the Name box, type the name you want to use for the link when the link displays in the Links catalog.
- 2 In the Description box, type a description for this link. A description is optional.
- 3 Click **Next**.

To change this information later, right-click the link in the Link catalog and select **Properties** from the menu. See [Viewing and Changing a Link's Properties](#) on page 508 for more information.

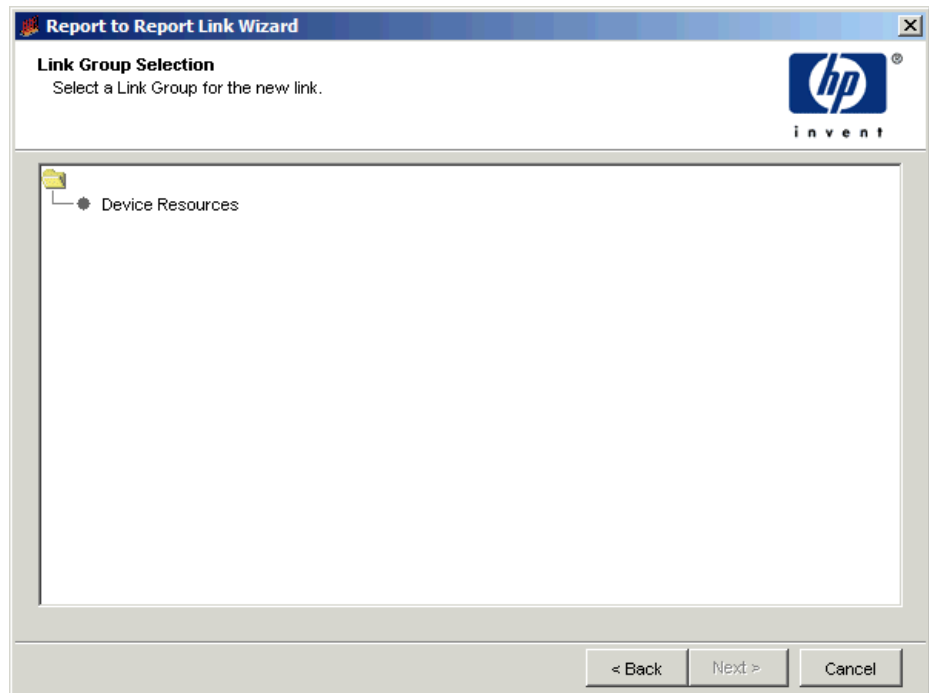
Task 3: Select the link group in which Catalog Manager should create the link.

The following figure shows the Select Link Group window.



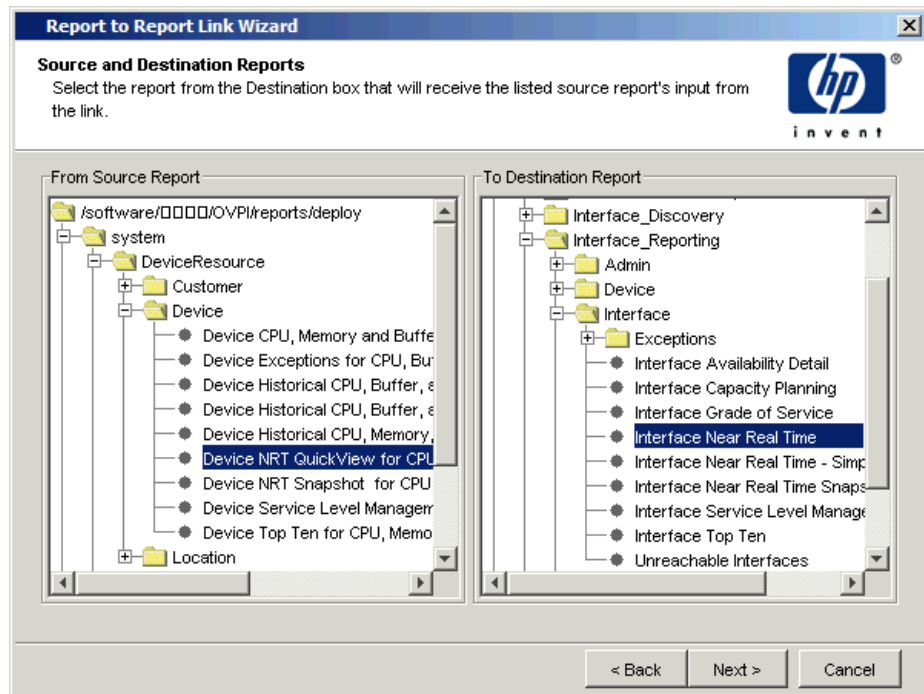
You will only see this window if you started the Report to Report Link wizard from Management Console's File menu.

If you do not see this window, go to [Task 4](#) on page 487.



- 1 Select the folder in which you want to create the link.
- 2 Click **Next**.

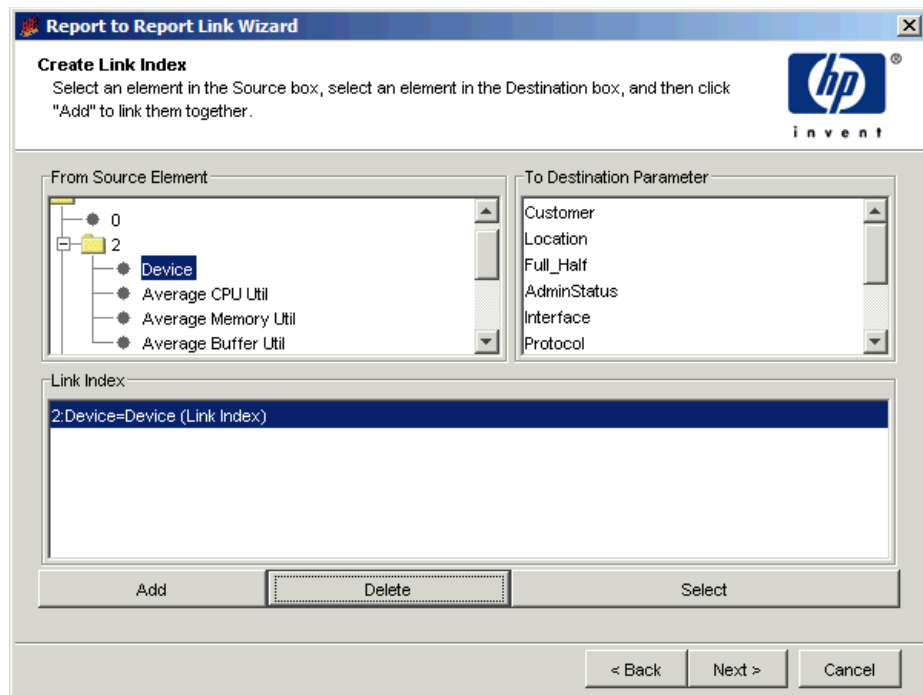
Task 4: Specify the source report and the destination report.



You will only see this window if you started the Report to Report Link wizard from Management Console's File menu.

- 1 In the From Source Report box, select the report that will contain the report link.
- 2 In the To Destination Report box, select the report that should open from the link in the source report.
- 3 Click **Next**.

Task 5: Specify the link index.



The *link index* is the place in the source report from which the link is triggered and, if applicable, information is passed to the destination report.

A link index in a report to report link can function in either of the following ways:

- An element (image or label) in the source report triggers a link to the destination report. The destination report opens but no information is passed.

For example, if a company logo is selected in the source report, the destination report opens when a user clicks the logo.

- A statistic in the source report triggers a link to a selected parameter in the destination report, passing the statistic's information and affecting the information displayed in the destination report.

For example, suppose the link index is made from the Device statistic from the Near Real Time Quickview report to the Device parameter in the Near Real Time report. When a device is selected in the Near Real time Quickview report, the link opens the Near Real Time report, which displays all the interfaces for the selected device in the Near Real Time report.

- 1 In the From Source Element box, select an element (image or label) or a statistic.



If you selected an image or a label, skip [step 2](#), and click **Add**. The Link Index box of the window will look similar to the following if a logo is selected: HP Logo (Link Index)

- 2 In the To Destination Parameter box, select the parameter in the destination report that will receive information from the selected statistic in the source report.



If the destination report does not contain any parameters, the To Destination Element box will be empty, **Next** is greyed out, and you will be unable to continue. At this point, you can either click **Cancel** or select an image or a label from the From Source Element box.

To make sure a destination report has parameter before you try to create a link to it, open the report using Report Builder and expand the `Parameters` folder in the Components tree. See the *HP Performance Insight Guide to Building and Viewing Reports* for more information about Report Builder and report parameters.

- 3 Click **Add**.

The link appears in the Element Links box. For example, a link may appear like the following:

```
2:Device=Device (Link Index)
```

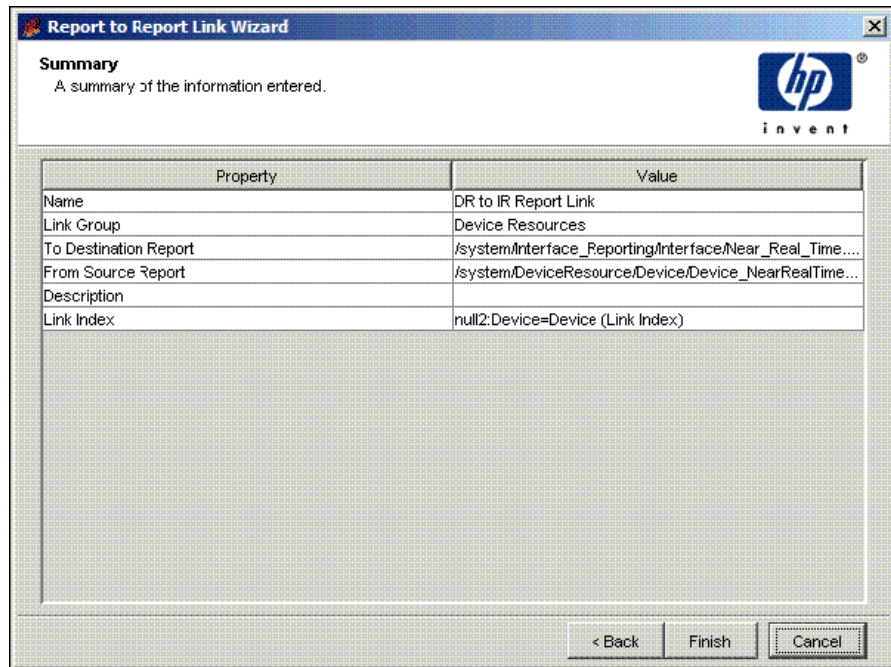
This shows that the device statistic in table number two in the source report is linked to the device parameter in the source report. Thus, when the link is triggered, the device statistic from the table is passed to the device parameter in the destination report.

You can create multiple link indexes. That is, you can pass information from numerous elements in the source report to the destination report; however, only one element can function as the link index. The link index indicates the element in the source report that will trigger the link. For

example, suppose you are passing a customer name and customer ID to the destination report; if customer name is the link index, the user will click it to trigger the link.

- 4 If necessary, you can do either of the following:
 - a To remove a link index, select it in the Link Index box, and click **Delete**.
 - b To change the key link index when there are link indexes, select the link index in the Link Index box, and click **Select**.
- 5 Click **Next**.

The Summary window opens.



Task 6: Review the properties of the link.

Click **Finish** to create the link.

A graphical representation of the link is displayed in Catalog Manager's viewing pane. The link will also appear in the folder that you specified in the Link catalog.

Linking a Report to an Application

You can create a link in a deployed report (source report) to an application. When a user clicks the link in the source report, the link opens the application using any parameters you specified.

Using the deployed reports in the Deployed Items catalog, you can create a link in a report (the source report), which will take selected information from the source report and supply it to an application that you have specified. Specifically, you create a link index, where you select an element or a statistic from the source report, that becomes the input to the specified application.

For example, you might want to use a report to application link when you have a report that lists all of your devices, and when you select a certain device in the report, you would like to be able to run a telnet utility to log on to that device.

To create a report to application link:

Task 1: Start the Report to Application Link wizard.

1 Do one of the following:

- Right-click a link group in the Link catalog, select **New Report Link**, and select **New Report to Application Link** from the menu. By default, `Links` is the only link group; however, you can create other link groups. For more information, see [Creating Link Groups](#) on page 480.
- From Management Console's File menu, select **New**, and select **Report to Application Link**.

The Link Name and Description window opens as shown on [page 485](#).

Task 2: Specify a name and a description for the link.

- 1 In the Name box, type the name you want to use for the link when the link displays in `Links` folder in the Links catalog.
- 2 In the Description box, type a description for this link. A description is optional.
- 3 Click **Next**.

To change this information later, right-click the link and select **Properties** from the menu. See [Viewing and Changing a Link's Properties](#) on page 508 for more information.

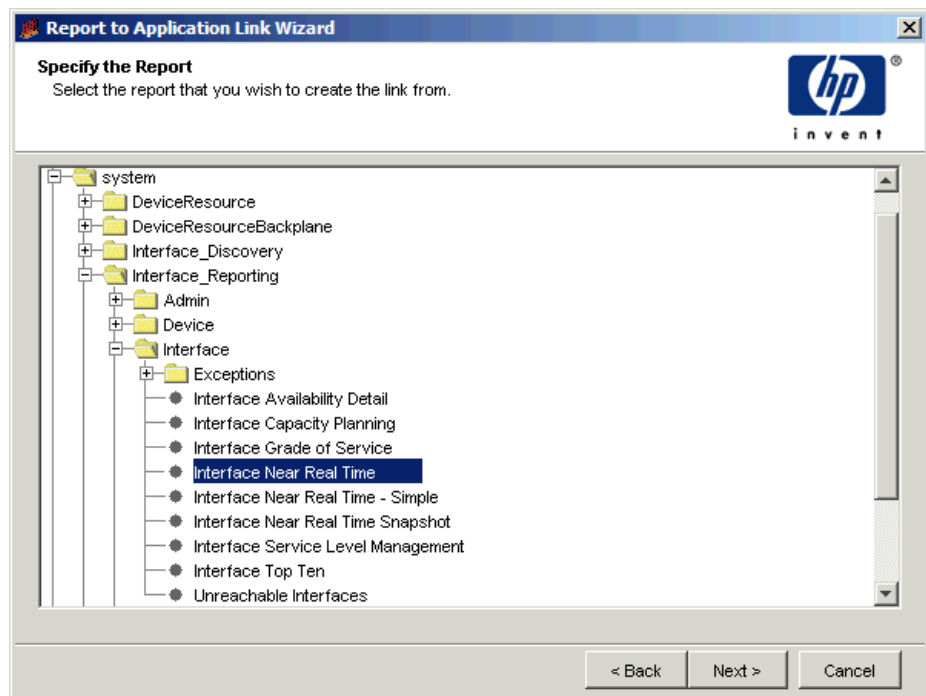
Task 3: Select the Link Group in which the link should be created.

The Select Link Group window opens, which is shown on [page 485](#).

▶ You will only see this window if you started the Report to Application Link wizard from Management Console's File menu. (If you do not see this window, go to [Task 4](#) on [page 492](#).)

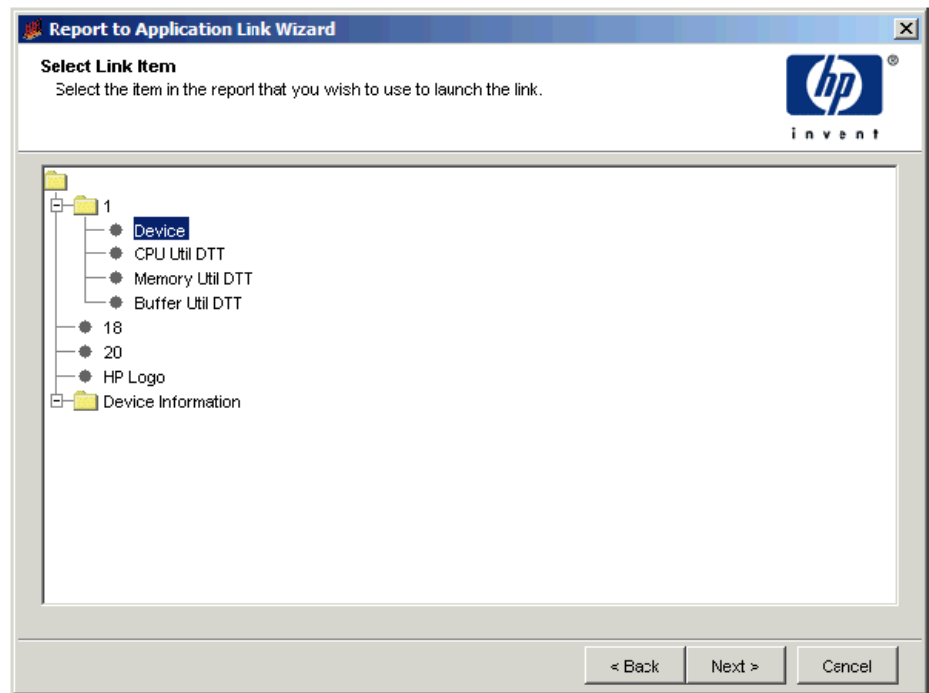
- 1 Select a link group.
- 2 Click **Next**.

Task 4: Specify the source report.



- 1 Select the source report.
- 2 Click **Next**.

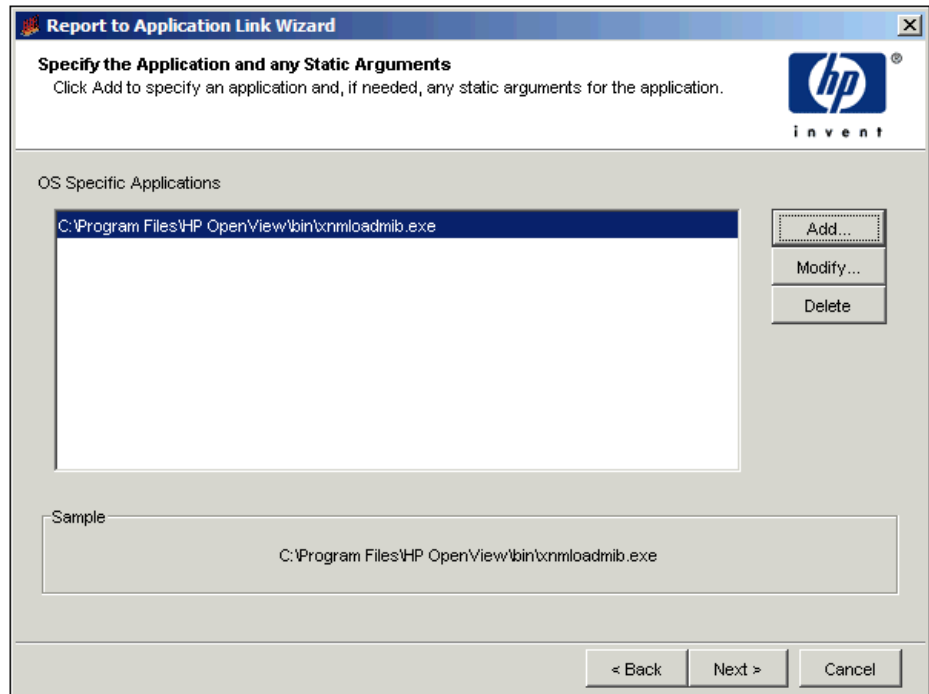
Task 5: Select the link item.



Select the element or statistic in the source report that will trigger the link to the destination application.

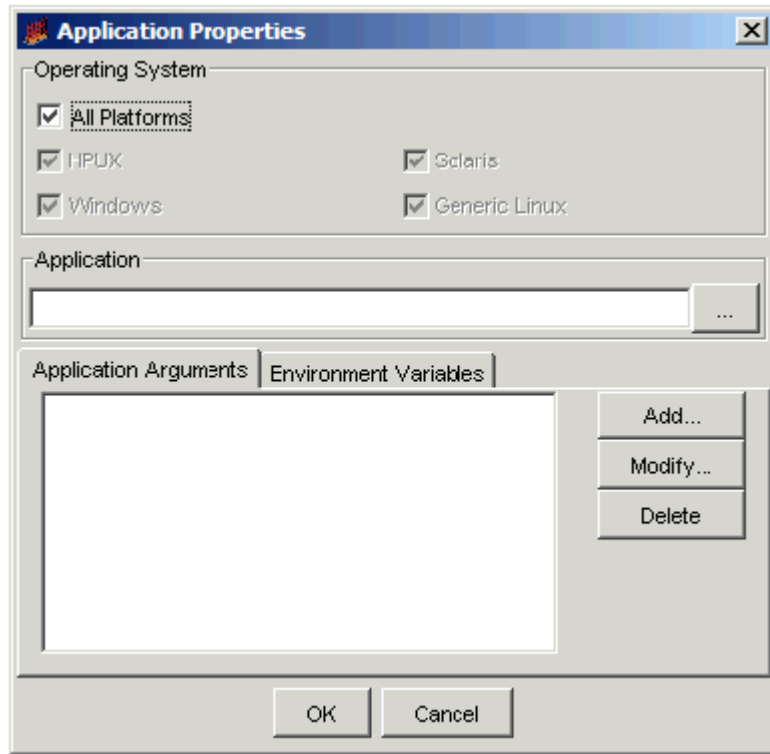
- 1 Select an element or a statistic in the source report.
- 2 Click **Next**.

Task 6: Specify the application and any arguments.



- 1 Click **Add** to specify an application and, if needed, supply any static arguments for the application.


The Application Properties dialog box opens.



- 2 In the Operating System box, select the operating system on which the application runs.
 - 3 In the Application box, do either of the following:
 - Type the location of the application's executable.
 - Click **Browse** to open the Open dialog box. Navigate to the application's executable, and click **Open**.
- ▶ The application specified must be installed on the system where the client application (Report Viewer or Report Builder) that is used to view the report is run.
- 4 To specify any arguments or environment variables you want used when the application is invoked, click the Application Arguments tab or the Environment Variables tab, respectively, and click **Add**.

This is optional. For example, you might supply the arguments necessary to log on to the application if this is required.

The Command Line Option Editor dialog box opens.



- a In the Argument box, type an argument.
- b To specify a value, do either of the following:
 - In the Value box, type a value for the statistic
 - Click **Statistic** to open the Select an Element or Statistic dialog box, which displays the elements and statistics in the source report. Select a statistic and click **OK**.

To modify an argument, select the argument, and click **Modify**. Make any changes, and click **OK**.

To delete an argument, select the argument, and click **Delete**. Make any changes, and click **OK**.

- 5 Click **Next**.

Task 7: Review the properties of the link.

Click **Finish** to create the link.

A graphical representation of the link is displayed in Catalog Manager's viewing pane. The link also appears in the Link Catalog.

The figure on [page 492](#) shows how a link might display in the Link Catalog and in the viewing pane of Catalog Manager.

Linking a Report to a URL

You can create a link in a deployed report (source report) to a URL. When a user clicks the link in the source report, the link opens the URL using the parameters, statistics or both that you specified.

For example, suppose you are managing HP procure switches with the procure datapipe. You might want to create a link from a node in a report to the URL that will log the user onto the switch.

To create a link from one report to a URL:

Task 1: Start the Report to URL Link wizard.

Start the Report to URL link wizard using either of the following methods:

- Right-click a link group in the Link catalog, select **New Report Link**, and select **New Report To URL link** from the menu.
- From the File menu, select **New**, and select **Report to URL Link**.

The Link Name and Description window opens as shown on [page 485](#).

Task 2: Specify a name and a description for the link.

- 1 In the Name box, type the name you want to use for the link when the link displays in `Links` folder in the Links catalog.
- 2 In the Description box, type a description for this link. A description is optional.
- 3 Click **Next**.

To change this information later, right-click the link and select **Properties** from the menu. See [Changing Link Properties](#) for more information.

Task 3: Select the Link Group in which the link should be created.

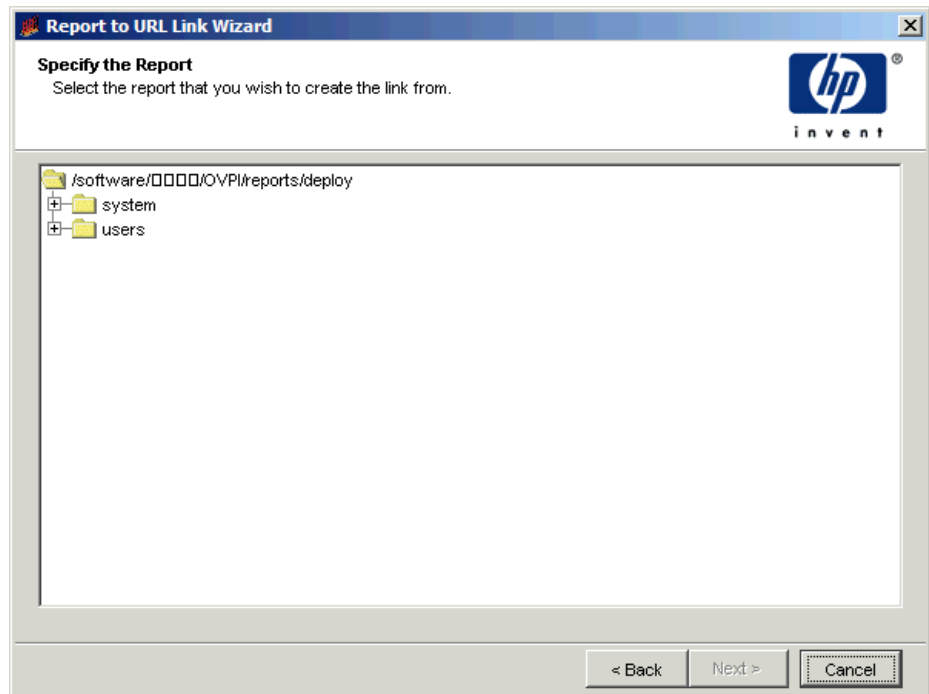
The Link Group Selection window opens, which is shown on [page 485](#).



You will only see this window if you started the Report to URL Link wizard from Catalog Manager's File menu. (If you do not see this window, go to [Task 4](#).)

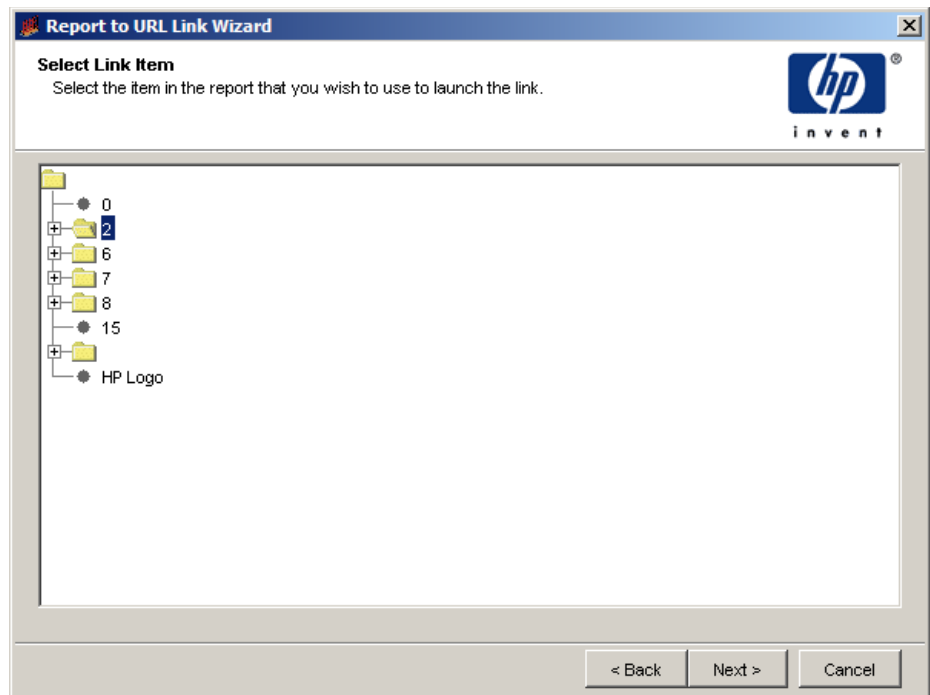
- 1 Select a link group.
- 2 Click **Next**.

Task 4: Specify the source report.



- 1 Locate and select the report you want to create the link in.
- 2 Click **Next**.

Task 5: Specify the item in the source report that will trigger the link.



You need to select the element or statistic in the source report that will trigger the link. This is called the *link anchor*. For example, if you select an element like a logo, the user will click the logo in the report to display the destination URL.

- 1 Select an element or a statistic in the source report.
- 2 Click **Next**.

Task 6: Type the URL information.

Report to URL Link Wizard

Add the URL
Enter the information necessary to build the URL.

Build this URL as a relative URL:

Protocol:

Host: Use Statistic

(Optional) Port: Use Statistic

(Optional) Path:

< Back Next > Cancel

- 1 Select the **Build this URL as a relative URL** check box when you want to link to a page on the current server.
After you select this option, the Protocol, Host, and Port boxes in the window become unavailable.
- 2 Select the protocol you want to use (http or https) from the Protocol box.
- 3 Do one of the following:
 - Type the host name of the destination web server in the Host box.
 - Click **Use Statistic** if you want to select a statistic in the source report that will provide the host information.
The Select an Element or Statistic dialog box opens. Select a statistic, and click **OK**.
- 4 Do one of the following:

- Type the port in the Port box.
- Click **Use Statistic** if you want to select a statistic in the source report that will provide the port information.

The Select an Element or Statistic dialog box opens. Select a statistic, and click **OK**.

- ▶ If you use a statistic for the host or port, the Host and Port boxes will be empty; however, the information will appear at the bottom of the Add the URL window.

- 5 Type a path in the Path box. This is optional.

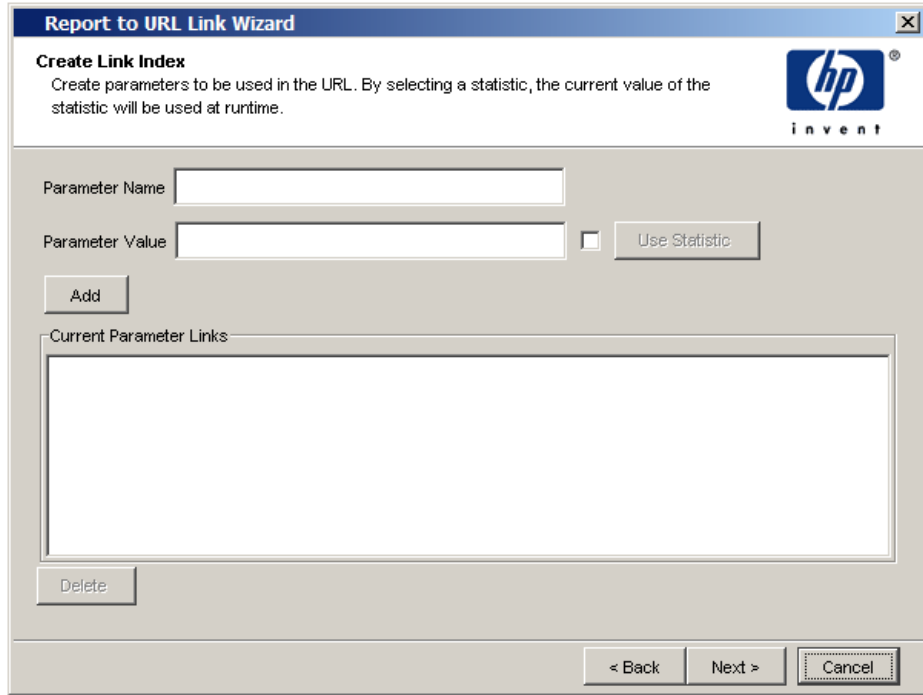
For example, suppose you are creating a link from a report to a report on another Web Access Server system. You would specify a path similar to the following:

/reports/webview

- 6 Click **Next**

Task 7: [Create the link index.](#)

The following figure shows the Create Link Index window.



You can specify parameters that should be used when the URL is invoked. This is a *link index*. This is the way information is passed from the source report to a parameter. A link index is optional. Click **Next** to continue if you do not want to create one. See *Creating a Link Index* for an example link index scenario.

- 1 Type the name of the parameter you want used when the URL is invoked in the Parameter Name box.
- 2 Do one of the following:
 - In the Parameter Value box, type the parameter value.
 - Select the check box next to **Use Statistic**, and click **Select Statistic** to select a statistic for the parameter value.
- 3 Click **Add** to create the link index. The link index appears in the Current Parameter Links box



You can create multiple link indexes. That is, you can pass information from numerous elements in the source report to the destination URL.

- 4 If necessary, to remove a link key, select it in the Current Parameter Links box, and click **Delete**.
- 5 Click **Next**. The Summary window opens.

Creating a Link Index

For example, suppose you want the URL to open a report on another Web Access Server system. In [Task 6](#) on page 500, you specified the host and port. In this step, you need to supply the parameters that specify where the report is located. If you wanted to open the Device Near Real Time Snapshot report and select a particular device that should be shown in the report, you would do the following:

- 1 In the Enter a Parameter box, type **rn** as the parameter.
- 2 In the Select a Parameter Value box, type the following:
**system/DeviceResource/Device/
Device_NearRealTime_Snapshot.rep**
- 3 Click **Add**.
- 4 In the Enter a Parameter box, type **p_Device**.
- 5 Select the check box parameter next to **Select Statistic**, and click **Select Statistic**.
- 6 Select the statistic for the device (for example, Target Name) and click **OK**.
- 7 Click **Add**.

At this point, the URL will appear at the bottom of the Create Link Index window and look like the following:

```
http://ackcat:80/reports/webview?rn=/system/DeviceResource/Device/  
Device_NearRealTime_Snapshot.rep&p_Device=<Target Name&
```

Task 8: [Review the properties of the link.](#)

Click **Finish** to create the link.

A graphical representation of the link is displayed in Catalog Manager's viewing pane. The link also appears in the Link Catalog. If you started the link wizard from the File menu, the link appears in the folder you specified in the Select Link Group window of the wizard. If you started the wizard from a

folder in the Link Catalog, the link appears in that folder. The figure on [page 492](#) shows how a link might display in the Link Catalog and in the viewing pane of Catalog Manager.

Opening Report Links

Depending on the application a user views the source report with and the type of element used to create the link, the link will appear as follows:

- Web Access Server

The link is treated just like a URL. It will display using the same method that the browser uses to display URLs. For example, if the browser underlines URLs, any links in the report will be underlined.



Users cannot open report to application links using the Web Access Server. Instead, they must use Report Builder or Report Viewer to open these kinds of links. Additionally, the application must reside on the system from which Report Builder or Report Viewer are run to successfully invoke the application linked to the report.

- Report Builder and Report Viewer

The cursor will change from an arrow to a hand when it is placed over the link. Additionally, depending on the type of element used for the link, the following occurs:

- Elements that are the destination of a report link are surrounded by a blue boundary.
- Text in a label or in a table column is underlined.

This behavior is illustrated in the following example:

Blue border—

Hand cursor—


Underlined text

Device Summary			
Mon, May 12 12:00 AM			
Device	Busy Hour CPU Util	Busy Hour Buffer Util	Busy Hour Memory Util
<u>Cisco_02</u>	85.00	98.23	94.67
<u>Cisco_03</u>	94.75	97.97	87.56
<u>Cisco_04</u>	82.50	97.95	82.70
<u>Cisco_01</u>	84.50	97.41	86.97

To open a link, click (Web Access Server) or double-click (Report Builder and Report Viewer) the link in the report.


Modifying Links in the Viewing Pane

When you select a link from the Link catalog, a graphical representation of the link appears in Catalog Manager's viewing pane.



You can modify the link by doing any of the following:

- Change the way the links are displayed.

Right-click in the navigation pane or click  in the toolbar. From the menu, select one of the following options:

- Circular. Places the links in a circular pattern.
- Hierarchic. Places the links in a hierarchal pattern.
- Organic. Places the links in an organic pattern.
- Orthogonal. Places the links at 90 degrees (right angles).
- Random. Places the links in a random pattern.
- Tree. Places the links in a tree pattern.

- Move linked objects (reports, URLs, and applications).

Click the object; handles will appear. Drag the object to its new location.

- Modify the link.

Select the arrow that connects the linked objects. It is highlighted yellow. Right-click, and a menu containing the following options displays:

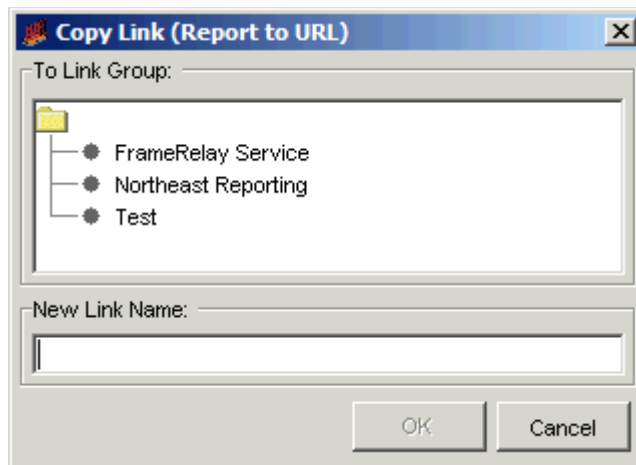
- **Edit.** Select this option to view and change the element's link key. See [Viewing and Changing a Link's Properties](#) on page 508 for more information. To change the link's display name, description, or both, right-click the link in the catalog, and select **Properties** from the menu.
- **Enable.** Select this option to disable the link in the report. The link remains in the catalog but is not active. You can enable it at any time.
- **Open In New Browser.** Select this option to make the destination of a link (report or URL) open in a new browser window when the report is viewed with the Web Access Server.
- **Delete.** Select this option to remove the link from the catalog.

Copying Links

To copy a link:

- 1 Locate the link in the catalog that you want to copy.
- 2 Right-click the link, and select **Copy** from the menu.

The Copy Link dialog box opens, displaying the link group folders to which you can copy the link.



- 3 Do the following:
 - a In the To Link Group box, select the folder to which you want to copy the link.
 - b In the New Link Name box, type in a new name for the link, if needed.
 - c Click **OK**.


Moving Links

You can move a link from one link group to another.

To move a link:

- 1 Locate the link in the catalog that you want to move.
- 2 Right-click the link, and select **Move** from the menu.

The Move Link dialog box opens, displaying the link group folders to which you can move the link.



- 3 Select the folder to which you want to move the link, and click **OK**.

Deleting Links

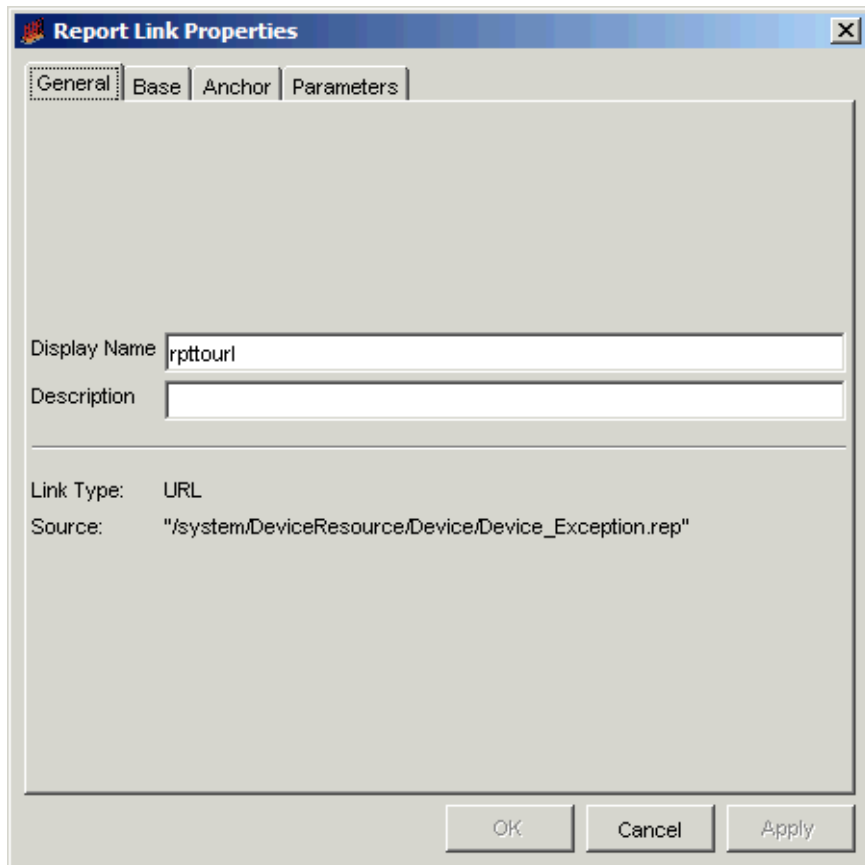
To delete a link from the link catalog, right-click the link, and select **Delete** from the menu. A confirmation dialog box opens, asking if you are sure you want to delete the link. Click **Yes** to delete the link.

Viewing and Changing a Link's Properties

To view or change a link's properties:

- 1 Right-click the link in the navigation pane, and select **Properties** from the menu.

The Report Link Properties dialog box opens.



- 2 Click the General tab to view or change the following:
 - Display Name. The name of the link in the catalog.
 - Description. The description of the link, if you have supplied one. A description is optional.
 - Link Type (Report, URL, or Application)
 - Source.
- 3 *Report to URL Link only:* Click the Anchor tab to change the element or statistic (the link anchor) in the source report that triggers the link. This tab is identical to the Select Link Item window in the Report to URL wizard. See [Task 5](#) on page 499 for details.
- 4 *Report to URL Link only:* Click the Base tab to change any of the URL information. This tab is identical to the Add the URL window in the Report to URL wizard. See [Task 6](#) on page 500 for details.
- 5 Click the Parameters tab. Depending on the type of report link, see the following for information about how to change the link parameters:
 - Report. This tab is identical to the Create Link Index window in the Report to Report wizard. See [Task 5](#) on page 488 for details.
 - Application. This tab is identical to the Create Link Index window in the Report to Application wizard. See [Task 7](#) on page 501 for details.
 - URL. This tab is identical to the Create Link Index window in the Report to URL wizard. See [Task 7](#) on page 501 for details.
- 6 Click **Apply** to make your changes take effect and click **OK** to close the dialog box.

Exporting and Importing Links

You can export links in a link group or import them. This is useful in any of the following situations:

- You want to include report links in a report pack you create.

Report links are included using the `reportlink` directive in the `install.pkg` file. See the *HP Performance Insight Reference Guide* for more information about the `install.pkg` file and its directives.
- You want to save your report links for backup purposes on another system.

- You have multiple Web Access Server systems, and you want to export the links from one system and import them to your other systems.

Exporting a Link Group

To export a link group:

- 1 Right-click a link group in the Link catalog.
- 2 From the File menu, select **Export**.
The Save dialog box opens.
- 3 Do the following:
 - a Navigate to the directory in which you want to save the link file.
 - b In the File name box, type a filename for the links (for example, `deviceresources_links.tlg`).
 - c Click **Save**.

Importing a Link

To import a link group:

- 1 Right-click the link group in the Link catalog to which the links should be imported.
When the link group is imported, it will appear as a subfolder of this link group.
- 2 From the File menu, select **Import**.
The Open dialog box opens.
- 3 Do the following:
 - a Navigate to the directory where the .tlg file is located.
 - b Select the file, and click **Open**.

Refreshing the Links Catalog

To refresh the a folder in the links catalog, right-click the folder, and select **Refresh** from the menu.

15 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, PI system attributes, mail server, and Secure Socket Layer (SSL) communications.
- View 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 OVPI HTTP Server is running. If not, see [Starting and Stopping HP Performance Insight Processes](#) on page 103.
- You have one of the following Web browsers:
 - Internet Explorer 5.5 or later
 - Mozilla Firefox 1.5 or later

To access the Web Access Server:

- 1 Start your Web browser.
- 2 Type the URL in the following format in the Address or Location field of your browser, and 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.



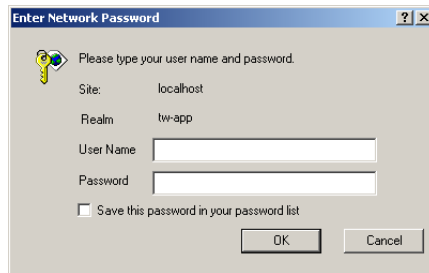
You do not need to type a port number if your Web Access Server is using port 80.

The Log On page opens.



- 3 Click **Log on**.

The Enter Network Password dialog box opens.



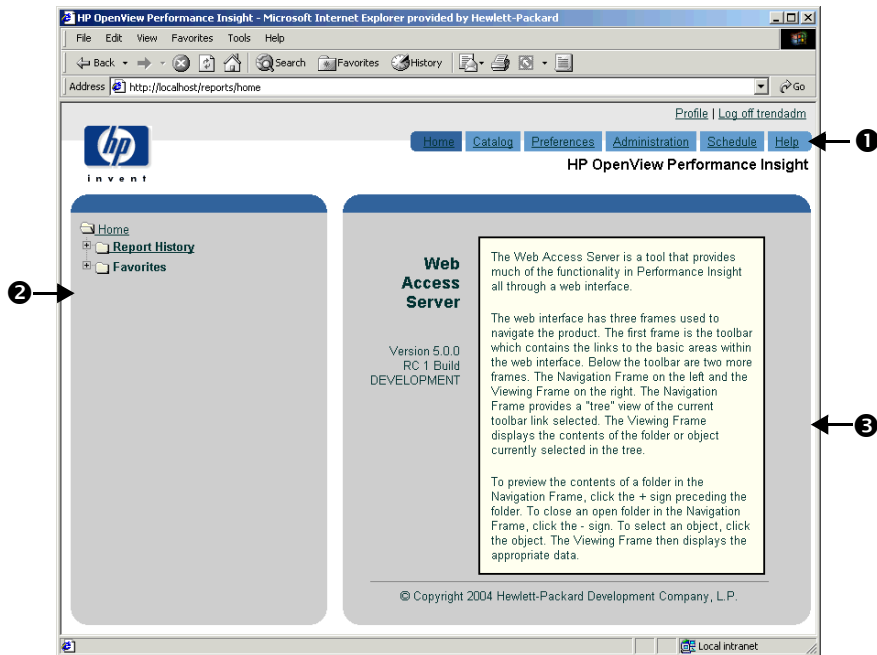
- 4 Type your Administrator username and password, and click **OK**.

PI 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 deployed item settings, and control the appearance of the Web Access Server.


For information about the links that are available to non-Administrator users, see the *HP 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.
- **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 deployed item settings.

The navigation tree for the Administration link appears as follows:

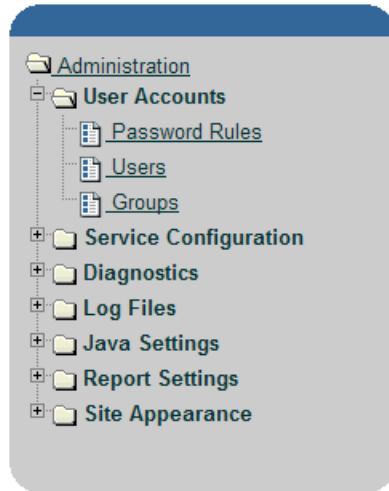


It contains the following folders: User Accounts, Service Configuration, Log Files, Diagnostics, Java Settings, Deployed Items, and Site Appearance.

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 PI's client applications (Management Console, Web Access Server, Report Builder, and Report Viewer).

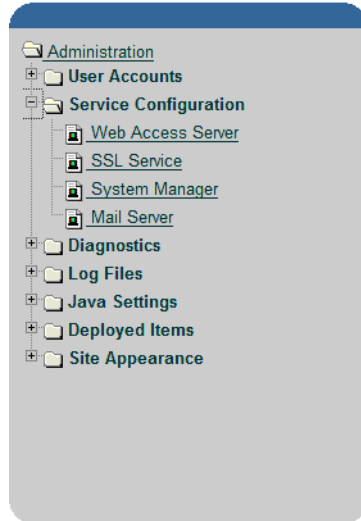
The User Accounts folder contains the following:

- 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 525 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 525 for details.
- Groups folder
Use this folder to create, modify, or delete groups of user accounts. See [Maintaining User Groups](#) on page 533 for details.

Service Configuration

Use this folder to modify settings for the Web Access Server and the system on which you installed PI, set up SSL communications, and specify the attributes for mailing a report link.

The navigation tree for the Service Configuration folder appears as follows:



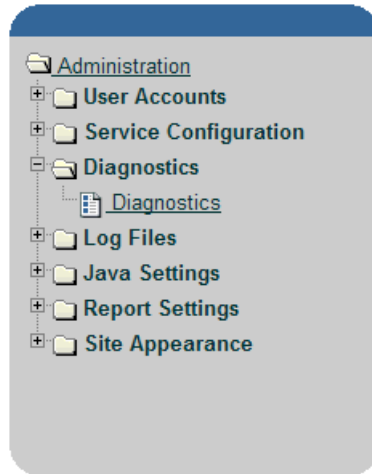
The Service Configuration folder contains the following:

- **Web Access Server page**
Use this page to change the current settings for the Web Access Server. The *Web Access Server* controls the PI client's access to the PI database. See [Changing the Web Access Server Settings](#) on page 565 for details.
- **SSL Service page**
Use this page to enable SSL communication between the Web Access Server and any supported Web browser. See [Changing the SSL Service Settings](#) on page 570 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 575 for details.
- **Mail Server page**
Use this page to specify the attributes for mailing a report link. See [Changing Mail Server Settings](#) on page 583 for details.

Diagnostics

Use this folder to view Java and operating system information for the Web Access Server and information pertaining to database connections.

The navigation tree for the Diagnostics folder appears as follows:



The Diagnostics folder contains the following:

- Diagnostics page

Lets you view the following information for the Web Access Server and Database:

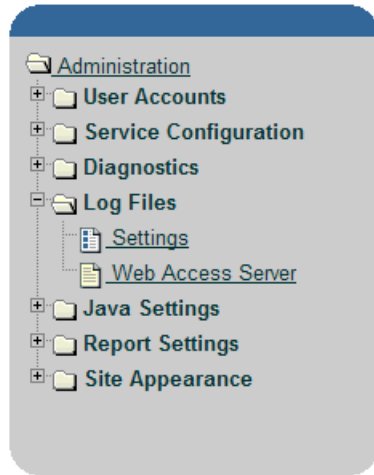
- Web Access Server. Lets you view the Java version, the Java vendor, the Java vendor's URL, and the operating system's architecture, name, and version number.
- Database. Lets you view the name, maximum number of connections, connection status, and current connections.

See [Displaying Diagnostics](#) on page 543 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 544 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 546 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 Java Settings page. This page lets you change heap size and arguments. See [Changing Java Settings](#) on page 548 for details.

Deployed Items

Use this folder to change settings for deployed items (reports and URLs) and set up catalog views. The navigation tree for the Deployed Items folder appears as follows:



It contains the following:

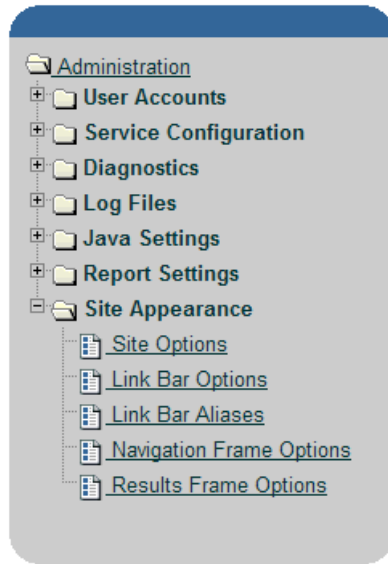
- Deployed Items page

This page lets you set attributes for reports and URLs. For more information, see [Changing Deployed Item Settings](#) on page 549.

- 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 552 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 556 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 556 for details.

- Link Bar Aliases page

Lets you show or hide and change the name of any of the links. For more information, see [Changing Link Bar Aliases](#) on page 559.

- **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 561 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 562 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 563 for details.

- **Deploy URL page**


Lets you deploy URLs to the Web Access Server's catalog to quickly access a website that you use frequently. For example, if your company has more than one Web Access Server system, you may want to deploy a URL for a particular report that you view frequently on one of those systems.

When you select a URL, the specified website displays in your current browser window or in a new browser window, depending how you have configured the Web Access Server. See the *HP Performance Insight Guide to Building and Viewing Reports* for more information about how to use this feature.

- System folder

All users can access the reports deployed to the `System` folder. 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 PI Administrator can deploy reports to this folder.

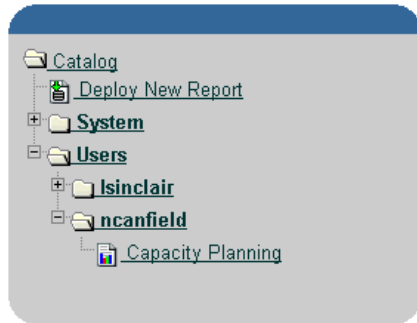
Following is an example of how the navigation pane might appear when reports are deployed to the `System` folder:



- Users folder

This folder appears only after reports have been deployed to the Web Access Server by a user or by the PI 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 PI clients (Report Builder, Report Viewer, Web Access Server, and the Management Console) require users to log on.

By default, PI 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 PI client applications.

You can perform most of the tasks described in this section using the `userimport` and `groupimport` commands. For more information about these utilities, see the *HP Performance Insight Reference Guide*.



The procedures contained in this section assume that you are logged on to the Web Access Server using the default PI Administrator user account or a user account that has PI 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 n values.	3
<input type="checkbox"/> User can continue to login for n 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 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`, `arcaes`, 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 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 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 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 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

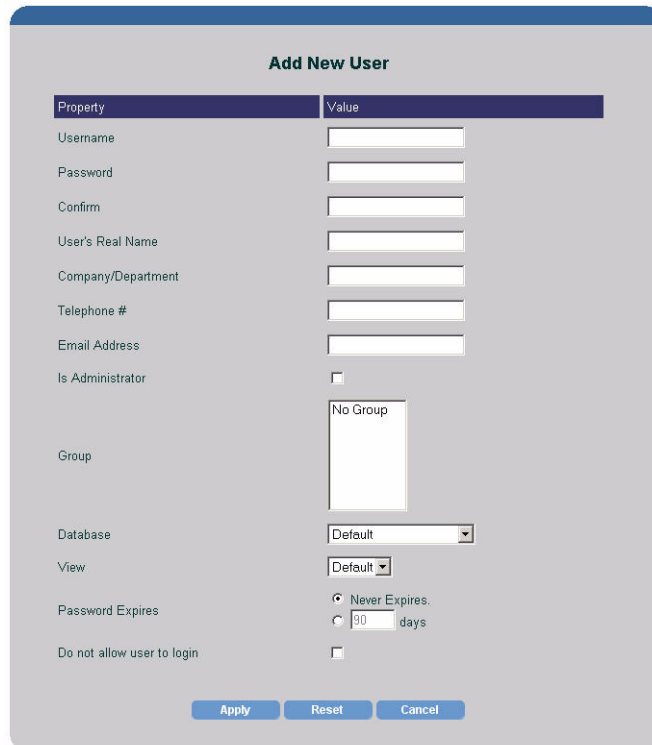
To create a user account:

- 1 From the Links bar, select **Administration**.

- In the Administration tree, expand User Accounts, and 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 Type 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 (they can be blank) and can be a maximum of 255 characters.
 - The password can be a minimum of four 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 select 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 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 575.
 - 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 552.

- 9 To select 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 type 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 525 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 565 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. 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.

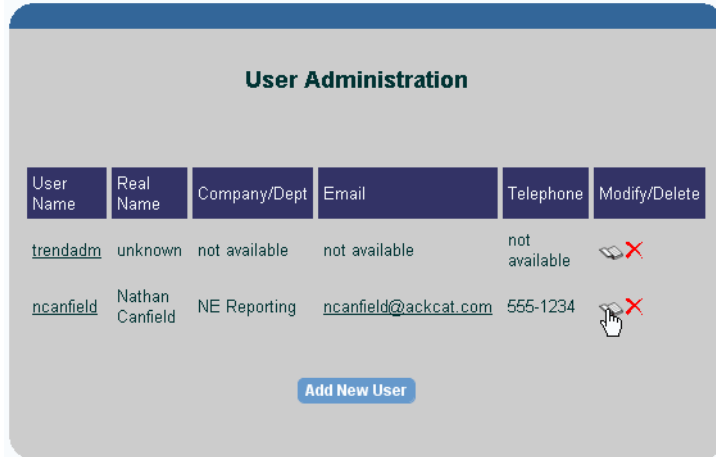
- 11 Click **Apply**.
- 12 Click **OK** to confirm the new user account.


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

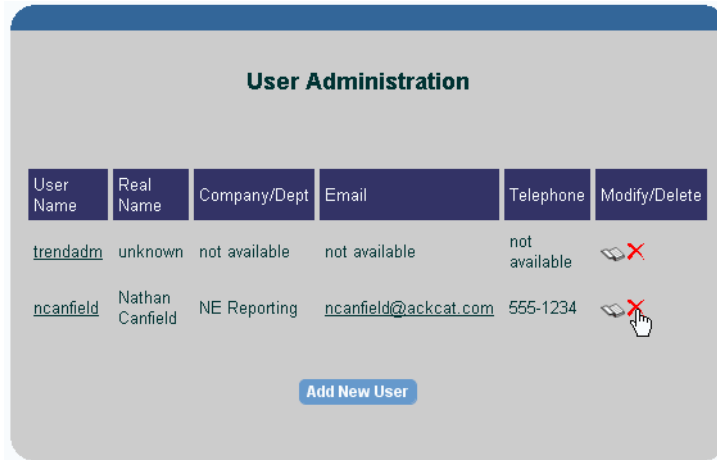
After you confirm your changes, the User Administration page lists the updated user account's information.

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 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**.
- 2 In the Administration tree, expand User Accounts, and click **Users**. The User Administration page opens, which lists the user accounts.
- 3 Click the email address for the user account. Your email application appears with a mail message addressed to the specified user.
- 4 Complete the mail message, and 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 PI Administrator.

You can perform most of the same tasks described in this section using the **userimport** and **groupimport** utilities. For more information, see the chapter about Account Administration Utilities in the *HP 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 563 for details

Creating a User Group

To create a group:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and click **Groups**. The Contents of Groups page opens.

Group Administration

Contents of Groups

Groups

Group Name	# Members	Modify/Delete
No Groups Found		

Users

	User Name	Modify User
👤	trendadm	📄
👤	ncanfield	📄

[Add New Group](#)

- 3 Click **Add New Group**. The Add New Group page opens.

The screenshot shows a web form titled "Add New Group". At the top, there is a header "Add New Group". Below the header, there are two input fields: "Group:" with a placeholder "NAME" and "Description:". Below these fields, there are two side-by-side lists: "Available Users/Groups" containing "ncanfield" and "trendadm", and "Members" which is empty. Between the lists are two buttons: "Add >>" and "<< Remove". At the bottom center is a "Next" button.

- 4 To create a new user group, do the following:
 - a In the Group box, type 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 click **Add**. Note that groups can belong to other groups.

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 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.

05is 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 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 select 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 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 click **Remove**. The user or user group is removed from the Available Users/Groups list.
- 7 Click **Next**, and 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 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


Filters are applied during report generation and applied to any report you run while logged on.

You can filter on any common column or columns in the property tables. If you were to create multiple filters on a single group, the filters are ANDed together; that is, you create a `cust_id=3` filter and another `location_id=2` filter in the same group. The SQL created to retrieve the data for any user belonging to this group would be as follows:

```
where cust_id=3 AND location_id=2
```

You do have know the property table columns names and they data that they contain. If you are using the property import/export feature for a report pack, you should have an idea of the columns and data are is in those columns.

To create a user group filter:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand User Accounts, and 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, type 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 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 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 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 click **Delete**.
- 6 Click **Finish**.

Example: Data Filtering Using Groups

You do not need to move reports around into user folders to limit the data a user can access. User folders and the catalog views are used only to limit the reports that a user can run. They do not limit the data users can see when they run a report. Instead, you can use groups and filters to limit the data that users can see.

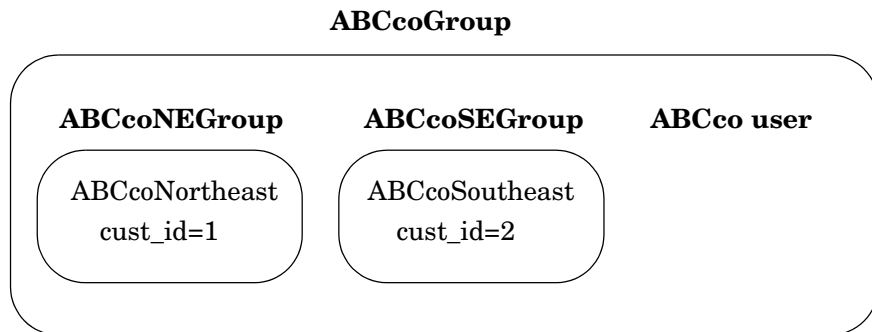
This example uses groups that belong to other groups to show how to set up hierarchical groups. Additionally, it shows how the filters are ORed or ANDED together.

To set up data filtering using groups:

- 1 Create the following users:
 - ABCco
 - ABCcoNortheast
 - ABCcoSoutheast
- See [Creating a User Account](#) on page 527 for details.
- 2 Create two groups that each have a different filter; each group will only filter for the data for the user assigned to them. For example, you could set up these groups as follows:
 - a Create ABCcoNEGroup with a filter of `cust_id=1`.
 - b Assign the ABCcoNortheast user to the ABCcoNEGroup group.

- c Create ABCcoSEGroup with a filter of `cust_id=2`.
 - d Assign the ABCcoSoutheast user to the ABCcoSEGroup group.
- 3 Create a master group called ABCcoGroup. Do not create any filters for it.
 - 4 Assign the ABCcoNEGroup and ABCcoSEGroup groups to ABCcoGroup.
 - 5 Assign the ABCco user to ABCcoGroup.

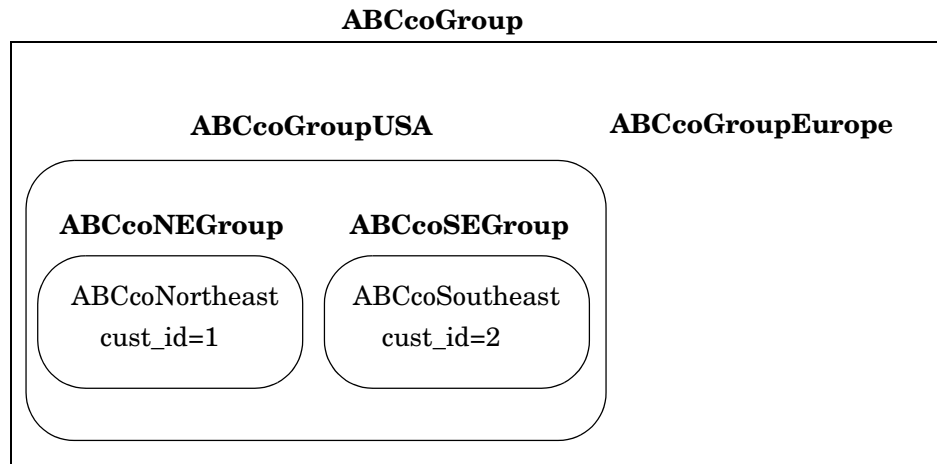
When you associate a group with a group, the filters for the master group are inherited by its subgroups and ORed together. Thus, any user, in this case ABCco, associated with the ABCcoGroup would have the filter where `cust_id=1 OR cust_id=2`. Thus, your grouping would look like the following:



Similarly, you could have added the ABCco user to both the ABCcoNEGroup and ABCcoSEGroup groups to have the same results. Though, that may be less desirable because it is easier to track grouping hierarchies than to determine the groups to which a user belongs to when a user belongs to multiple groups.

Because this is a hierarchal structure, you could have more than two layers of groups. For example, you could have the ABCcoGroupUSA and ABCcoGroupEurope groups. In this instance, ABCcoNEGroup and ABCcoSEGroup would belong to ABCcoGroupUSA, and ABCcoGroupUSA and

ABCcoGroupEurope would belong to the ABCcoGroup group. This would allow you to have users assigned to regions, countries, and company wide filters. This grouping would look like the following:



The filters are applied during report generation and is applied to any report the user runs while logged on. Thus, if you log on as ABCcoNortheast and run a report from the `system` folder, the ABCcoNortheast user will only see data where `cust_id=1`. Similarly, if you log in as ABCcoSoutheast and run the same report from the `system` folder, you will only see data where `cust_id=2`.

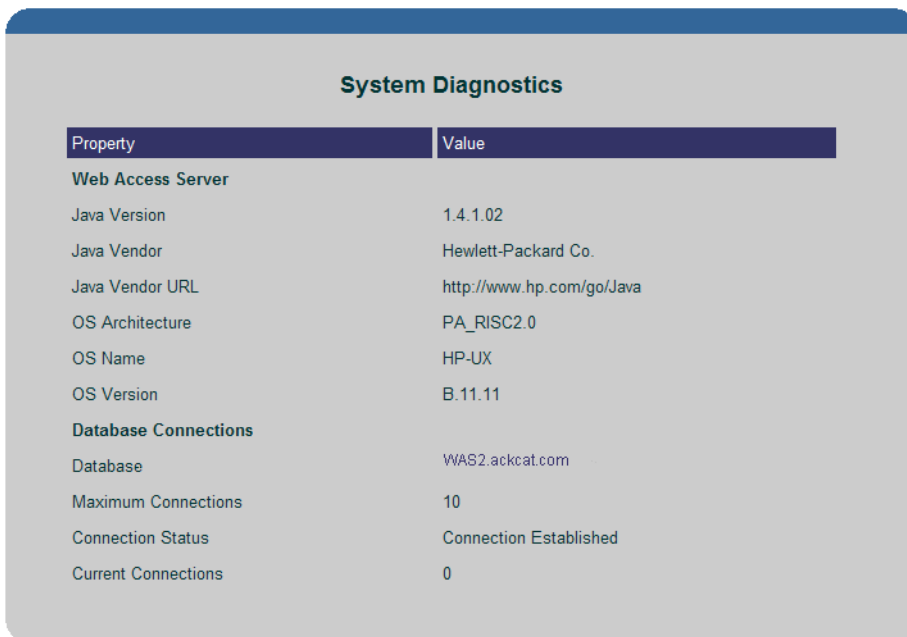
Displaying Diagnostics

You can view the following information for the Web Access Server and Database by using the Diagnostics page:

- **Web Access Server.** Lets you view the Java version, the Java vendor, the Java vendor's URL, and the operating system's architecture, name, and version number.
- **Database.** Lets you view the name, maximum number of connections, connection status, and current connections.

To view diagnostics:

- 1 On the Links bar, click Administration.
- 2 In the Administration tree, expand **Diagnostics**.
- 3 Click **Diagnostics**. The System Diagnostics page opens, which will appear similar to the following:



Property	Value
Web Access Server	
Java Version	1.4.1.02
Java Vendor	Hewlett-Packard Co.
Java Vendor URL	http://www.hp.com/go/Java
OS Architecture	PA_RISC2.0
OS Name	HP-UX
OS Version	B.11.11
Database Connections	
Database	WAS2.ackcat.com
Maximum Connections	10
Connection Status	Connection Established
Current Connections	0

Changing Log File Settings

To change log file settings:

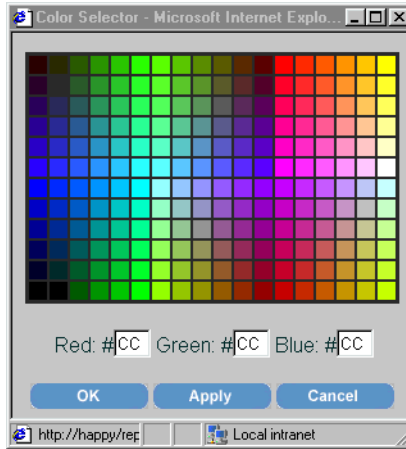
- 1 On the Links bar, click Administration.
- 2 In the Administration tree, expand **Log Files**.
- 3 Click **Settings**. The Log File Settings page opens.

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:
 - Type 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**.



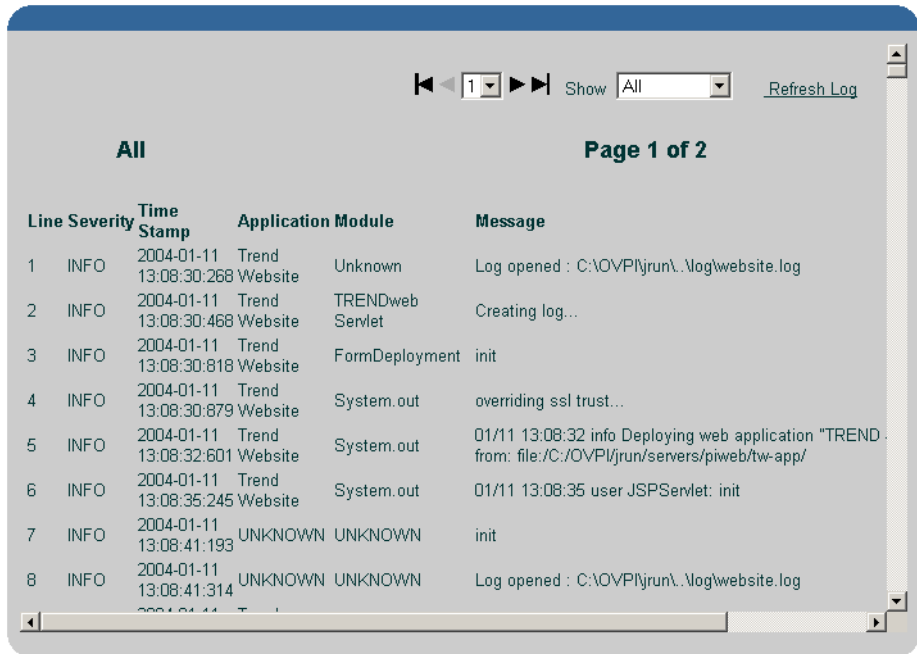
- 5 Click **Apply** again.

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.



When Rotation is enabled, the logger creates files with names like website0.log, website1.log etc. in the logging directory. The default logging configuration for builder, viewer and piadmin would be No-Rotation.

- 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 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 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}

- 3 Do any of the following:
 - a In any of the Heap Size boxes, type 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, type 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, see 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 Deployed Item Settings

To change deployed item settings:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Deployed Items, and click **Deployed Items**. The Deployed Items page opens.

Attribute	Value
Generated Reports Location	C:\OVPI\data\schedules
Deployed Items Location	C:\OVPI\reports\deploy
Image Cache Location	C:\OVPI\run\servers\piweb\tw-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"/>

Apply Reset

- 3 Do one or more of the following:
 - a In the Generated Reports Location box, type the name of the folder in which you want to generate your schedules.
 - b In the Deployed Items Location box, type the name of the folder in which you want to keep deployed reports and URLs.
 - 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 PI 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. 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 [Starting and Stopping HP Performance Insight Processes](#) on page 103 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 Deployed Items 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
















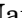
Name:



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 type 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  (**Select All**) to select all the reports in the folder.To deselect all the reports in a folder, click  (**Deselect All**).

- 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.

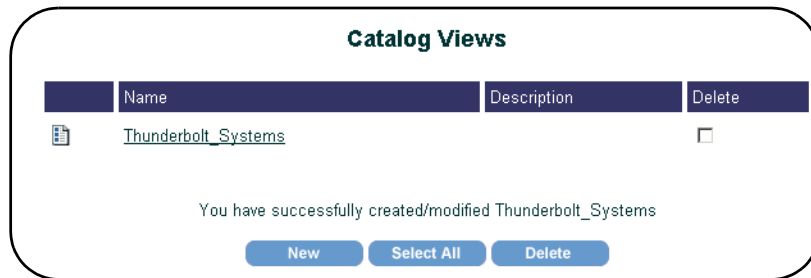
- 8 To specify the database from which the report should get its data, select a database from the *Database* list. There is one database located next to each report.

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 **Select All** 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 Deployed Items 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  (**Edit**) 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 click **Site Options**. The Site Options page opens.

Attribute	Value
Site Title	HP OpenView Performance Insight
Show Profile	<input checked="" type="checkbox"/>

Apply Reset

- 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 **Refresh** in the browser.

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 click **Link Bar Options**. The Link Bar Options page opens.

Attribute	Value
Title	HP OpenView, Performance Insight
Title Size	4
Align	right
Height	100
Background Color	#FFFFFF Select Color
Background Image	/webview/images/backgroundtile.gif
Image	/webview/images/hpinvent.png
Image Width	74
Image Height	61
Image Alt	Hewlett Packard
Image X Location	25
Image Y Location	25
Image Background Width	200
Link Background Focus On Color	#336699 Select Color
Link Background Focus Off Color	#6699CC Select Color

[Apply](#) [Reset](#)

- 3 Do any of the following:
- a In the Title box, type a title for the link bar. This title appears below the link bar.
 - b Click the Title Size arrow, and select the point size you want for the title.
 - c Click the Align arrow, and 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:
 - Type 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 type specific color values. After you select a color, click **Apply** then **OK**.

f In the Background Image box, type 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 type 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:

Windows:

```
%OvInstallDir%\nonOV\PiJBoss\405\server\piweb\deploy\tw-app.war\images
```

UNIX :

```
/opt/OV/nonOV/PiJBoss/405/server/piweb/deploy/tw-app.war/webview/images
```

g In the Image box, type the location for the image that will appear to the left of the link bar. See [step f](#) of this procedure for information about where the image should be located.

h In the Image Width and Image Height boxes, type the width and height for the image in pixels.

i In the Image Alt box, type 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, type 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, type 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:

- Type 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 type 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 finish making changes, click **Apply** and then click **Refresh** in the browser.

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 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 checked="" 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 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, type a new name in a ***Link_Name* Link Alias** box.
- 4 When you finish making your changes, click **Apply** and then click **Refresh** in the browser.

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 click **Navigation Frame Options**. The Navigation Frame Options page opens.

Navigation Frame Options

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, type 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:
 - Type 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 type specific color values. After you select a color, click **Apply** then **OK**.
- 4 When you finish making changes, click **Apply** and then click **Refresh** in the browser.

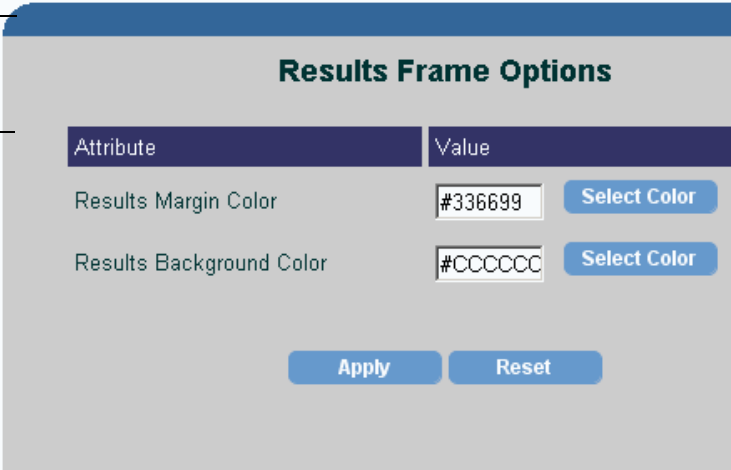
Changing Results Frame Options

To change results frame options:

- 1 From the Links bar, select **Administration**.
- 2 In the Administration tree, expand Site Appearance, and 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:
 - Type 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 type specific color values. After you select a color, click **Apply** then **OK**.
- 4 When you finish making changes, click **Apply** in the Results Frame Options page and then click **Refresh** in the browser.

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 type="button" value="Browse..."/>		Local report file to deploy
Save To	<input type="text" value="/users"/> <input 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, type 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, type 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, type an optional description for the report.

- 7 Click **Deploy**.

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 565](#))
- SSL ([page 570](#))
- System Manager ([page 575](#))
- Mail Server ([page 583](#))

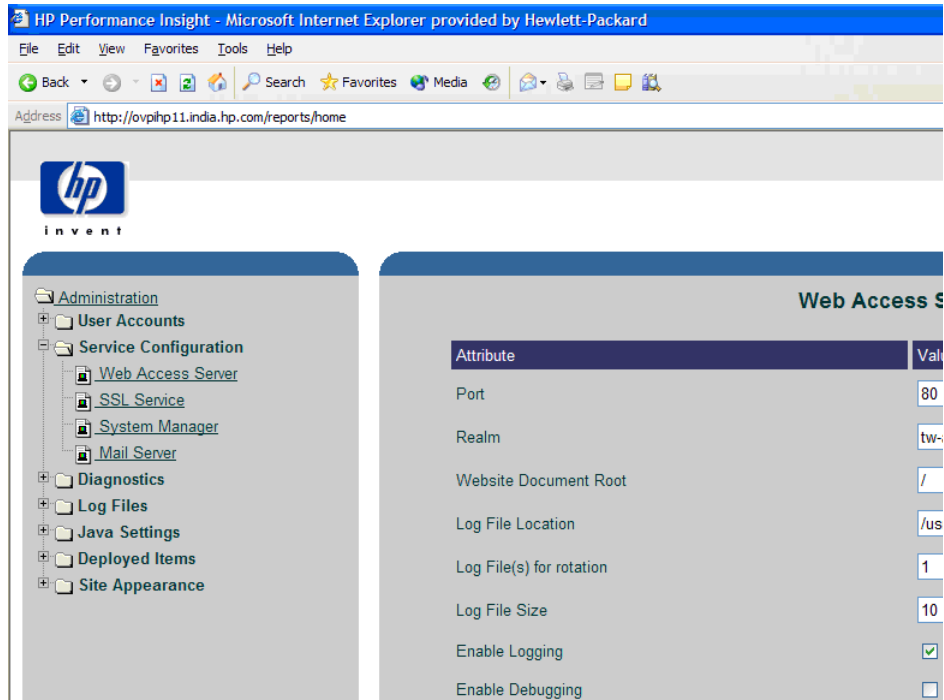
Changing the Web Access Server Settings

You can change the following settings for the Web Access Server: port number, realm, 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.



- 3 Do one or more of the following:
 - a In the Port box, type the port number of the Web Access Server.
 - b In the Realm box, type 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 **Log on** on the Web Access Server.



You cannot type any value in the Website Document Root box.

- c To change the path of the server's log file, type the path in the Log File Location box.



You can only use ASCII characters in the path.

- d In the Log File(s) for Rotation box, type the number of files to rotate. This number should be greater than zero.
- e In the Log File Size box, type the size of the Log file. This number should be an integer greater than zero.
- f 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.
- g 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.
- h *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 PI, your system is configured for headless support. *Headless support* uses software to render graphics. 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
```

- i In the Timeout box, type 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.

- j 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, see the *HP 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 568.

- **Account ID.** Type 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 finish your changes, stop and start the Web Access Server. See [Starting and Stopping HP Performance Insight Processes](#) on page 103 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 specified a user account without Administrator privileges in the Account box, you must reset the Administrator privileges to access the Administration link of the Web Access Server. You can do this by editing the `config.prp` file or by using a URL to log on to the Web Access Server.

Editing the config.prp File

To reset Administrator privileges by editing the `config.prp` file:

- 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 PI.

- 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:

```
JBOSS_HOME/server/piweb/deploy/tw-app.war/WEB-INF.
```

In this instance, `JBOSS_HOME = opt/OV/nonOV/PiJBoss/< 405 or 4.23.000 > for UNIX,`

and, `JBOSS_HOME= C:\< Program Files or Program Files (x86)>\HP OpenView\nonOV\ PiJBoss\< 405 or 4.23.000 > for Windows.`

- 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 [Starting and Stopping HP Performance Insight Processes](#) on page 103.

Using a URL

To log on to the Web Access Server with Administrator privileges, type the following URL:

```
http://hostname/reports/home?ov_user=trendadm
```

In this instance, *hostname* is the name of the Web Access Server.

Changing the SSL Service Settings

You can enable or disable 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. 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. See [Importing Your Digital Certificate](#) on page 573 for details.

To enable or disable SSL:

- 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"/>
Enable non-secure HTTP Port	<input checked="" type="checkbox"/>
Secure Port	443
Keystore Password	•••••
Certificate Information	<input type="checkbox"/> Generate a New Self-Signed Certificate Key Expires: Mon Feb 28 16:08:36 EST 2005

Certificate Identifying Data	
Name	izzybell.americas.cpqcorp.net
Organization	ackcat
Organization Unit	reporting
Location	Nashua
State/Province	NH
Country Code	US

Apply Reset



If you enabled SSL communication when you installed PI, the Keystore Password and Certification Information options will not appear.

- 3 To enable SSL, do the following:
 - a Select the **Enable SSL** check box.
 - b Click to clear the **Enable non-secure HTTP Port** box if you want to turn off the non-secure port for the Web Access Server.

When you disable this option, users will have to access the Web Access Server using `https://hostname`.
 - c Leave the value in the Secure Port box as is (the default is 443) or change the port number in the Secure Port box.

- d In the Keystore Password box, type a password of six characters or more.



Record the password and keep it in a secure place. You will need it to generate a signed certificate.

After you generate your self-signed certificate, it will be stored in the certificate repository. This password lets the Web Access Server access the certificate repository.

The remaining fields request information that is necessary when you generate an SSL key that must be verified and signed by a certificate-signing authority. This is information used by the person who checks by the certificate to ensure that the certificate is valid.



The following fields (Country, Organization, Organization Unit, Location, State/Province, and Country Code) can only contain ASCII characters.

- e Type your two-letter country code in the Country box (for example, “US”).
- f Type the name of your organization (for example, “ABC Systems, Inc.”)
- g Type the name of your department or division (for example, “Northeast Sales”).
- h Type the name of your city or locality (for example, “Los Angeles”).
- i Type the name of your state or province (for example, “California”).
- j Click **Apply**.

After you click **Apply**, the **Generate a New Self-Signed Certificate** check box appears on the SSL Service Configuration page along with the expiration date for your certificate.

When you need to generate a new certificate, change the information in the Certificate Identifying Data section of the page, if necessary, select this option, and click **Apply**.

- k Stop and restart the Web Access Server to make the change take effect. See [Starting and Stopping HP Performance Insight Processes](#) on page 103 for this procedure.



This self-signed digital certificate will allow you to enable SSL. 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. See [Importing Your Digital Certificate](#) for details.

- 4 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 [Starting and Stopping HP Performance Insight Processes](#) on page 103 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 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 Type the following at the command line to generate your CSR:

```
keytool_path -certreq -file filename.csr -keystore  
JBOSS_HOME/lib/keyStore -alias ovpi
```

keytool_path is the location of the keytool utility in the `DPIPE_HOME` directory.

JBOSS_HOME = opt/OV/nonOV/PiJBoss/405 on UNIX and
C:\Program Files\HP OpenView\nonOV\PiJBoss\405 on
Windows.

filename is the name you want to give to the .csr file.

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.

The following prompt will display:

```
Enter keystore password:
```

- 3 Type your keystore password, and press **ENTER**.

This is the same password that you provided when enabling SSL during the installation of PI or using the Web Access Server.



When you type your password, it is echoed to the screen. 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:* Type the following at the command line to import the server certificate:

```
keytool -import -file filename -keystore JBOSS_HOME/lib/  
trustStore
```

- 3 Type the following at the command line:

```
keytool -import -file filename -keystore JBOSS_HOME/lib/  
keyStore -alias ovpi -trustcacerts
```

Where JBOSS_HOME= opt/OV/nonOV/PiJBoss/405 on UNIX and
C:\Program Files\HP OpenView\nonOV\ PiJBoss\405 on
Windows.]

The following prompt will display:

Enter keystore password:

- 4 Type your keystore password, and press **ENTER**.



When you type your password, it is echoed to the screen. 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 PI 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:

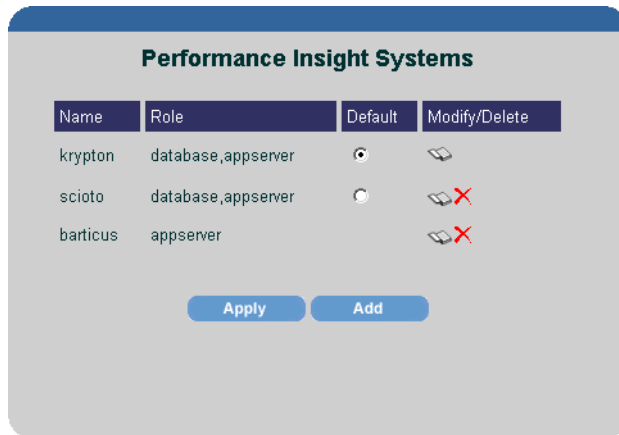
- Select a default system ([Choosing the Default System](#) on page 576).
- Add systems ([Adding a System](#) on page 576).
- Modify the settings of current systems ([Modifying a System's Settings](#) on page 579).
- Delete systems ([Deleting a System](#) on page 583).

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 575.
- 2 Click **Add**.

The System Information page opens.

Property	Value
Name	<input type="text"/>
Host Name	<input type="text"/>

- 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"/>

- 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 opens.

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 opens.


Property	Value
Admin Server Port	80
Authorization Realm	tw-app

Back Finish Cancel

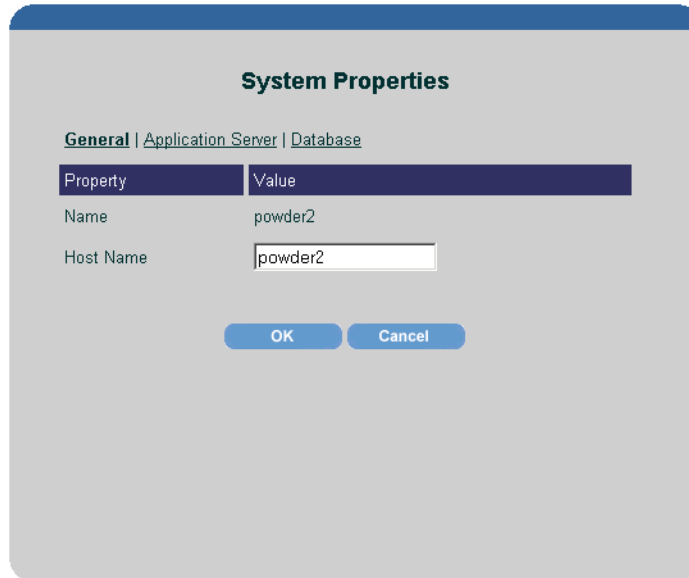
- 9 If applicable, do any of the following in the Database Properties page:
 - a In the Admin Server Port box, change the port number at which the Web Access Server listens for connections.
 - b In the Authorization Realm box, change the realm name that the database server uses for authentication.
- 10 Click **Finish**.

Modifying a System's Settings

To modify the settings of a system:

- 1 Open the Performance Insight Settings page. See [Displaying the Performance Insight Systems Page](#) on page 575.
- 2 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.

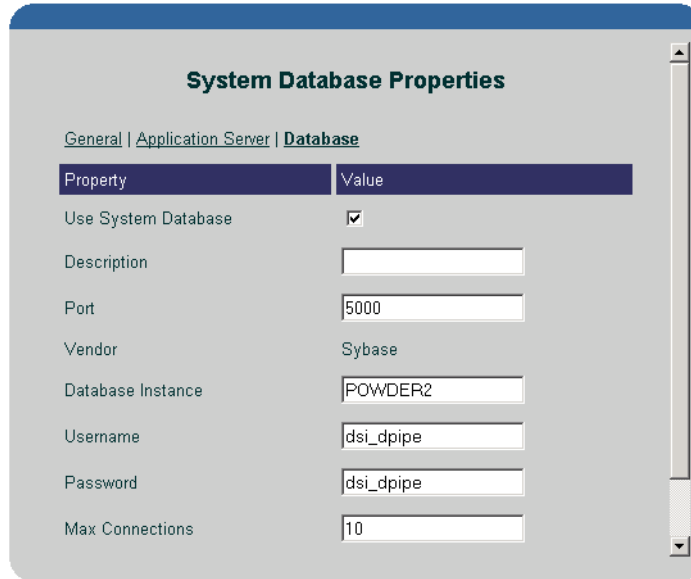
- 3 To change the general properties of the system, do the following:
 - a Click the **General** link. The System Properties page opens.
 - b In the Host Name box, change the host name of the system.
 - c Click **OK**.
- 4 To change Web Access Server properties:

- a Click the **Web Access Server** link. The **System Web Access Server Properties** page opens.

Property	Value
Use System Application Server	<input checked="" type="checkbox"/>
Port	80
Context	/
Connect Using SSL	<input type="checkbox"/>
SSL Port	443
Authorization Realm	tw-app

- 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.



System Database Properties

[General](#) | [Application Server](#) | **[Database](#)**

Property	Value
Use System Database	<input checked="" type="checkbox"/>
Description	<input type="text"/>
Port	<input type="text" value="5000"/>
Vendor	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, type a description for the database.
 - In the Port box, type 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, see the *HP Performance Insight Installation Guide*.


- In the Username box, type the username for the database.

- In the Password box, type 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 finish making your changes, stop and start the Web Access Server.


For information about how to stop and start PI services, see [Starting and Stopping HP Performance Insight Processes](#) on page 103.

Deleting a System

To delete a system:

- 1 Open the Performance Insight Settings page. See [Displaying the Performance Insight Systems Page](#) on page 575.
- 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** icon  when you view a report. (See the *HP 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 View Settings page. To display this page, from the Links bar, select **Preferences**. In the Preferences tree, expand Deployed Items 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 type="text" value="localhost"/>
Port	<input type="text" value="25"/>
Header Message	<input type="text" value="*** This message originated from HP OVPI User Account: <username> ***"/> Use the <username> tag to substitute the username within the phrase.
URL Message	<input type="text" value="Report Link: <url>"/> Use the <url> tag to substitute the report URL within the phrase.

[Apply](#)

- 3 Do any of the following:
 - a In the Host box, type the host name of the SMTP mail server.
 - b In the Port box, type the port number of the SMTP mail server.
 - c In the Header Message box, type 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, type 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). The following table lists the available variables.

Table 12 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 12 Graph X-Axis Time Format Variables

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 PI.

The time formats displayed in the `config.prp` file may appear similar to the following:

```

axis.hourly.time.format=E hh:mm a
axis.daily.time.format=E MMM d
axis.weekly.time.format=MMM d
axis.monthly.time.format=MMMMM
axis.summary.hourly.time.format=E hh:mm a
axis.summary.daily.time.format=E MMM d
axis.summary.weekly.time.format=MMM d
axis.summary.monthly.time.format=MMMMM

```

- 2 Change the formats you want using the variables listed in [Table 12](#).
- 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 [Starting and Stopping HP Performance Insight Processes](#) on page 103 for details.

16 Archiving Data

Storing data in data tables for a long period causes the size of the database to increase. This adversely affects both summarization and reporting performance. The data archiving function of PI helps prevent such performance bottlenecks. The data archiving function uses archive tables instead of data tables to store data for longer durations.

The data archiving function helps to store the data for longer duration without affecting the summarization and reporting performance of PI. This function enables you to store PI data in a different tablespace in the PI database schema.

Archiving data involves the concepts of archive table, archive flag, archive retention time, active table, and data retention profiles.

Archive Tables

An archive table stores the data present in a data table for a longer period of time. The `datapipe_manager` command creates archive tables for existing and new data tables. One data table can have only one archive table associated with it. A data table to which an archive table is associated is called active table.

An archive table is partitioned in the same way as its associated data table. Whenever the structure of a data table is changed, the `datapipe_manager` command makes the same changes in the archive table of that data table. If a data table is deleted, the `datapipe_manager` command deletes the associated archive table.



An archive table always has a table type `Archive`.

Data Retention Profiles

A data retention profile for a data table is created when the data table is created in the database. The `datapipe_manager` command uses the information specified in the TEEL file to create a retention profile. The `datapipe_manager` command deletes a retention profile if no data table has reference to the retention profile.

A retention profile contains the following information:

- **Name:** Specifies the name of a retention profile. By default, the `datapipe_manager` command uses a combination of table category, retention time, and time type for the name of a profile. For example, If the parameters specified in the TEEL file are as follows:
 - `TableTimeType`: hourly
 - `TableCategory`: DevResCiscoSwitch
 - `TableRetention`: 12

The retention profile name will be DevResCiscoSwitch12hourly.

- **Retention:** Specifies the number of days to age the data in a data table.
- **Archive Retention:** Specifies the number of days to age the data in an archive table. The default archive retention value is based on the time type of the table. For example, if the time type of a table is hourly, then the archive retention time is 30 days.
- **Archive Flag:** Indicates whether an archive table is required or not. By default, the value of the archive flag is set to OFF. You must set the value of the archive flag to ON to retain the data of a data table.
- **Description:** Specifies the description of a retention profile.

Archiving Data Process

While installing a report pack, the `datapipe_manager` commands creates a retention profile for each data table based on the retention time specified in the TEEL file. The `datapipe_manager` command uses default values for the profile name and archive retention time, based on the data table time type, to create a retention profile. This profile is registered into a PI dictionary.



If a retention profile of the same name and the same retention time already exists, the `datapipe_manager` command uses that profile and does not create a new profile.

To archive the data of a data table, you must set the value of the archive flag to ON in the retention profile. The `datapipe_manager` command creates the archive table for a data table when its archive flag is set to ON. The `datapipe_manager` command assigns the retention profile of the data table, which is also called active table, to the archive table also.

If the archive flag is ON for a data table, the aging process moves the data from the data table to the archive table when the specified retention time for that data table is over. For example, if the retention time for a data table is set to 10 days and the archive retention time is set to 60 days, the `db_delete_data` command (the aging process) moves the data from the data table to the archive table after 10 days. In the archive table, data is available for the next 60 days. The `db_delete_data` command deletes the data from the archive table after 60 days. This means that a total of 70 days data is available in the active table and the archive table together. See [Managing Retention Profiles](#) on page 593 for more information about changing the archive retention time and archive flag for a data table.

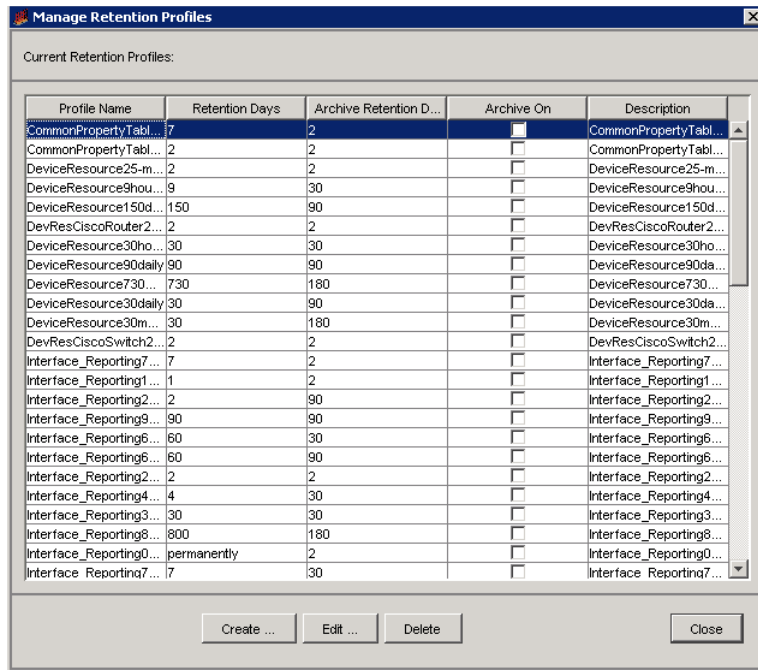
Managing Retention Profiles

To archive data efficiently, you must effectively manage retention profiles.

View Retention Profiles

To view all current retention profiles:

- 5 Open the HP Performance Insight Management Console window.
- 6 Select **Administration->Tables** option.
- 7 Select **File->Retention Profiles->Manage Retention profiles...** The Manage Retention Profiles dialog box opens.



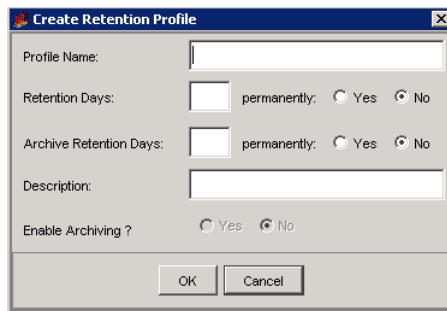
The Manage Retention Profiles dialog box enables you to create, modify, and remove retention profiles.

Add Retention Profiles

After installing a report pack, you can create and add new retention profiles in the PI system. If an error or problem occurs while creating a retention profile, an error message is displayed.

To add a retention profile:

- 1 Open the Manage Retention Profiles dialog box. See [View Retention Profiles](#) on page 593 to know how to open the Manage Retention Profiles dialog box.
- 2 Click **Create...** in the Manage Retention Profiles dialog box. The Create Retention Profile dialog box opens.



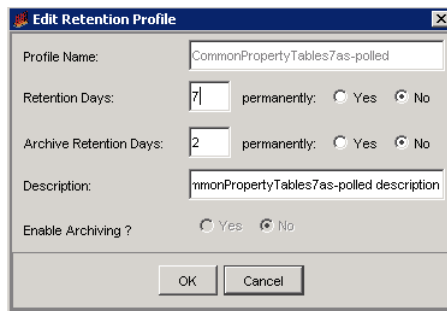
This dialog box allows you to configure the following information:

- **Profile Name:** Name of the retention profile. The maximum length for the profile name is 64 characters.
 - **Retention Days:** Number of day for which the data will be available in the data or active table. To specify a permanent retention, you need to select Yes from the permanently radio button.
 - **Archive Retention Days:** Number of days for which the data will be available in the archive table. To specify a permanent archive retention, you need to select Yes from the permanently radio button.
 - **Description:** Description of the retention profile. The maximum length for the profile description is 255 characters.
 - **Enable Archiving?:** A flag to indicate whether an archive table will be created or not.
- 3 Click **OK** in the Create Retention Profile dialog box to create and save the retention profile in the system.

Modify Retention Profiles

PI does not allow you to modify the default profiles created during the product installation. To modify a retention profile, you must complete the following steps:

- 1 Open the Manage Retention Profiles dialog box. See [View Retention Profiles](#) on page 593 to know how to open the Manage Retention Profiles dialog box.
- 2 Select the profile that you want to modify from the *Profile Name* list.
- 3 Click **Edit...** The Edit Retention Profile dialog box opens.



You can modify all parameters other than the name of the retention profile.

- 4 Click **OK** to update the retention profile. .



You can verify if the retention profile has been updated from the *Current Retention Profiles* list.

Remove Retention Profiles

To remove a retention profile, follow these steps:

- 1 Open the Manage Retention Profiles dialog box. See [View Retention Profiles](#) on page 593 to know how to open the Manage Retention Profiles dialog box.
- 2 Select the profile that you want to delete from the *Current Retention Profiles* list.
- 3 Click **Delete**. A user confirmation dialog box opens.

4 Click **Yes** to remove the profile from the system.



You cannot remove a default profile or a profile that is in use.

Viewing Data and Internal Table Details

You can view detailed information about any data or internal table by double-clicking on the table in the Table Manager viewing pane.

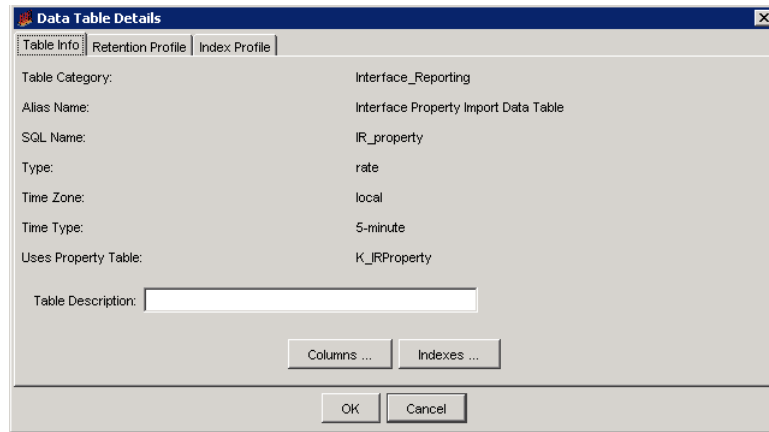
For example, to view detailed information about the rate_IR_property table, you must complete the following steps:

- 1 In the Table Manager navigation pane, select Interface_Reporting from the **Category** combobox.

The viewing pane displays all the tables associated with the Interface_Reporting category.

- 2 Double-click the **IR_property** table in the viewing pane.

The Data Table Details dialog box opens as shown in the following figure.



The Data Table Details dialog box lets you do the following:

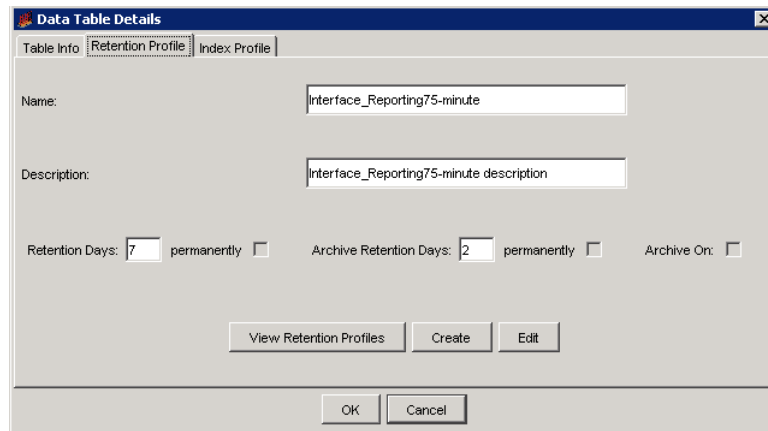
- View information about the table's columns.
- Add a description.
- Set the retention profile.
- View information about the table's columns.
- View information about the corresponding indexes.

Viewing Data Table Column Information

The Table Info box in the Data Table Details dialog box 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 Retention Profile

The Retention Profile tab in the Data Table Details dialog box shows the current data retention settings for the selected data table.



The screenshot shows the 'Data Table Details' dialog box with the 'Retention Profile' tab selected. The dialog has three tabs: 'Table Info', 'Retention Profile', and 'Index Profile'. The 'Retention Profile' tab contains the following fields and controls:

- Name:** A text box containing 'Interface_Reporting75-minute'.
- Description:** A text box containing 'Interface_Reporting75-minute description'.
- Retention Days:** A spin box set to '7', followed by a 'permanently' checkbox (unchecked).
- Archive Retention Days:** A spin box set to '2', followed by a 'permanently' checkbox (unchecked).
- Archive On:** A checkbox (unchecked).
- Buttons:** 'View Retention Profiles', 'Create', and 'Edit' buttons are located below the retention settings. 'OK' and 'Cancel' buttons are at the bottom of the dialog.

The Retention Profile page allows you to do the following tasks:

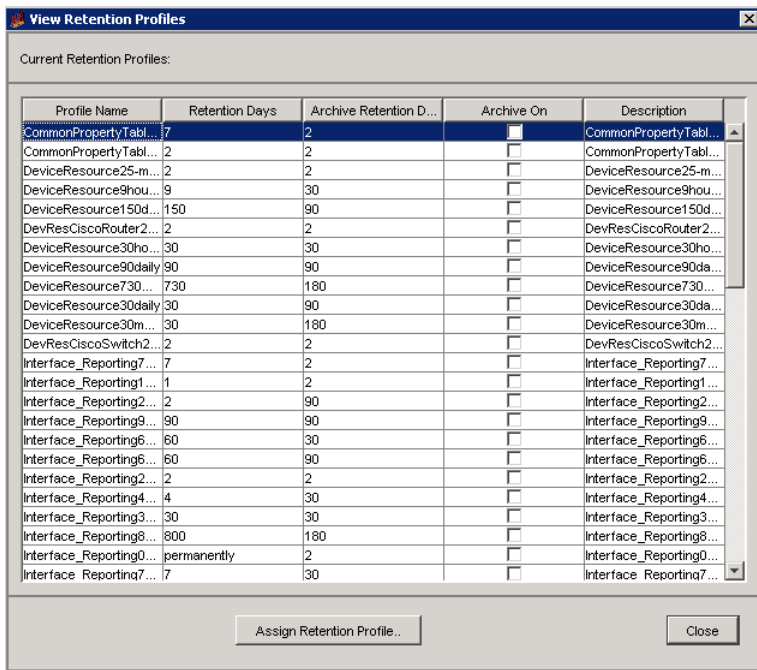
- Assign a different retention profile to a data table.
- Create a new retention profile.
- Modify an existing retention profile.

Assigning a Different Retention Profile to a Data Table

You can assign a different retention profile to a data table to change archiving related attributes. If archiving is on for the newly assigned profile, the `datapipe_manager` ensures the availability of the archive table for the data table to which the profile is assigned. A message is displayed after the archive table availability check is completed. The `datapipe_manager` command creates an archive table if the archive table does not exist.

To assign a different retention profile to a data table:

- 1 Open the Data Table Details dialog box. See [Viewing Data Table Column Information](#) on page 598 to know how to open the Data Table Details dialog box.
- 2 Select the Retention Profile tab.
- 3 Click **View Retention Profiles**. The View Retention Profiles dialog box opens.



- 4 Select the retention profile that you want to assign to the data table from the *Current Retention Profiles* list.
- 5 Click **Assign Retention Profile...**

Assigning a Retention Profile to Multiple Data Tables

You can assign single retention profile to multiple data tables. If archiving is on for the newly assigned profile, the `datapipe_manager` command ensures the availability of the archive table for the data table to which the profile is assigned.

- 1 Open the HP Performance Insight Management Console window.
- 2 Select **Administration->Tables** option.
- 3 Select the Type from the viewing pane of the Database Table Management window. The list of data tables is displayed in the navigation pane of the Database Table Management window.
- 4 Select the tables to which you want to assign a retention profile from the navigation pane by pressing SHIFT or CTRL key.
- 5 Select **File->Retention Profiles->Assign Retention Profile....** The View Retention Profiles dialog box opens.
- 6 Select the retention profile that you want to assign from the *Current Retention Profiles* list.
- 7 Click **Assign Retention Profile...** to assign the selected profile to all the selected tables.



The **Assign Retention Profile** option will be disabled when you select **View** as the **Object Type**.

Adding a Data Table Description

The Description box in the Data Table Details dialog box lets you add a description of the table. Enter text in the Description box, and click **OK**. The description must not 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. 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	A logical, unique, descriptive name of the column in the table.
SQL Name	The SQL name of the column in the table.
Data Type	The PI data type for the column.
SQL Type	The database type of the data for the column.
Asset Column	If selected, specifies that the column contains an asset property for use in an asset summary.
Object Identifier	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, specifies that the column can be NULL.
Default Value	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. The Indexes window opens.



The Indexes window shows the following information about each index for the table.

- Index Name** The name of the index for the table.
- Unique** If selected, specifies that the index is unique.
- Clustered** If selected, 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 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.

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 PI 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.
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.
Keep archive data for (days)	This is the number of archive retention days for the data in the table. Data is discarded when it is age + 1 days old. The available archive 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 archive retention period for the time type. This value is set using the Archive Retention box in the Set Table Details dialog box.
Archive	This is a flag to indicate whether an archive table will be created or not.

Guidelines to set up Archiving

The following are the guidelines for setting up archiving:

- Create a new retention profile with appropriate retention time for active and archive table.
- Archiving can be enabled for a profile during profile creation itself or can be turned on after assigning this profile to a data table.
- Data tables with low granularity time types like hourly and below, are the good candidates for archiving.

The following table shows recommended minimum retention times, when archive is “on”.

Table 13 Recommended Retention Time

Table Time Type	Minimum retention time (days)
as-polled	1
sub-minute	1
5-minute	1
Hourly	2
Daily	Default, set by Report Packs.
Weekly	Default, set by Report Packs
Monthly	Default, set by Report Packs
Quarterly	Default, set by Report Packs
Yearly	Default, set by Report Packs

- Assign the above created retention profile to the data table(s) that needs to be archived.
- Create different profiles depending on the table’s time type, table type, retention times and table category and so on. This way it is easy to manage and maintain not only retention profiles but also corresponding active and archive tables.

- If `db_delete_data` is taking more than 1 hour to finish, increase the number of threads by adding `-c` option in the `trendtimer.sched` file. Ex. `db_delete_data -c 5`
- When archive is turned “on”, make sure enough storage is available for data files to grow, especially that are associated to “`dpipe_archive_seg`” and “`dpipe_archive_ind_seg`” tablespaces.
- The following is the list of archive reports and corresponding data tables on which reports are built. This will guide you in selecting a right archive report when a data table is set for archive.

Interface Reporting

Table 14 Reports with Corresponding Data Tables

SN	Report	Data Tables
1	<code>IR_log.rep</code>	<code>RIR_ProcedureLog</code>
2	<code>System_Performance.rep</code>	<code>SHIRPerformance</code> , <code>SDIRPerformance</code>
3	<code>Adhoc_Interface_Search.rep</code>	<code>SHIRDevPorts</code>
4	<code>Availability_Detail.rep</code>	<code>SRIRDevPorts</code> , <code>SHIRDevPorts</code> , <code>SDIRDevPorts</code>
5	<code>Capacity_Planning.rep</code>	<code>SHIRDevPorts</code> , <code>SDIRDevPorts</code> , <code>SMIRDevPorts</code> , <code>SD42SDIRDevPortsFORE</code> , <code>SD42SDIRDevPortsDOW</code>
6	<code>Discard_Exception.rep</code>	<code>SHIRDevPorts</code> , <code>SDIRDevPorts</code> , <code>SMIRDevPorts</code>
7	<code>Error_Exceptions.rep</code>	<code>SHIRDevPorts</code> , <code>SDIRDevPorts</code> , <code>SMIRDevPorts</code>
8	<code>Exception_Hot_Spots.rep</code>	<code>SHIRDevPorts</code> , <code>SDIRDevPorts</code> , <code>SMIRDevPorts</code>
9	<code>Utilization_Exception.rep</code>	<code>SHIRDevPorts</code> , <code>SDIRDevPorts</code> , <code>SMIRDevPorts</code>

Table 14 Reports with Corresponding Data Tables

SN	Report	Data Tables
10	Grade_Of_Service.rep	SHIRDevPorts, SDIRDevPorts, SMIRDevPorts
11	Near_Real_Time.rep	SRIRDevPorts, SHIRDevPorts, SDIRDevPorts, SMIRDevPorts
12	Near_Real_Time_Simple.rep	SRIRDevPorts, SHIRDevPorts
13	Near_Real_Time_Snapshot.rep	SRIRDevPorts, SHIRDevPorts, SDIRDevPorts, SMIRDevPorts
14	Service_Level_Management.rep	SHIRDevPorts, SDIRDevPorts
15	Top_Ten.rep	SDIRDevPorts
16	Unreachable.rep	SHIRDevPorts
17	Utilization_Volume.rep	SHIRDevPorts, SDIRDevPorts, SMIRDevPorts

System Resource**Table 15 Reports with Corresponding Data Tables**

SN	Report	Data Tables
1	Adhoc_Search.rep	SD_SR_SysVolXcep
2	Appl.rep	SD_AppOVPA_byApp, SD_AppOVPA_SysApp
3	Capacity_filesystem.rep	SD_SR_SysXcep, SD_SR_SysFSXcep, SD42SD_SR_SysFSXcep_FORE
4	Capactiy_memory.rep	SH_SR_SysExec, SD_SR_SysExec, SM_SR_SysExec
5	Execsum.rep	SH_SR_SysExec, SD_SR_SysExec, SM_SR_SysExec

Table 15 Reports with Corresponding Data Tables

SN	Report	Data Tables
6	Forecast_cpu.rep	SD42SD_SR_SysXcep_Fore, SD_SR_SysXcep
7	Forecast_filesystem.rep	SD_SR_SysXcep, SD_SR_SysFSXcep, SD42SD_SR_SysFSXcep_Fore
8	Forecast_memory.rep	SD_SR_SysXcep, SD42SD_SR_SysXcep_Fore
9	Hotspots.rep	SR_SR_Kernel, SD_SR_SysXcep
10	NRT_quickview.rep	SR_SR_Kernel
11	NRT_snapshot.rep	SR_SR_Kernel
12	Quickview.rep	SH_SR_SysXcep, SD_SR_SysXcep
13	Snapshot.rep	SH_SR_SysXcep
14	SystemAvailability.rep	SH_SR_SysUpGShift, SD_SR_SysUp
15	Topten.rep	SD_SR_SysVolXcep, SD_SR_SysXcep, SD_SR_SysFSXcep
16	Toptentran.rep	SDTran_ByTransaction, SDTran_Transaction, SHTran_ByTransaction
17	TransactionBySystem.rep	SHTran_Transaction, SDTran_Transaction
18	CPUQuickView.rep	SH_SR_CPU(A), SD_SR_CPU
19	SystemResourceDiskDetails.rep	SH_SR_Disk, SD_SR_Disk

Table 15 Reports with Corresponding Data Tables

SN	Report	Data Tables
20	SystemLogicalVolumeDetails .rep	SH_SR_LogicalVolume, SD_SR_LogicalVolume
21	SystemNetInterface.rep	SH_SR_NetInterface, SD_SR_NetInterface
22	ProcessDetails.rep	SH_SR_Process, SD_SR_Process

Limitations of Archiving

Archiving has the following limitations:

- It is currently supported only on Oracle, and not supported on Sybase.
- Archive reports are currently available for IR and SR report packs only.

17 Location Independent Reporting and Copy Policies

Location Independent Reporting (LIR)

PI is typically deployed in a distributed network environment with central and satellite server configurations. Satellite servers poll data and perform summarization tasks. To report on the polled data, the data is copied to a centralized repository, usually the central server. This leads to enormous duplication of data and puts additional load on the network PI is monitoring.

LIR is designed to address the above issues and provide distributed data management. In the PI distributed environment, LIR makes reporting independent of the location of data on the network. Data from multiple databases can automatically be included in a single report, solving the problem of frequent copying of data from the satellite server to the central server.

The advantages of setting up LIR for reporting are:

- Reduce network traffic. Configure data management such that a minimal amount of data is transferred from the satellite servers to the central server for report generation. An ideal scenario is to allow the satellite servers to store most of the data (raw, rate, and perhaps summarizations) locally.
- Integrate data distributed on the network into a single report.
- Create LIR mappings and copy policies for customizing data copy functionality on the network.

LIR removes the requirement to update the central server with data from each satellite server using point-to-point **trendcopy** procedures.

LIR Mapping

When LIR mapping is configured on the central server, a mapping of the type of data contained on each satellite server is recorded. The satellite server should be registered with the central server. The registration can be done during PI installation (see [Install Wizard Updates for LIR Support](#)) or using the Add Database Reference Wizard. The wizard is available in the System and Network Administration application on the PI Management Console. For more information, see [Managing Your System](#) on page 363.

LIR mapping information includes the following:

- **Data type:** Single property table and all property tables in a report pack or table category
- **Aggregation type:** Rate, hourly, daily, weekly, monthly, quarterly and yearly

When reports are to be generated, the reporting engine will use LIR mappings to determine the satellite server that contains the requested data. The query is done on a property table basis. If a property table is not listed in the data mapping table, the local or default database is queried for data.

LIR queries are applicable for a report element (graph or table) that is created for an individual node or set of nodes. For more information about selecting nodes, see Chapter 6 of *HP Performance Insight Guide to Building and Viewing Reports*. LIR queries are also applicable if two or more report elements are linked by using drilldowns. For more information about drilldowns, see Chapter 5 of *HP Performance Insight Guide to Building and Viewing Reports*.

LIR Mappings and Copy Policies

When you create an LIR mapping, based on the Aggregation type, copy policies are created. An LIR mapping may have zero or more associated copy policies. The copy policy is auto-generated on the central server. A copy policy defines the data table type, the source satellite server which has data to be copied, and the data aggregation type.

LIR Implementation Example

Rate data on satellite servers, hourly data copied to central server

This example can be applied when Near Real Time (NRT) reports are required on the central server. Earlier versions of PI provided this facility by copying rate data from all the satellite servers. But this might not be feasible in most implementations purely due to amount of rate data collected every polling cycle. Moreover, there will be redundant rate data maintained between satellite and central servers.

With LIR, NRT reports can be viewed without having to copy rate data on the central server. Customers have to configure satellite server to perform hourly summarization and copy only hourly data on to the central server. Hourly data is smaller in volume when compared to rate data.

- Install PI and the desired report packs on a central server and one or more satellite servers.
- Register the satellite server with central server.
- Configure each satellite server to poll for a specific list of devices and perform hourly summaries only; higher level aggregations (daily, weekly, and so on) will be performed on the central server.
- Configure LIR on the central server and add LIR mappings.
- Run reports from the central server. The report will now automatically include data from the satellite servers.

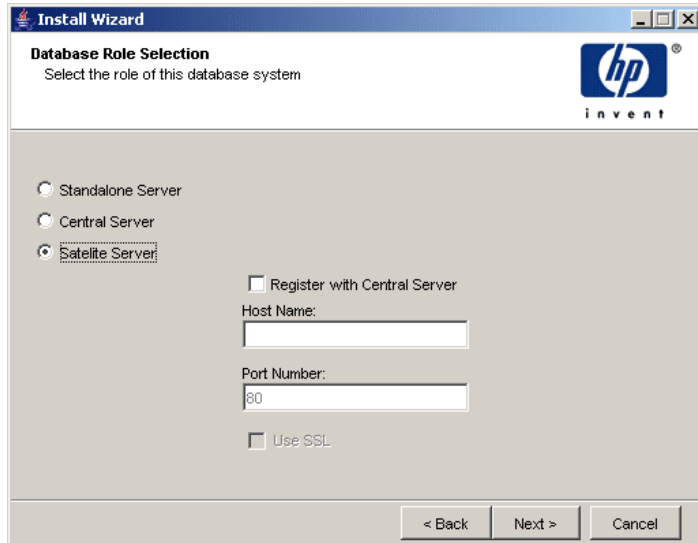
Install Wizard Updates for LIR Support

During PI installation, you will be prompted to select a role for the PI server system. For complete installation instructions, see Chapter 4 of the *HP Performance Insight Installation Guide for Oracle Databases*.

The Server Role Selection window presents three options:

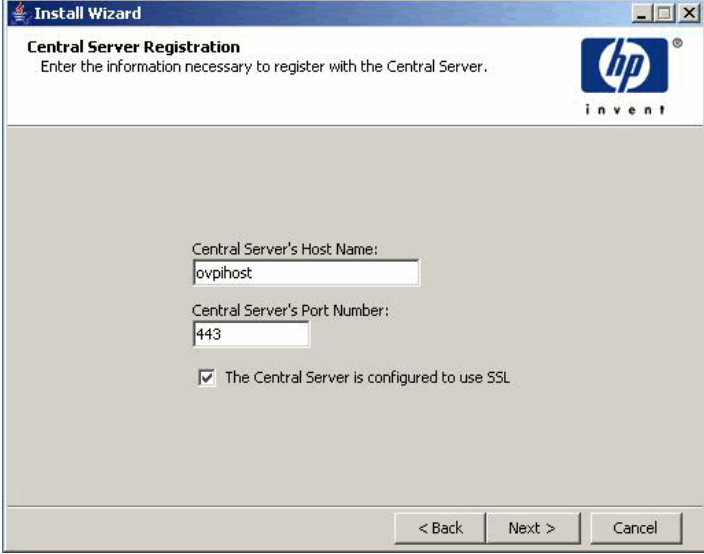
- Standalone Server
- Central Server

- **Satellite Server**



When you select the Satellite Server option, the Central Server Registration page opens.

The Central Server Registration page provides the option to specify the host name and port number (default value is 80) of the central server to register with. This is required to synchronize the information stored in the `systems.xml` file. For secure communication, you can enable SSL. The registration is completed at the end of the installation process.



The screenshot shows a dialog box titled "Install Wizard" with a sub-header "Central Server Registration". Below the sub-header, it says "Enter the information necessary to register with the Central Server." and features the HP logo with the word "invent" underneath. The main area contains two text input fields: "Central Server's Host Name:" with the value "jovpghost" and "Central Server's Port Number:" with the value "443". Below these fields is a checked checkbox labeled "The Central Server is configured to use SSL". At the bottom, there are three buttons: "< Back", "Next >", and "Cancel".

System.xml file updates

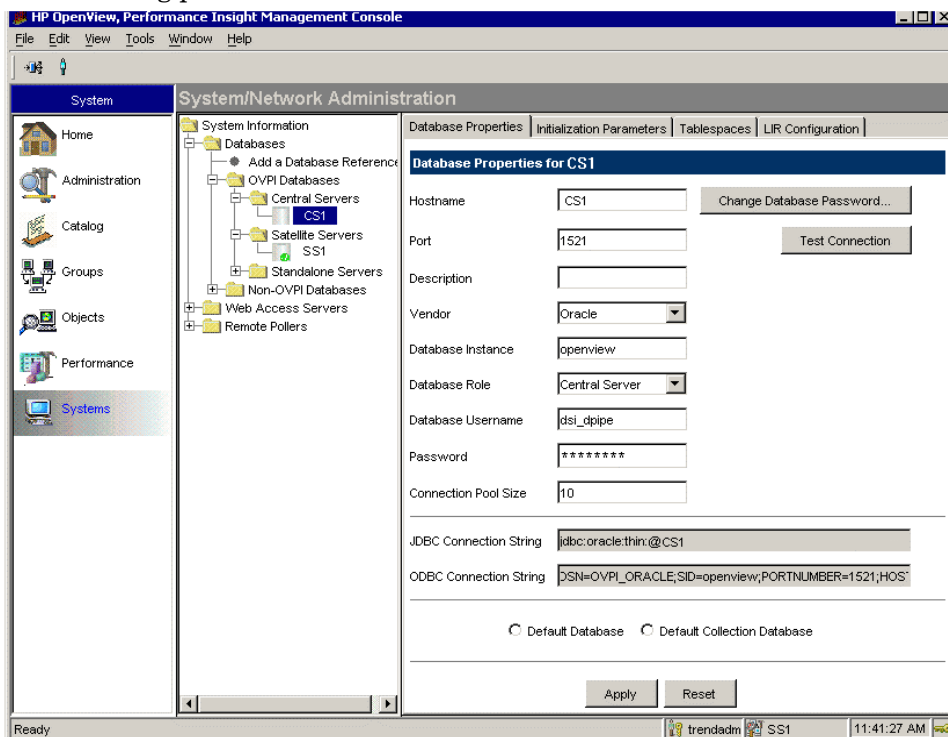
With the introduction of LIR, two new tags are added in the `DatabaseProperties` block of the `systems.xml` file. The first is the `Role` tag with acceptable values of Standalone (denoted numerically as 0), Central (denoted numerically as 1), Satellite (denoted numerically as 2). The second is the `LIREnabled` tag with acceptable values of true and false. The `LIREnabled` tag is only applicable for central servers.

With these changes in the `systems.xml` file, the application server will be responsible for the management of LIR access. When a client logs into the application server, it is told which database to connect to and will now also be told whether or not to use the LIR functionality. This single point of configuration will prevent the need to configure client-only systems for LIR.

Managing Performance Insight Server System Roles

To view the System and Network Administration window, click the **Systems** icon in the navigation pane of the Management Console. The System and Network Administration window opens.

The OVPI Databases folder in the navigation pane is separated into three sub-folders: Central Servers, Satellite Servers, and Standalone Servers. When you select the central server system name, the following window is displayed in the viewing pane.



The window allows you to configure the following information:

- Database Properties
- Initialization Parameters
- Tablespaces
- LIR Configuration (applicable only for central server)

To change the database system role:

- 1 Start the Management Console.



You must log on to the Management Console using an account that has Administrator privileges.

- 2 Click the **Systems** icon in the navigation pane of the Management Console. The System and Network Administration window opens.

- 3 Navigate to the OVPI Databases folder, and select the desired database system.

- 4 Click **Database Properties**.

The Database Properties window opens.

- 5 From the Database Role list, select the new role to be set and then click **Apply**.

The system will immediately move to the appropriate folder in the OVPI Databases folder and the change will also propagate to the system whose role was changed if it is different from the system that the PI Administrator is logged in to. The `systems.xml` file will be automatically updated to reflect the change.



To add a new database reference, you can use the Add Database Reference Wizard in the System and Network Administration application.

Manage LIR Mappings

This section discusses the following topics:

- [Enable LIR](#)
- [Disable LIR](#)
- [Add a New LIR Mapping](#)
- [View LIR Mappings](#)
- [Delete LIR Mappings](#)
- [Modify LIR Mapping](#)

Enable LIR

To enable LIR, do the following:

- 1 Start the Management Console.



You must log on to the Management Console using an account that has Administrator privileges.

- 2 Click the **Systems** icon in the navigation pane of the Management Console.
The System and Network Administration window opens.
- 3 Navigate to the OVPI Databases folder, and select the central server name.
- 4 Click the **LIR Configuration**.
The LIR Configuration window opens.
- 5 Select the **LIR enabled** check-box located at the bottom of the window.

This will enable LIR for the central server and set the `LIRenabled` tag to `true` in the `systems.xml` file. You can then move on to [Add a New LIR Mapping](#).

Disable LIR

To disable LIR on a central server, do the following:

- 1 From the System and Network Administration window, navigate to the OVPI Databases folder, and select the central server name.
- 2 Click the **LIR Configuration**.
The LIR Configuration window opens.
- 3 Deselect the **LIR enabled** check box located at the bottom of the window.
A message box appears.
- 4 Click **Yes** if you want to remove the LIR mappings and the associated policies.

Add a New LIR Mapping

To add a new LIR mapping, do the following:

- 1 From the System and Network Administration window, navigate to the OVPI Databases folder, and select the central server name.

- 2 Click **LIR Configuration**.

The LIR Configuration window opens.



Make sure that you have enabled LIR for the central server. See [Enable LIR](#) on page 618.

- 3 Click **Add Mapping**.

The Add LIR Mapping window opens.

- 4 From the Select Satellite Server list, select a satellite server to add a mapping to.

- 5 Select one of the following data table options:

- **Category** (the report pack table category)

- **Property table** (individual property table)

- 6 Select the name from the drop down list.

- 7 Select the desired data types to map.

When you select a copy type, all the lower level aggregation types are selected automatically. For example, if you select daily option, hourly and rate mappings are automatically selected.

- 8 Click **Add to List**.

The LIR mapping is created and displayed in the Added mappings list.

- 9 Repeat [step 4](#) through [step 8](#) to add multiple LIR mappings.

- 10 After adding the mappings, click **OK**, and then click **Apply**.

The LIR mappings that were added are displayed in the viewing pane. A corresponding copy policy is also auto-generated on the central server.

View LIR Mappings

- 1 From the System and Network Administration window, navigate to the OVPI Databases folder, and select the central server that has LIR enabled.
- 2 Click the **LIR Configuration**.
The LIR mappings are listed in a table.
- 3 To filter the mappings, select a satellite server name from the Select Satellite Server list to view the corresponding mappings.

Delete LIR Mappings

- 1 From the System and Network Administration window, navigate to the OVPI Databases folder, and select the central server that has LIR enabled.
- 2 From the Edit menu, click **Delete LIR mappings**.
The Delete LIR mappings window opens.
- 3 From the list of mappings, you can filter by satellite server, data table type or both.
- 4 From the filtered list, select the mapping to be deleted.
 - To select consecutive objects, select the first item, press and hold down the **Shift** key, and select the last object.
 - To select nonconsecutive objects, press and hold down the **CTRL** key, and click each object.
- 5 Depending on what you selected, click **Delete Selected** or **Delete All**.
The LIR mappings are deleted.
- 6 Click **Apply**.
- 7 Click **OK** to return to the System and Network Administration window.

Modify LIR Mapping

You can modify the data type value that is mapped to a satellite server.

- 1 In the LIR Configuration window, select the satellite server corresponding to the LIR mapping to be modified.

- 2 The LIR mapping associated with the satellite server is displayed.
- 3 Select or deselect the data type option for the mapping and then click **Apply**.

Copy Policies

The LIR functionality supports the Copy Policy wizard to create copy policies for trendcopies from the satellite servers to the central server. You must launch the Copy Policy wizard from the central server.

Copy policies lets you customize and manage data copy functionality on the network. You can create policies to copy data for a particular data type from satellite servers to central server.

Each LIR mapping will have a corresponding copy policy. For more information, see [LIR Mapping](#) on page 612. Use the Copy Policy Manager application to create and modify the policies.

Copy Policy Implementation Examples

Example 1: Custom copy policy for a table category.

This example creates a custom copy policy for a table category. Daily data is copied from different satellite servers on to the central server.

- Install PI on a central server and one or more satellite servers.
- Install desired report packs on both central and satellite servers.
- Create custom copy policy for a table category. Set summarization as Daily and copy type as Data Only.

Because copy type is data copy, time variant data is copied and property data is ignored.



The procedure contained in the above example assumes that you have already created another copy policy with property only copy for the selected table category.

Example 3: Custom copy policy for an individual data table.

This example illustrates copy policy for an individual data table. Property and data tables for all hourly tables are copied from different satellite servers on to the central server.

- Install PI on a central server and one or more satellite servers.
- Install desired report packs on both satellite and central server.

- Create copy policy for an individual hourly data table. Set summarization as Property only and Data only. Do not select Property and Data option.

Start Copy Policy Manager

The Copy Policy Manager application is accessible from the PI Management Console. The application is available only when the Management Console is started on the central server.

The application allows you to perform the following tasks:

- Create copy policies.
- Edit copy policies
- View copy policies

To start the Copy Policy Manager application, do the following:

- 1 Start the Management Console.



You must log on to the Management Console using an account that has Administrator privileges.

- 2 Click the **Administration** icon in the navigation pane of the Management Console.

The Administration window opens. The 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. The Copy Policy Manager supports two view options. Depending on the view, the pane shows a tree. The tree contains folders. The folders contain items, which may include other items or additional folders.

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 (-).


View Configured Copy Policies

The Copy Policy Manager provides two view options:

- **Table category view:** Copy policies are displayed based on the table category in the viewing pane of the Copy Policy Manager window. You also have the option to further drill down in the navigation pane to retrieve copy policies filtered by table category.
- **Server view:** Copy policies are displayed based on satellite servers in the viewing pane of the Copy Policy Manager window. You also have the option to further drill down in the navigation pane to retrieve copy policies created for each satellite server. The satellite server view is the default view when you start Copy Policy Manager application.

View Copy Policies Associated with a Satellite Server

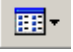
To view copy policies associated with a satellite server, do one of the following:

- From the View menu, select **Server View**.
- Click  (the View icon) and then select **Server View**.

The navigation pane displays the satellite servers. If you select a satellite server, the associated copy policies are displayed in the viewing pane.

View Copy Policies of a Table Category

To view copy policies of a table category, do one of the following:

- From the View menu, select **Table Category View**.
- Click  (the View icon) and then select **Table Category View**.

The navigation pane displays the table categories. If you select a table category, the associated copy policies are displayed in the viewing pane.

The items in each column are as follows:

Copy Policy ID	A unique identifier for copy policies.
Table Category	The name of the table category.
Source Table Alias	The name of the source table on the satellite server.

Target Table Alias	The name of the target table on the central server.
Copy Type	The data type to be copied. The possible values are property table, data table, or both.
Source Server	The server where the data is stored.
Destination Server	The name of the central server.
Copy Policy Enabled	Indicates whether data is copied from source server to the central server.

Manage Copy Policies

The Copy Policy manager application gives you the ability to construct new copy policies or modify existing copy policies.

Typically, you will create new policies in the following scenarios:


- When you install a new report pack.
- When you create a new collection wherein the central server processes some part of the aggregation.

Create a Copy Policy

To create a copy policy, complete the following tasks:

Task 1: Start the Copy Policy wizard.

From the Copy Policy Manager window, do one of the following to start the Copy Policy Wizard.

- Click  (the New Copy Policy icon).
- From the **File** menu, select **New Copy Policy**.

The Copy Policy Wizard opens. Click **Next**.

Task 2: Select the satellite server and data tables.

- 1 From the drop-down list, select the source satellite server from which data has to be copied. The target server for the data transfer is the central server to which you have logged in.
- 2 Do one of the following:
 - To select all the tables, click **Table Categories** and then select a table category from the drop-down list.
 - To select an individual table, click **Single Table** and then select a table from the drop-down list.
- 3 Click **Next**.

If you have selected table categories go to [Task 3](#); otherwise go to [Task 4](#).

Task 3: Define copy policy details for table categories.

The Copy Type Selection window opens.

Summarization Type	Copy Type
<input type="checkbox"/> Property Only Copy	Property only
<input type="checkbox"/> Rate Data Copy	Data only
<input checked="" type="checkbox"/> Hourly Data Copy	Property and Data
<input type="checkbox"/> Daily Data Copy	Data only
<input type="checkbox"/> Weekly Data Copy	Data only
<input type="checkbox"/> Monthly Data Copy	Data only

< Back Next > Cancel

- 1 Do one of the following:

- Select the **Default Copy Policy** option if you want to create a policy with default settings. The default policy will copy property and data tables for all hourly tables.

The default interval (hourly) is stored in the `$DPIPE_HOME/data/config.prp` file with the following name (DPIPE_HOME is the directory where you have installed PI).

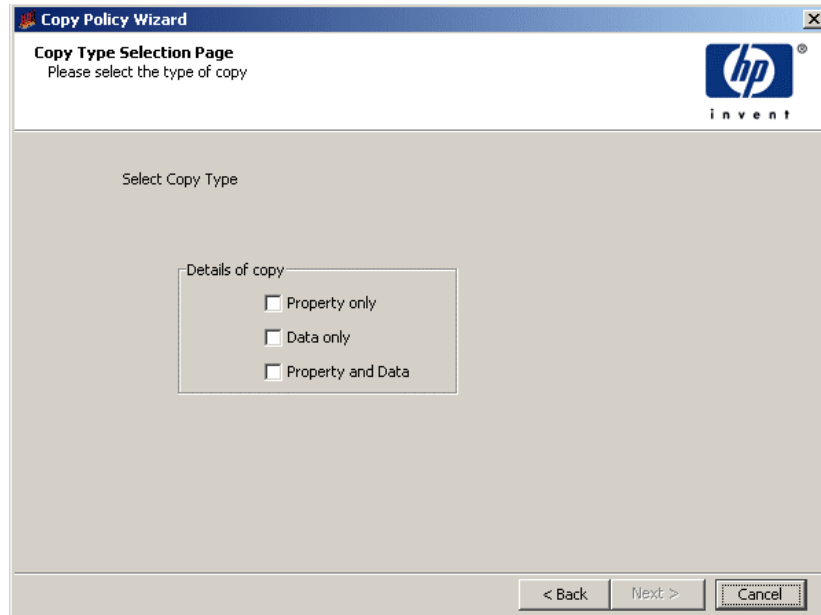
```
copy.timetype.default = hourly
```

- Select **Custom Copy Policies** option if you want to create a customized policy and then select the desired summarization type and its copy type.

2 Click **Next** and go to [Task 5](#).

Task 4: Define copy policy details for a single table.

The Copy Type Selection window opens.



1 Select the copy type for the table.

- If you select **Property only** option, only the property and keymap tables are copied to the central server.

- You can select both **Property only** and **Data only**. This is best option for base and hourly tables because property only copy happens frequently and data copy will happen at larger intervals.
- The **Property and Data** option is the typical **trendcopy** operation for the table.



For detailed documentation about the `trendcopy` command, see the *HP Performance Insight Reference Guide*.

- 2 Click **Next** and go to [Task 5](#).

Task 5: [Verify copy policy information and complete the procedure.](#)

- 1 Verify the information displayed in the Summary window.
- 2 Do one of the following:
 - Verify if the information is correct.
 - If necessary, click **Back** to return to a previous window to make a change. After you make a change, click **Next** until you return to the Summary window.
- 3 Click **Finish**.

Edit Copy Policy

To edit a copy policy, do the following:

- 1 From the Copy Policy Manager window, double-click the copy policy you want to modify. See [View Configured Copy Policies](#) on page 624.
- 2 You can modify the value of copy type from the Copy Type list.
- 3 Clear the **Enabled** check box if you do not want to copy data to the central server.

The **Enabled** check box is selected by default.

- 4 Click **OK**.



You cannot modify copy policies at the report pack or satellite server level. Every policy has to be manually edited.

Delete Copy Policy

The following options are available to delete copy policies.:

To delete all copy policies associated with a table category.

- 1 Set the view in the Navigation pane to the table category view. See [View Configured Copy Policies](#) on page 624.
- 2 Right-click the table category name and then select **Delete Copy Policy**.
- 3 Click **Yes** to delete all copy policies from the database.

To delete all copy policies on a satellite server.

- 1 Set the view in the Navigation pane to the server view. See [View Configured Copy Policies](#) on page 624.
- 2 Right-click the server name and then select **Delete Copy Policy**.
- 3 Click **Yes** to delete all copy policies from the database.

To delete an individual copy policy.

- 1 Depending on the copy policy you want to delete, set the view in the Navigation pane to either server view or table category view. See [View Configured Copy Policies](#) on page 624.
- 2 Right-click the copy policy you want to delete.
- 3 Select **Delete Copy Policy** and then click **OK** to delete the copy policy from the database.

18 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 Performance Insight (PI) MIB Browser to create a custom collection table.

This chapter describes how to do the following using the PI MIB Browser:

- Load a standard, vendor, or device-specific MIB. If the MIBs provided with PI are not sufficient for your data collection needs, you can load a custom MIB definition.

See [Loading MIBs](#) on page 640.

- 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 648.

- Set SNMP values.

See [Setting SNMP Values](#) on page 653.

- Create a custom collection table by selecting objects from the loaded MIBs.

See [Creating Collection Tables](#) on page 654.



In this chapter, the terms objects and variables are synonymous in the context of MIBs. See 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 118 for details.
 - b From the Tools menu, select **MIB Browser**. The MIB Browser window appears.
- Command line

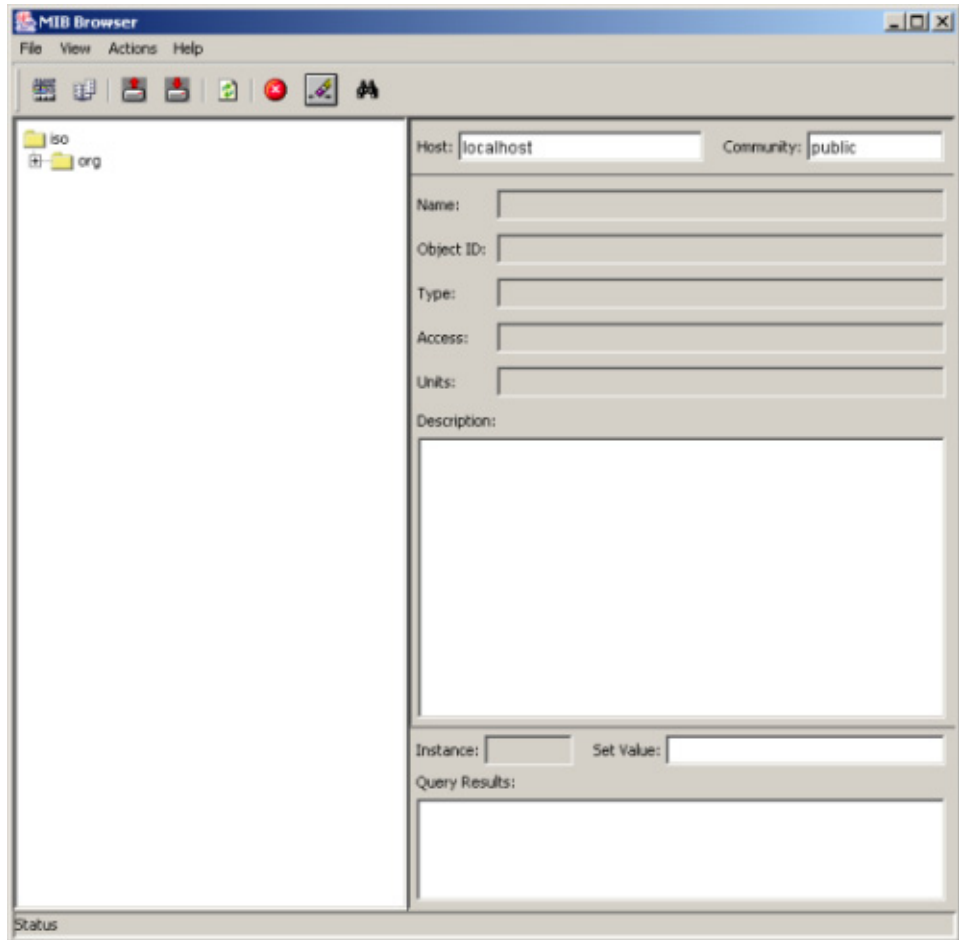
Type the following command:

mibbrowser



If **brass** utility (see [Using the brass_v3_seed Utility](#) on page 65) is installed in your machine, ensure **brassagt.exe** (for Windows) or **brassd** (for UNIX) is running.

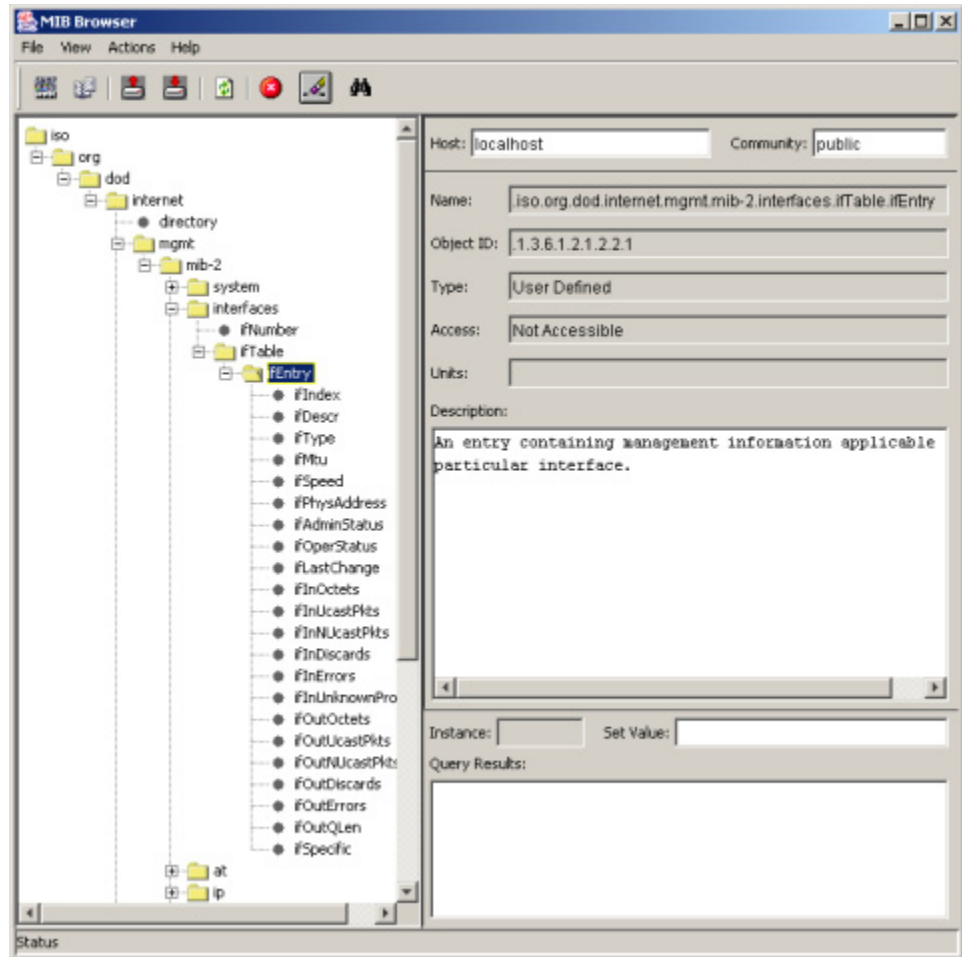
The following figure shows an example of the initial MIB Browser window.



In Windows system, you can also start the MIB Browser from the Start menu.



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 PI 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.

- SMNP 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 16 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 654 for details. This option is not required in standard PI installations. You do not need it if you are using PI solutions reporting.	
Open	This option lets you load MIB definition files into the MIB Browser. See Loading MIBs on page 640 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 SNMP Get request results section of the viewing pane.



Actions Menu

Table 17 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 648 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 653 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

PI stores MIBs in the default location: `DPIPE_HOME/mibs`; `DPIPE_HOME` is the environment variable that points to the directory where PI 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 642.

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:

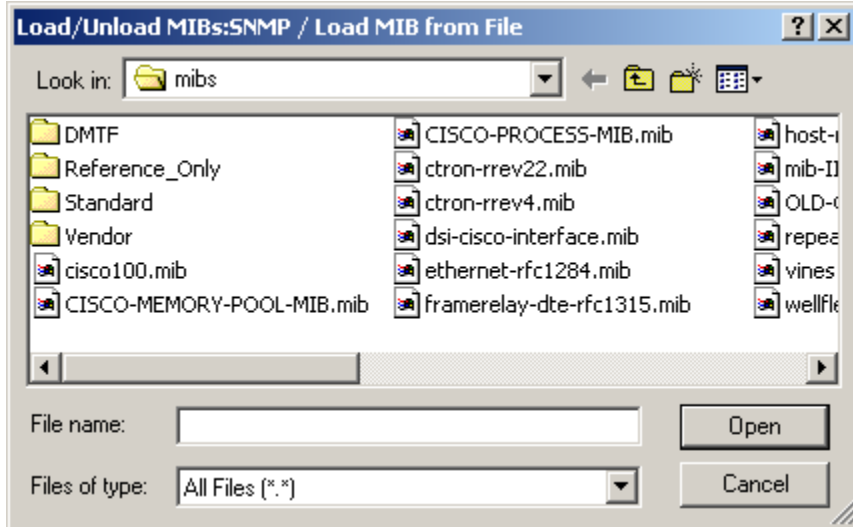
- 1 Do one of the following:
 - From the File menu, select **Open**.
 - Click .

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.

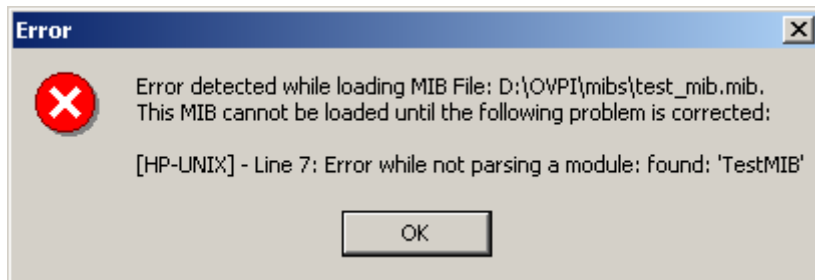
The Load/Unload MIBs:SNMP / Load MIB from File dialog box opens, which shows the mibs directory.



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.

▶ The warning messages are logged in `xnmlloadmib.log` within the following directories:


Windows:

%OvInstallDir%\log directory

Unix:

var/opt/OV/share/log directory


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 640 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 PI 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, community string, and SNMP Version.


Setting the Node

The current node, which appears in the Host box in the viewing pane, is the target for the SNMP functions. Type 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 PI 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:




Setting the SNMP Version

The MIB browser can communicate with SNMP Version 1 (SNMPV1) and SNMP Version 2 (SNMPV2). The default SNMP Version is SNMPV1. You can select SNMPV1 or SNMPV2 from the drop down list based on your requirement.

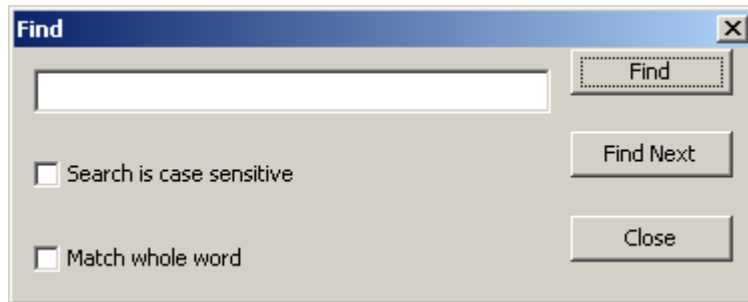
Finding a Group or Object in the MIB Map

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**.

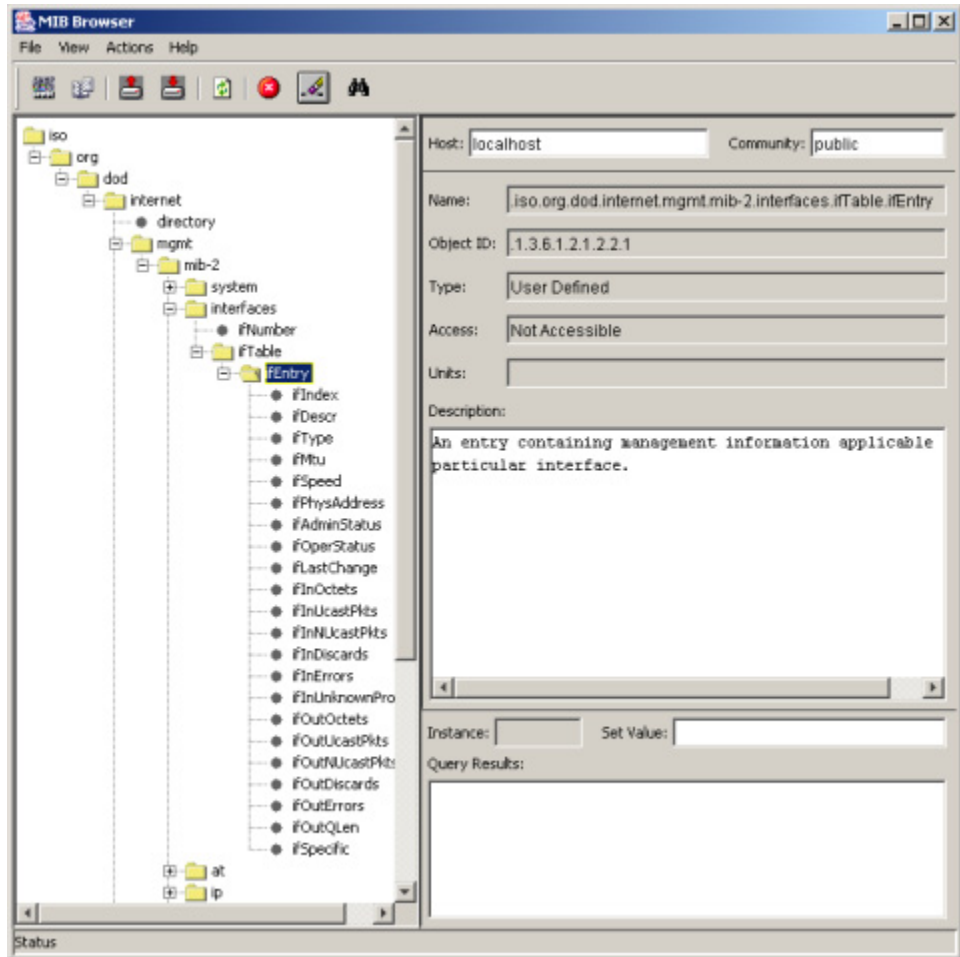
- If you want to search for only the name that is in the box, click on **Match whole word**; otherwise, the default search compares your entry to any portion of the MIB names.
- 4 Click **Find**.
 - 5 If you want to locate another instance of your entry, click **Find Next**.
 - 6 When you are finished 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.

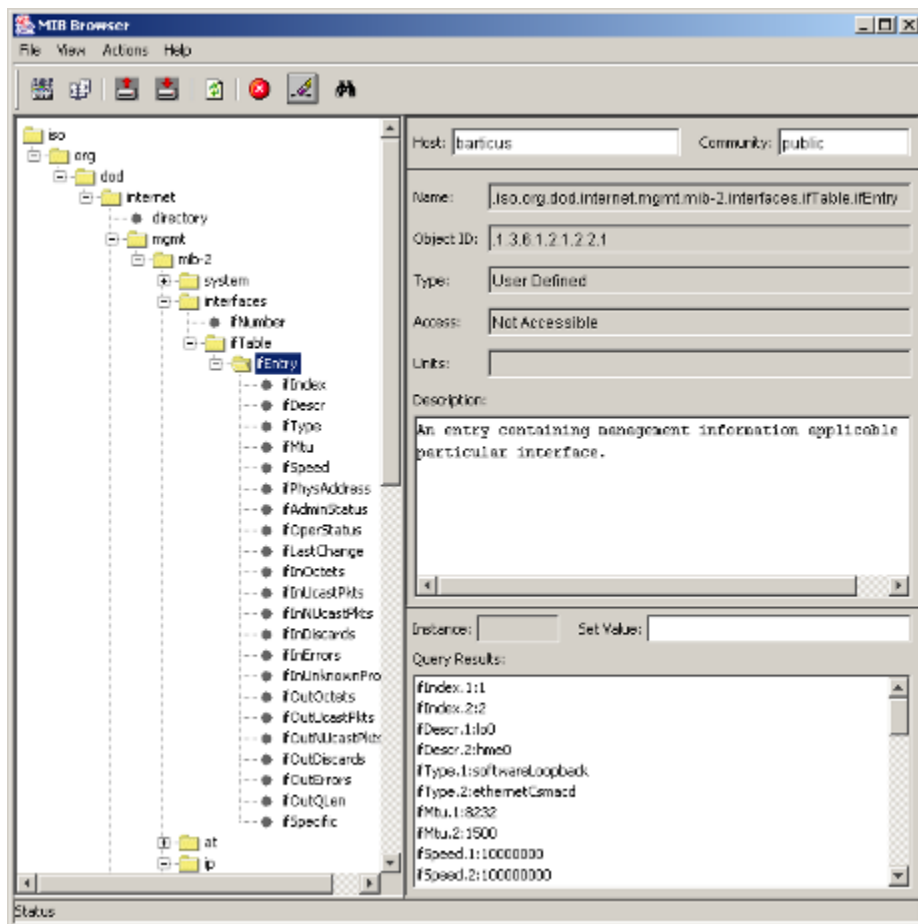
- 1 Select a MIB group or object by clicking the name in the MIB map.

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 Type 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 PI 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 654.

To create the collection table, you use the **Collection** option to select the objects to collect. When you create a collection table from the PI 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 PI. 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, PI 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
---------------------------	--------------------------

You define the pairing of high and low counters using the Define the Hi-Lo Counters Pairs window, which is available when you use the New Collection wizard. For details about creating a new collection, see [Creating a Collection Table](#) on page 656. [step 7](#) on page 661 of this procedure describes how to select pairs of counters to define 64-bit counters.

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. You must allow for 13 default columns 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 Software sales representative to obtain the *HP 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 **Create Collection...**

(A login dialog may appear at this point. Type the login information to continue.)

The Define Collection window opens.

Collection Definition Wizard

Define Collection

Enter a collection name that will be used to create the set of tables in the database. The name must be 30 characters or less. The collection category can be selected from the existing list or you can enter a new name to be used for grouping your collections. You may optionally enter a description up to 255 characters.

Collection Name:

Collection Category: default

Collection Description:

< Back Next > Cancel

- 2 Type the collection name, collection category, and description for the collection table.
 - α In the Collection Name box, type an alias name for the data table; this name is stored in the PI data repository.

➤ The wizard will assign a name to the collection based on the alias name you provide.

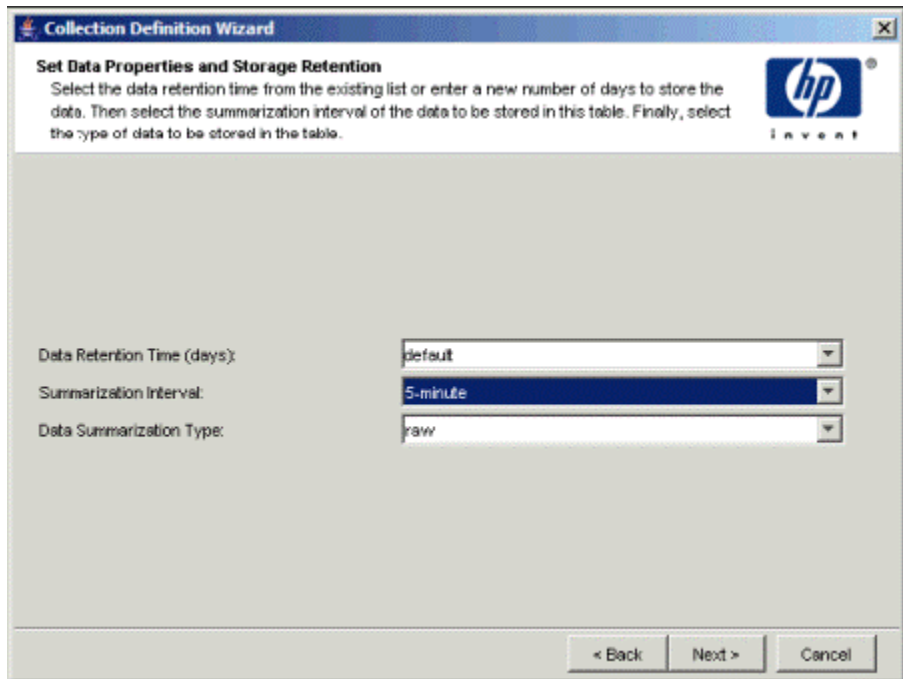
The name must be unique and can be up to 255 characters. Valid entries include uppercase and lowercase alphabetic characters, all numeric digits, and the special characters hyphen (-) and underscore (_).

- b In the Collection Table Category box, type or select the category for the collection table.

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.

- c In the Collection Description box, type a description for the table. A description is optional. If you specify a description, it cannot exceed 255 characters.

- 3 Click **Next**. The Set Data Properties and Storage Retention window opens.



- 4 Select the data retention time, summarization interval, and summarization type of the table.
 - a In the Data Retention Time box, type or select the number of days you want to keep the data in the table.
 - b In the Summarization Interval box, select a time characteristic for the data in the table. For example, a 5 minute table contains data for a 5 minute interval. Data collected at 5 minute intervals is called rate data.

- c In the Data Type box, select a table type.

The table type specifies the class of data contained in the table. The table type options are as follows:

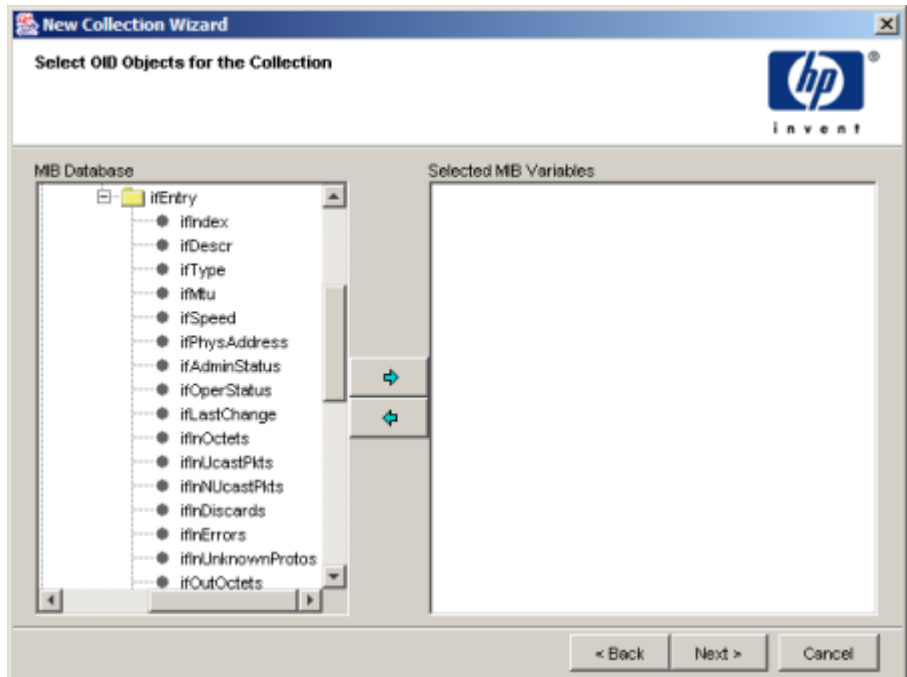
- **rate**

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.

- 5 Click **Next**. The **Select OID Objects for the Collection** dialog box opens.

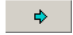


- 6 Select the MIB objects.

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. This means 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.

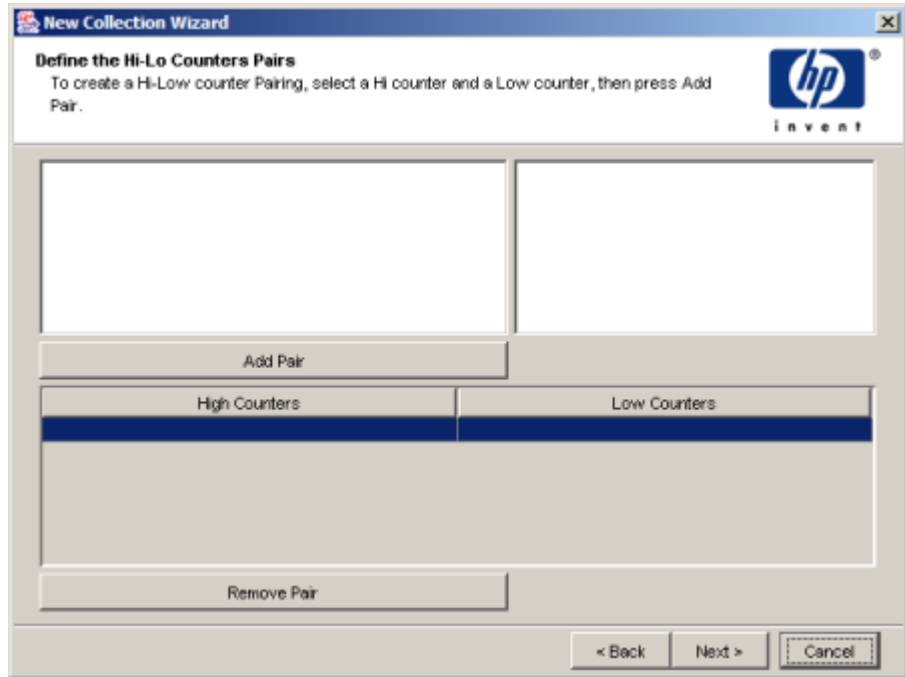


If you select objects that are not valid as columns, such as a folder, you will get an error message with the title, **Unable to add new column**, and it will list the objects; these objects will not appear in the **Selected MIB variables** list.



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 window opens, which shows all the selected counters in each column.



At this point, you can create an artificial 64-bit counter by ANDing two 32-bit counters 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 **Remove Pair** to remove it from the list.

- 8 Click **Next**. The Enter a MIB Index value window opens.

New Collection Wizard

Enter a MIB Index value

Enter a MIB index value if you wish to. A MIB index consists of any integer separated by periods. The value for each integer ranges from -255 to 255. For example: 1.2.-3.-4 A MIB index is not required to continue.

hp
invent

MIB Index:

< Back Next > Cancel

- ▶ The MIB index applies to tabled entries only. Non-tabled entries will have an index of zero.

- 9 Type the MIB index format; see your MIB to determine this.

For example, a MIB index format of 1.1 would indicate two indexes, each one an integer in length.

Supplying the MIB index format is optional; however, consider the following before you decide whether or not to provide one:

If do not supply the MIB index, the following occurs:

- Non-parsed index will be inserted into `dsi_table_key`.
- If there are multiple index values, they will be concatenated and decimal delimited (for example, 1.2).

If you supply the MIB index, the following occurs:

- **mw_collect** will insert each component associated with the index into the property table column(s) that you create.
- After `mw_collect` inserts the components, it will take each component and underscore delimit it and insert it into the table key (for example, 1_2).

Example: Determining the Format of a MIB Index

The following example shows how to determine the MIB index format for the `frCircuitEntry` object in the `frCircuitTable` group in the `rfc2115-FRAME-RELAY-DTE` MIB.

- Open the MIB file. Locate the entry and note that its index is frCircuitIfIndex, frCircuitDlci, which is shown in the following figure:

```

frCircuitEntry OBJECT-TYPE
SYNTAX FrCircuitEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "The information regarding a single Data Link
    Connection. Discontinuities in the counters contained
    in this table are indicated by the value in
    frCircuitCreationTime."
INDEX ( frCircuitIfIndex, frCircuitDlci )
 ::= { frCircuitTable 1 }

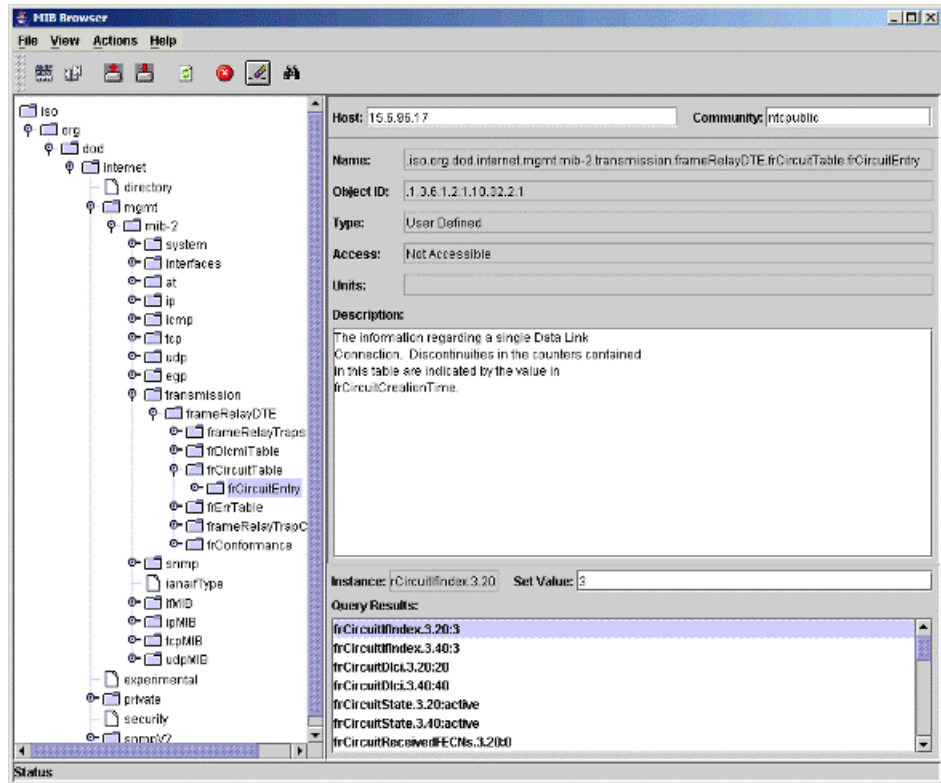
FrCircuitEntry ::=
SEQUENCE (
    frCircuitIfIndex          InterfaceIndex,
    frCircuitDlci             DLCI,
    frCircuitState            INTEGER,
    frCircuitReceivedFECNs    Counter32,
    frCircuitReceivedBECNs    Counter32,
    frCircuitSentFrames       Counter32,
    frCircuitSentOctets       Counter32,
    frCircuitReceivedFrames   Counter32,
    frCircuitReceivedOctets   Counter32,
    frCircuitCreationTime     TimeStamp,

```

The index comprises one integer value (frCircuitIfIndex) followed by another integer value (frCircuitDlci).

When viewed with the MIB Browser, this indicates that the index will be formatted using an integer value, separated by a period, followed by a second integer value (for example, *n.n*).

- b Open the MIB using MIB Browser and navigate to the frCircuitEntry object in the frCircuitTable group, which is shown in the following figure:



When you select the object and perform an SNMP Get request, the object is polled; the Query Results box shows the value of that object for all of its instances in that table. For example, for the frCircuitIfIndex instance, the MIB index format is 1.1 because 3 and 20 are each single integers. You would type 1.1 in the MIB index box in [step 9](#) on page 663 for the frCircuitEntry object.

If you supplied the MIB index format in [step 9](#) on page 663, you will need to do the following:

- Create the new property table columns; it is suggested that you give the tables the same names as the index components ([step 10](#)).
- Make sure you set the correct datatypes and, if applicable, the length ([step 11](#)).

- Specify each new column as an Object-By and Collection-By variable ([step 11](#)).
- After you specify the columns, click to clear the Object-By and Collection-By options for `dsi_table_key` ([step 11](#)).



If you do not perform the steps above, your collection definition will not work.

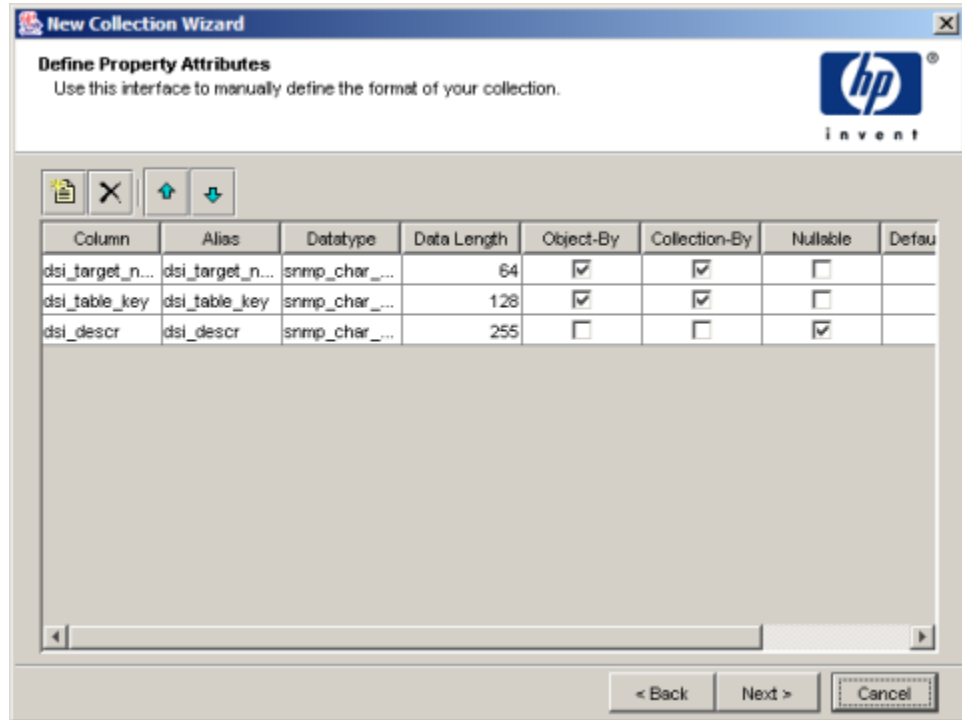
- 10 Click **Next**. The **Associate with a Property Table** dialog box opens.

The screenshot shows a dialog box titled "New Collection Wizard" with the HP logo. The main heading is "Associate with a Property Table" with the instruction "Enter a name for this property table that describes the data to be". There are two radio button options: "Property Table" (which is selected) and "Create a new property table.". Next to the "Property Table" option is a dropdown menu currently displaying "ksi_user". Below the "Create a new property table." option is an empty text input field. At the bottom right of the dialog are three buttons: "< Back", "Next >", and "Cancel".

- 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 444.


- 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:

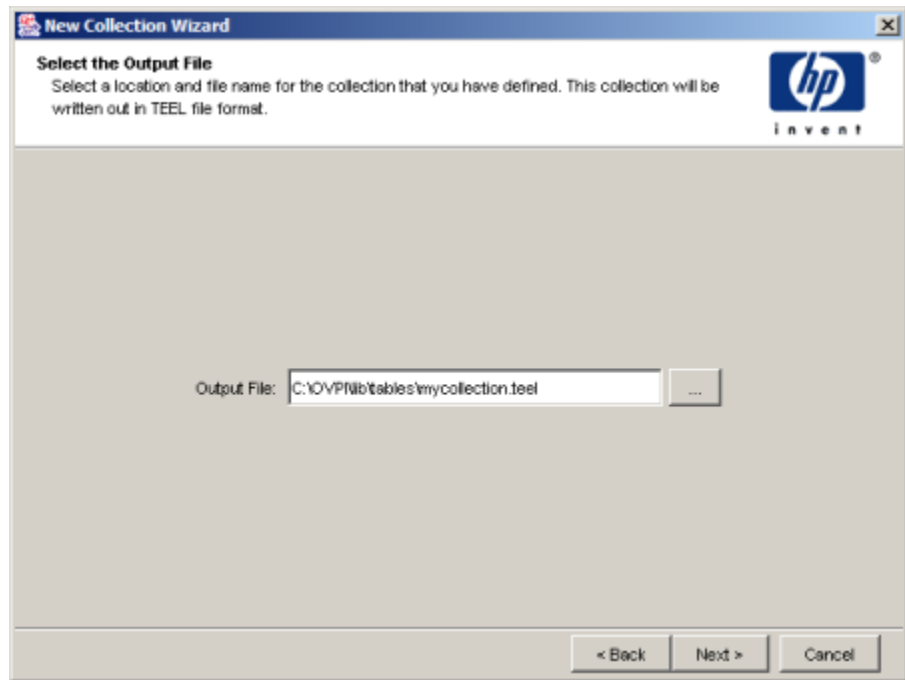
To move the column up, click the **Move Up** button, .

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 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.

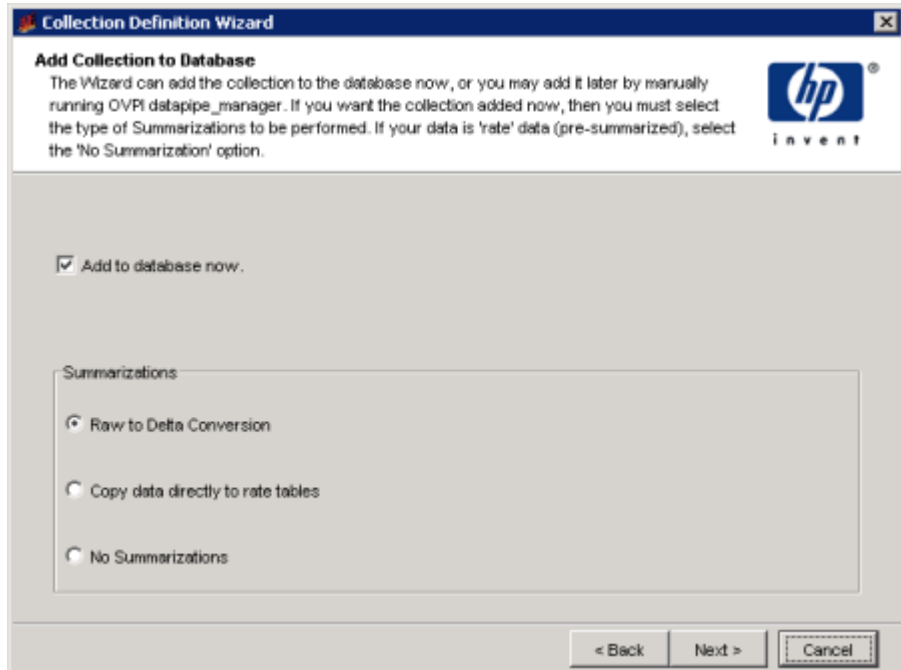


- 13 Verify that the name and location of the output file are correct.

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 type 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 670.

See the *HP Performance Insight TEEL Reference Guide* for additional information about the options for this command.

- 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](#).

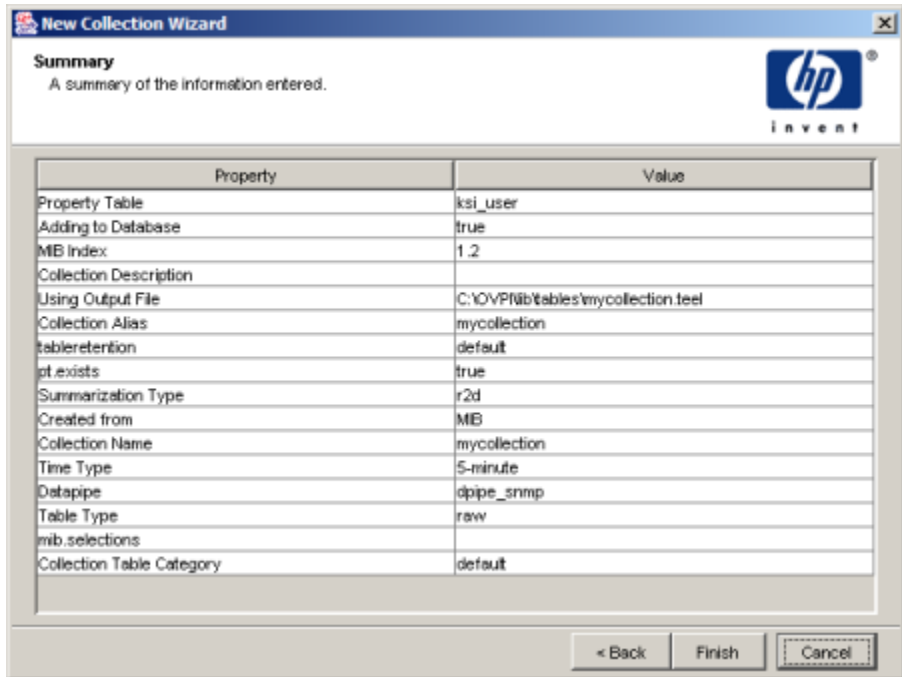
— **Copy data directly to rate tables**

The system will copy the data directly into the rate tables.

— **No Summarizations**

The system will not generate any additional database objects.

17 Click **Next**. The **Summary** dialog box opens.



18 Do the following:

- Review the summary.
- If necessary, click **Back** to return to a previous page to make a change.
- Click **Finish**. The TEEL file will be in the designated location.

19 Tuning the System

This chapter describes things you can do to improve the performance of your PI system. Many factors can positively or negatively affect PI's performance. These factors can be grouped into the following areas:

- Hardware
- Operating system configuration parameters
- Sybase configuration parameters
- PI 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.

This chapter does not describe Oracle configuration parameters. See Oracle documentation for Oracle specific tuning.

Hardware Considerations

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 PI processing.

Use the fastest available processor for the selected platform.

Number of Processors

PI 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 data cache and procedure cache buffers. The more available disk cache the less I/O a system does, thereby improving performance.

Table 18 Memory Guidelines

System Memory (per Processor)	Application
1 GB	Small installations
2 GB	Medium installations
More than 2 GB	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 two to three times the amount of physical memory in the machine.

Table 19 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 19 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.

Use hard drives that are 15,000 RPM.

Number of Physical Disk Devices

In general, the number of 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 medium installations where hardware budget is not a limiting factor.

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.

Configuring Operating System Parameters

Depending on the type of operating system that you are using, you may need to adjust various configuration parameters.

- 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 PI installation program provided defaults for several Oracle database parameters. (See Chapter 3 of the *HP Performance Insight Installation and Upgrade Guide for Oracle databases* for these parameters.) If you experience problems, consult your Oracle database administrator.

Configuring Sybase Parameters

This section 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:

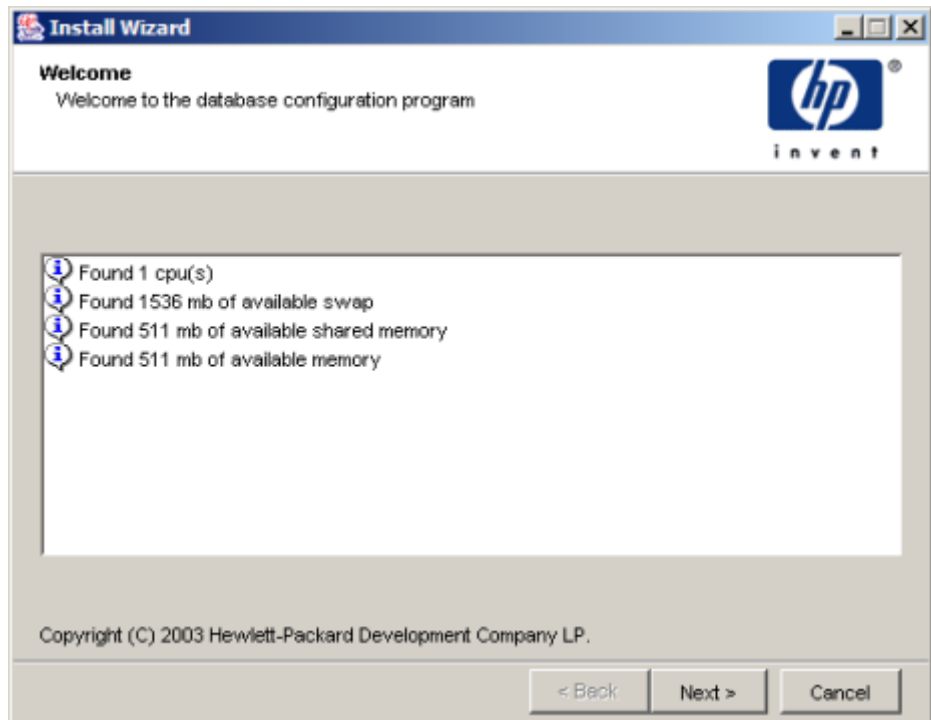
- Command line

Open a shell window (UNIX) *or* a Command Prompt window (Windows), and type **dbtuner**.

- Management Console

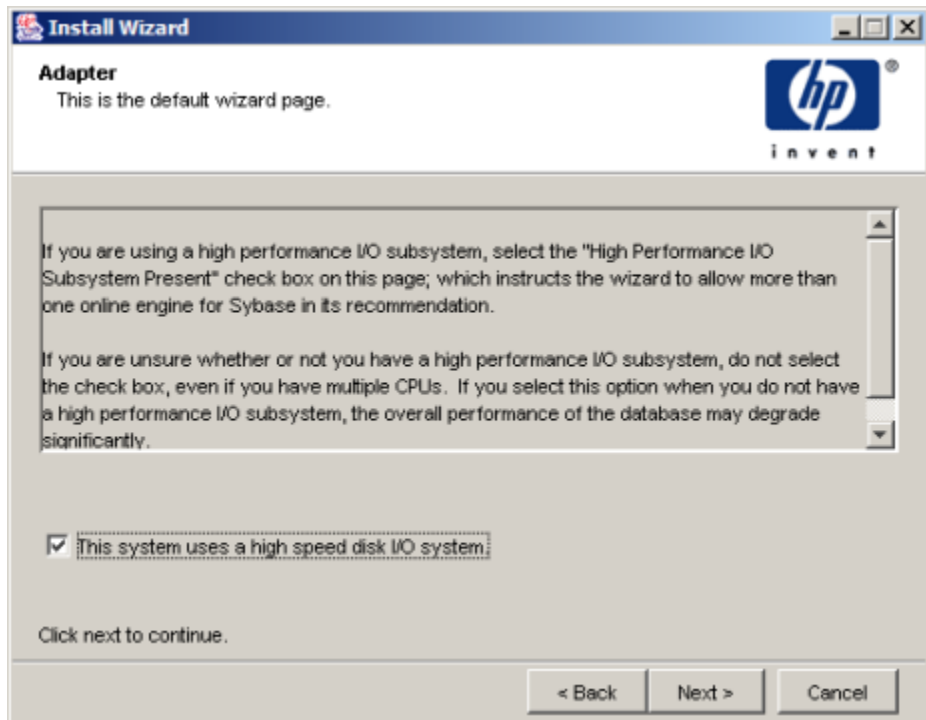
Start the Management Console. (See [Starting the Management Console](#) on page 118 for details.) Log on to the Management Console using an account with administrator privileges, and then choose **Database Tuner** from the Tools menu.

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.

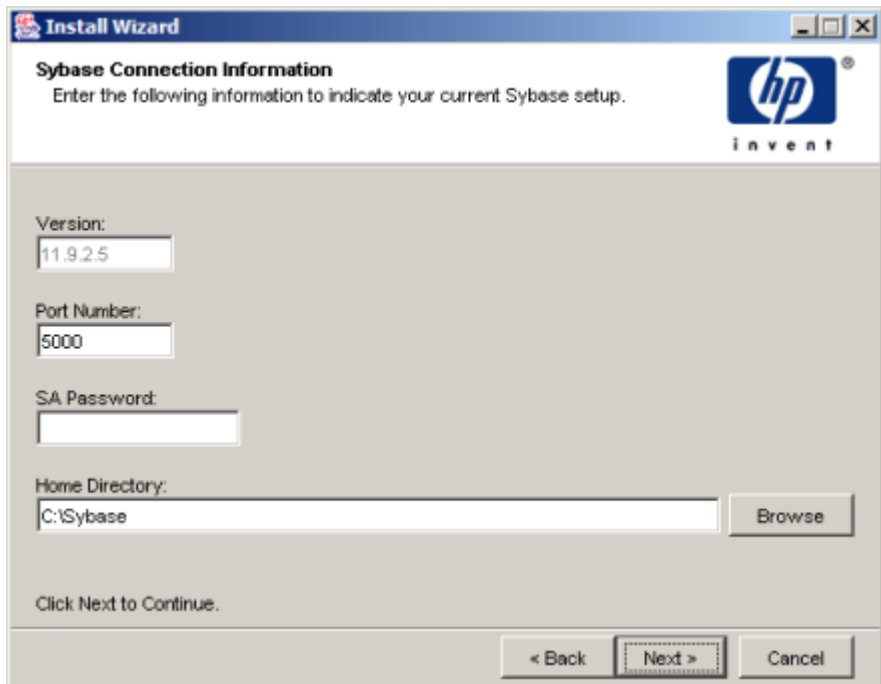


- 4 Do one of the following:
 - If you are using a high-performance I/O subsystem, such as RAID, select the **This system uses a high speed disk I/O system** check box, and click **Next**.

If you select this check box, it will affect the Max Online Engines parameter. See [Max Online Engines](#) on page 686 for details.

- If you are not using a high-performance I/O subsystem, click **Next**.

The Information window displays information about your Sybase setup.



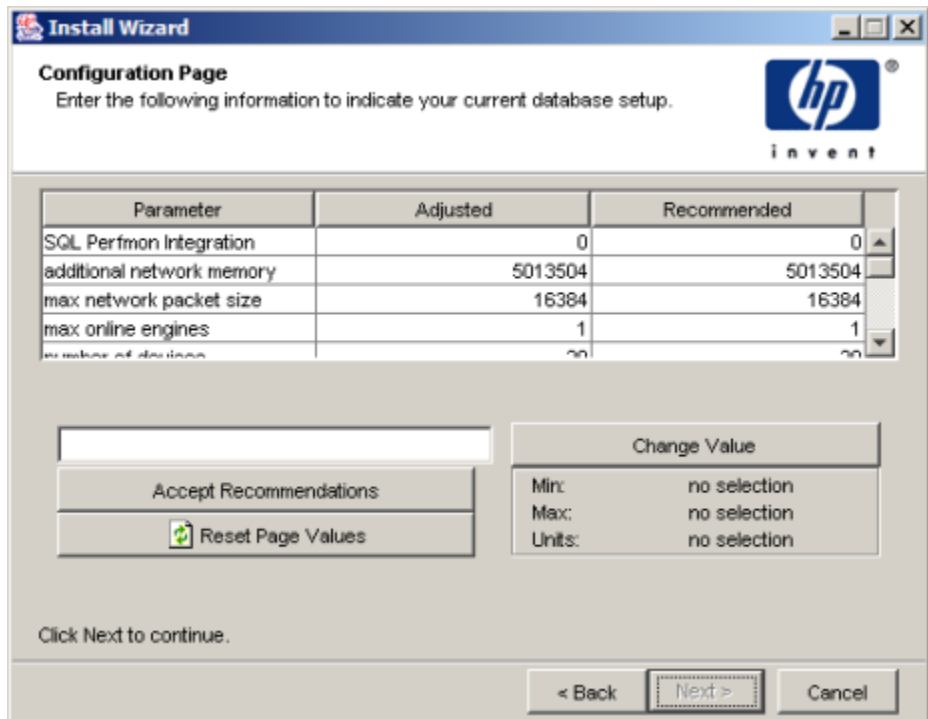
The screenshot shows a window titled "Install Wizard" with a sub-header "Sybase Connection Information". Below the sub-header is the instruction "Enter the following information to indicate your current Sybase setup." and the HP InvenT logo. The form contains the following fields and controls:

- Version:** A text box containing "11.9.2.5".
- Port Number:** A text box containing "5000".
- SA Password:** An empty text box.
- Home Directory:** A text box containing "C:\Sybase" and a "Browse" button to its right.

At the bottom of the window, there is a prompt "Click Next to Continue." and three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel".

- 5 Review the information, make any changes, and then click **Next**.

The Configuration window opens.



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 684. 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.

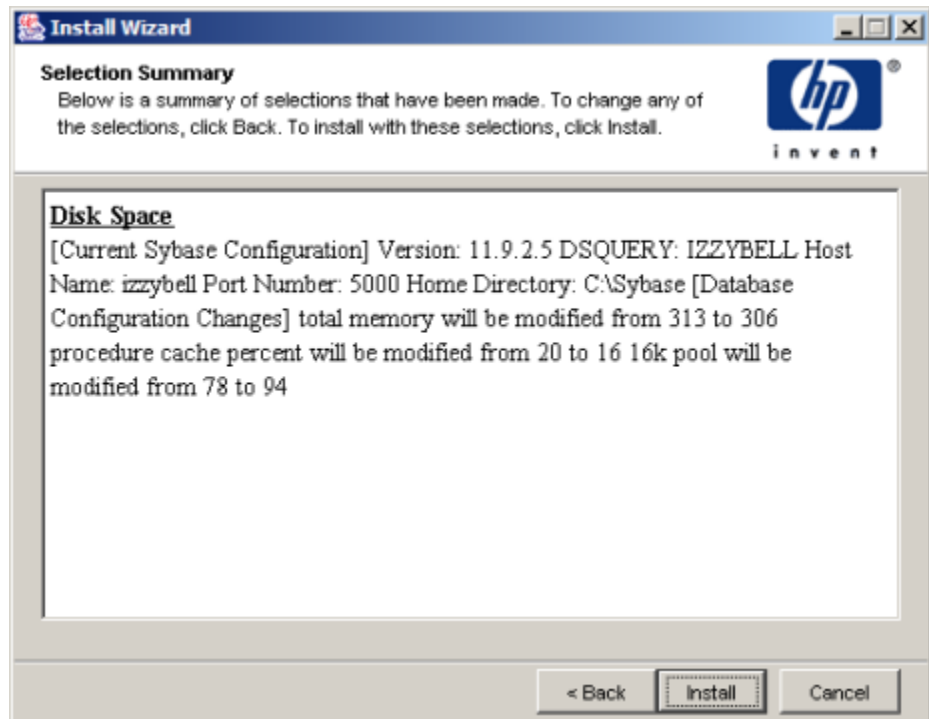
If needed, click **Reset Page Values** to reset parameters to their current values.

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 opens.



- 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 682 using the Sybase Configuration wizard. See the *HP Performance Insight Installation and Upgrade Guide for Sybase databases* 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. However, for optimum performance, use a minimum of 20,000 locks. Implementations that want to run a large number of child collectors or parallel instances of `trendcopy` 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, PI sets the value to 3000 to prevent most related errors.

Keep 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, PI sets the value to 3000 to prevent most related errors.

Keep 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 see 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.

Usually 100 user connections 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 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 or by using `sar` or `vmstat` on the various UNIX ports.

4K Pool and 16K Pool

The **4K pool** and **16K pool** see 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 PI database. Use the information in [Modifying the Default Data Cache](#) on page 692 and [Creating a Dedicated Log Cache](#) on page 692 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 PI 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 PI 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 PI 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 PI transaction log I/O size. If you are using a common log cache, you do not need to modify the PI transaction log I/O size.

Determine the PI 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 PI 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 PI. 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 PI 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 PI 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 PI 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 690 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 type the quotes.

For example, to bind the property tables **K_IFEntry_Disc**, **KIRDevPorts**, and **K_frswitch_port** to the cache named keycache, type 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.

20 Interpreting Log Files

HP Performance Insight's (PI) 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.

HP Performance Insight Log Files

PI 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 PI 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 PI.
- `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 PI'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 546 for details about viewing the `website.log` file.

You can find all but one of PI's log files in the `log` folder of the PI 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 707 and [The metrics.log File](#) on page 710, respectively.

Viewing Log Files

To view log files:

- 1 Navigate to the `log` folder in the PI 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 707 and [The metrics.log File](#) on page 710, 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.
- nnn 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 PI 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*.

See the *HP 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 PI 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 712, [bcp_gateway Messages](#) on page 713, or [trendcopy messages](#) on page 714, 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. 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, see the *HP 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 the transaction size, 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, 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.

A NNM-PI Integration

The NNM and PI Integration Module creates tight linkages between HP Network Node Manager (NNM) and HP Performance Insight (PI). By joining fault management with performance management, the Integration Module enhances problem diagnostic capabilities.

Features and Benefits

The following list outlines the features of the NNM and PI Integration Module and its benefits to you:

- It provides additional performance data from NNM, which contributes to faster and easier resolution of network-based service level problems.
- It shares and synchronizes detailed topology information between NNM and PI databases to better enable NNM and PI to monitor and manage your environment.
- It can forward PI threshold traps to a specified NNM management station (or set of NNM management stations).
- It enables you to launch PI reports directly from an NNM map or the NNM alarm browser. Reports display information pertinent to the node or alarm from which the action is invoked.
- It can integrate other NNM Smart Plug-in and HP Performance Insight products, such as the NNM Event Report Pack, to further enhance the management and monitoring of networks.

For detailed information on the integration of NNM with PI, see the *HP Network Node Manager and Performance Insight Integration User Guide*.

Index

Numerics

- 14K Pool parameter, 690
- 16K Pool parameter, 690
- 32-bit counter, 654
- 64-bit counter, 654, 655, 661
- 64-bit counter, higher precision values, 654

Symbols

- .ini files, saving, 129

A

- Actions menu
 - MIB Browser, 639
- adding
 - nodes, 436
 - property table descriptions, 410
- additional network memory parameter (Sybase), 685
- Administration link, 514, 515
 - Diagnostics folder, 518
 - Java Settings folder, 519
 - Log Files folder, 518
 - Report Settings folder, 520
 - Services Configuration folder, 516
 - User Accounts folder, 515
- administrator, resetting privileges, 569
- advantages of setting up LIR, 611

- Aggregation type
 - , 612
- aliases, changing, 559
- audit.log file, 702
- authentication
 - changing settings, 568
 - setting method, 111, 568

B

- backup script
 - creating, 99
 - running, 101
- backup server, 87
 - dumps and loads, 88
 - shutting down, 97
- backup server, creating
 - on UNIX systems, 88
 - on Windows NT systems, 91
- binding property tables, 697, 698
 - verifying, 699
- builder.log file, 702

C

- cache
 - data pools, 690
 - default data, 690
 - named, creating, 692, 697

- Catalog link, 514, 522
 - Deploy URL page, 522
 - System folder, 523
 - Users folder, 523
- Catalog Manager
 - changing
 - report properties, 477
 - URL properties, 477
 - copying directories, 475
 - copying reports, 475
 - creating new directories, 473
 - Deployed Items catalog, description of, 461
 - deploying URLs, 472
 - description of, 461
 - Links catalog, description of, 461
 - moving directories, 473
 - moving reports, 473
 - report links, 461
 - starting, 463
 - undeploying reports, 469
- Catalog Manger
 - deploying reports, 465
- categories
 - grouping by, 332
- change server
 - menu option, 121
- Changing, 219
- changing
 - database servers, 128
 - table column size, 394
- checklist, post-relocation, 107
- collectable data table
 - definition, 440
- collection table
 - common index, 656
 - definitions of, 654
 - maximum columns, 440, 656
 - maximum width, 440, 656
 - required tables, 654
- Collector
 - EE Collector, 293
 - MW Collector, 293
 - SR Collector, 293
- Collector for Performance Graphs, 293
- column count, 440, 656, 669
- column size, changing, 170
- common index, 656
- community string profiles
 - creating, 280
 - deleting, 281
 - editing, 280
 - managing, 279
- community strings
 - changing values for, 648
- Configuring LIR, 617
- connections, maximum number of, 688
- console.log file, 702
- constraints
 - creating, 432 to 434
 - definition of, 432
 - modifying, 435
- Copy Policies, 622
 - Implementation Examples, 622
- Copy Policy Manager, 623
- counter
 - 64-bit, 654
- CPU grace time parameter, 696

- create
 - polling groups
 - rule-based, 267
- creating
 - community string profiles, 280
 - constraints, 432 to 434
 - custom constraints, 433
 - derived groups, 339
 - enumerated list groups, 337
 - folder groups, 344
 - managed objects, 319
 - nodes, 272
 - polling groups, 261
 - nodes in same view, 264
 - nodes of same type, 263
 - single node, 262
 - specific instances, 267
 - type and view, 265
 - property groups, 342
 - SNMP profiles, 283
 - views, 552
- creating data tables, 440

- CSV file
 - create a collection from, 173

D

- database
 - backing up with script, 98
 - dump database command, 94
 - dumping, 94 to 97
 - examples of, 95 to 97
 - forms, used to modify, 318
 - loading, 97
 - restoring, 97
 - retention days, 396, 397, 604
 - sizing guidelines, 381
 - tuning
 - See performance*

- database servers
 - changing, 128

- Database Tuner
 - configuring Sybase, 678 to 684
 - See performance*

- Database Tuner (Sybase), 687
 - additional network memory parameter, 685
 - max network packet Size parameter, 685
 - max online engines parameter, 686
 - number of devices parameter, 686
 - number of locks parameter, 687
 - number of open indexes parameter, 687
 - number of user connections parameter, 688
 - procedure cache percent parameter, 688
 - recovery interval in minutes parameter, 689
 - SQL perfmon integration parameter, 685
 - total memory parameter, 689

- data cache pools, configuring, 690

- data collection
 - defined, 35

- datapipe installations
 - creating, 286
 - deleting, 287
 - editing, 287
 - managing, 285

- Datapipes
 - Property Import, 37

- datapipes, 135

- data processing
 - defined, 38
 - described, 38

- data reporting
 - defined, 40

- dataset report, definition of, 40

- Data Table, 293

- data table
 - collectable, definition, 440

- data tables
 - adding a description, 401, 410, 600
 - creating, 440
 - defined, 38
 - displaying data in, 414
 - retention days, 396, 397, 604
 - setting defaults, 457
- data tables, viewing
 - by category, 603
 - by type, 603
- dbmaint.log file, 702
- debugging
 - enabling and disabling
 - Web Access Server, 567
- deleting
 - community string profiles, 281
 - datapipe installations, 287
 - nodes, 274
 - polling policies, 259
 - SNMP profiles, 284
 - views, 554
- delta data, defined, 38
- deploying reports, 465
 - to multiple Web Access Servers, 142
- Deployment Wizard, 465 to 471
 - undeploying reports, 469
- Deploy URL page, 522
- derived groups
 - creating, 339
- devices, dumping
 - to disk, 94
 - to tape, 94
- Diagnostics page, described, 518
- digital certificates
 - CSR, generating, 573
 - importing, 573
- directed instance polling, 166

- discovering, nodes
 - using SNMP Discovery wizard, 152
- disk drives
 - effect on performance, 676, 677
- displaying
 - online documentation, 131
- display options
 - Group Manager, 353
- dpipe_db, 88
- dumping
 - database, 94 to 97

E

- editing
 - community string profiles, 280
 - datapipe installations, 287
 - groups, 350
 - polling policies, 257
 - SNMP profiles, 284
- editing polling groups, 270
- editing views, 554
- Edit menu
 - Management Console, 121
- ee_collect, 165
- enumerated list groups
 - creating, 337
- execution frequency, defined, 458
- exporting
 - groups, 348
 - nodes, 317
 - table data, 419

F

- File menu
 - Management Console, 121

- finding
 - groups, 336
 - nodes, 272
 - objects, 308
 - table alias names, 395
- flat files
 - defined, 35
 - obtaining data from, 37
- folder groups
 - creating, 344
- Folder Uniqueness value, 309
- foreign key, 412
- forms
 - using in Object Manager, 318

G

- groupimport utility, 525
- Group Manager, 329
 - deleting groups, 352
 - display options, 353
 - grouping methods, 331
 - group types, 330
 - managing polling policies, 356
 - starting, 333
 - viewing logs, 352

- groups
 - deleting, 352
 - derived
 - creating, 339
 - defined, 330
 - enumerated list
 - creating, 337
 - defined, 330
 - exporting, 348
 - finding, 336
 - folder
 - creating, 344
 - defined, 330
 - importing, 347
 - maintaining, 533
 - managed objects
 - defined, 329
 - editing, 350
 - types, 330
 - using in Report Builder, 361
 - viewing contents of, 349
 - property
 - creating, 342
 - defined, 330
 - user accounts
 - creating, 534
 - defined, 61, 533
 - deleting, 537
 - filters, 63, 85, 536, 538, 539
 - modifying, 84, 537

Groups folder, 516

H

- hardware effects on performance, 674
- headless support
 - disabling, 567
 - enabling, 567
- Help menu
 - Management Console, 123
 - MIB Browser, 639

High Part, defined, 655

HTTPS, 222

hysteresis units
 defined, 458
 setting, 458

I

importing

 groups, 347
 nodes, 275, 316
 users and groups, 525

install.log file, 702

installing packages, 139

Install Updates, 613

interfaces file, 87

intervals, polling, 254

Intimate Shared Memory (ISM), 676

J

Java settings

 changing, 548
 folder, 519

K

keymap table, 410

L

link bar

 changing aliases, 559
 changing options, 556

links

 report to application, 491

LIR

 advantages of setting up LIR, 611
 Implementation Examples, 613
 Mappings, 612

LIR Mappings and Copy Policies, 612

locking

 user accounts, 530

log

See transaction log

log files, 701 to 715

 audit.log, 702
 builder.log, 702
 changing settings, 544
 console.log, 702
 dbmaint.log, 702
 examples of, 706
 format of, 703 to 706
 install.log, 702
 location of, 702
 metrics.log, 702
 names of, 702
 ovpi_uninstall.log, 702
 report.log, 702
 trend.log, 81, 702
 viewer.log, 702
 viewing, 546, 703
 website.log, 702

Log Files folder, 518

logging, enabling and disabling

 Web Access Server, 567

logging on

 Management Console, 119

log information

 viewing, 352

Low Part, defined, 655

M

Mail Server page, described, 517

Manage database roles, 616

- managed objects
 - creating, 319
 - defined, 329
 - finding, 308
 - modifying data in database, 318
 - modifying Objects tree, 309
 - reports associated with, 306
 - viewing, 304
 - views, 306
 - See also* Group Manager
 - See also* Object Manager
 - See also* objects
- Management Console
 - exiting, 121
 - logging on, 119
 - menu bar, 121 to 123
 - Edit menu, 121
 - File menu, 121
 - Help menu, 123
 - Tools menu, 122
 - starting, 118
 - toolbar, 124
- managing
 - community string profiles, 279
 - datapipe installations, 285
 - nodes, 271
 - SNMP profiles, 282
- Managing Certificates, 223
 - Export, 228
 - Import, 227
 - Install, 226
 - Remove, 229
 - Update, 228
- master, 88
- max async I/Os per engine, 696
- max async I/Os per server, 696
- maximum columns, 440, 656
- max network packet Size parameter (Sybase), 685

- max number network listeners, 696
- Max Online Engines parameter, 681
- max online engines parameter (Sybase), 686
- maxusers parameter, 677, 678
- memory
 - guidelines, 675
 - shared, 675
- menu bar
 - for Management Console, 121 to 123
- menu options, MIB Browser, 638
- metrics.log file, 702
- MIB
 - create a collection from, 173
- MIB Browser, 631 to 672
 - Actions menu, 639
 - changing community strings, 648
 - Help menu, 639
 - interface, described, 635
 - menu options, 638
 - starting, 633
- MIBs
 - collection tables, 654
 - described, 631
 - loading, 640
 - managed objects, 654
 - object, inappropriate type, 641
 - setting object values, 653
 - unloading, 646
- modifying
 - constraints, 435
 - property tables, 288
- mw_collect, 165

N

- named cache, creating, 692, 697
- navigation frame options, 561

- node
 - defined, 631
- nodes
 - adding, 436
 - characteristics of, 271
 - creating, 272
 - custom groups, 253
 - deleting, 274
 - editing, 274
 - exporting, 317
 - finding, 272
 - grouping, 331
 - groupings, 253
 - importing, 275, 316
 - managing, 271
- number of devices parameter (Sybase), 686
- number of locks parameter (Sybase), 687
- number of open indexes parameter (Sybase), 687
- number of open objects parameter, 687
- number of open objects parameter (Sybase), 687
- number of user connections parameter (Sybase), 688

O

- Object Manager, 301
 - changing the view, 306
 - starting, 304
 - viewing logs, 308
 - See also* managed objects
- objects
 - See also* managed objects
- online documentation
 - displaying, 131
- operating systems
 - effect on performance, 677

- output
 - viewing
 - Object Manager, 309
- ovpi_uninstall.log file, 702

P

- Package Manager, 133
 - starting
 - from command line, 138
 - from Performance Insight, 137
 - from Windows desktop, 137
 - uninstalling packages, 148
- packages
 - dependencies for, 136
 - installing, 139
 - uninstalling, 148
- paging, caution about, 690
- password
 - setting expiration, 60, 530
 - setting validation rules, 525 to 527
- Password Rules page, described, 516
- performance
 - changing multiple tuning parameters, 678
 - disk drive I/O, 676
 - disk drive quantity, 677
 - effect of retention time, 459
 - factors affecting, 673
 - hardware considerations, 674
 - memory guidelines, 675
 - operating systems, 677
 - parameter definitions, 684
 - shared memory, 675
 - swap file size, 675
- Performance Agents
 - creating collections from, 201 to 222, ?? to 241, ?? to 241, ?? to 251, ?? to 251

- Performance Graphs
 - Process Duration, 294
 - Starting Performance Graphs, 292
 - Viewing Performance Graphs, 293
- performance tuning, 673
 - See* performance
- PI database
 - post-relocation checklist, 107
 - relocating, 106
- polling
 - frequency, 165
 - intervals, 254
 - purpose of, 166
- polling groups
 - creating, 261
 - nodes in same view node, 264
 - nodes of same type, 263
 - rule-based, 267
 - single node, 262
 - specific instances, 267
 - type and view, 265
 - editing, 270
 - managing, 260
- Polling Health
 - Create Polling Health Graphs, 296
- Polling Interval, 293
- polling policies
 - collections, creating, 173
 - from a CSV file, 173
 - from a MIB, 173
 - manually, 173
 - creating, 252
 - deleting, 259
 - editing, 257
 - for groups, 356
 - creating, 356
 - deleting, 359
 - modifying, 358
 - viewing, 171

- Polling Policy Manager
 - changing column size, 170
 - importing nodes, 275
 - intervals, 254
 - refreshing policy list, 170
 - sorting information, 169
 - starting, 168
 - viewing functions, 169
- port number
 - Web Access Server, 566
- procedure cache percent parameter (Sybase), 688
- Process Duration
 - Create, 294
- Process Duration graphs, 294
- property groups
 - creating, 342
- Property Import Datapipe, 37
- Property Table
 - column count, 449, 669
 - default columns, 449, 669
 - maximum columns, 449, 669
- property table
 - maximum columns, 440, 656
- property tables
 - adding descriptions, 410
 - binding to named cache, 697, 698
 - creating, 453
 - displaying details, 409
 - keymap table, 410
 - modifying values, 288
 - verify binding, 699
 - viewing, 398
- proxy servers, setting up, 108

R

- RAID
 - effect on database tuning, 681

- realm name, Web Access Server, 566
- recovery interval in minutes parameter (Sybase), 689
- refreshing
 - polling policy list, 170
- report.log file, 702
- Report Builder
 - creating constraints, 432 to 434
- Report file, definition of, 40
- report link, definition of, 461
- report links
 - creating, 482 to 504
 - exporting, 509
 - importing, 509
 - opening, 504
 - report to application, creating, 491
 - report to report, creating, 482
 - report to URL, creating, 496
- report pack
 - provisioning, 37
- report packs, 134
- reports
 - activating, 563
 - demo packages, 134
 - deploying, 563
 - managed objects, 306
 - types of files, 40
 - undeploying, 469
- reports, deploying, 465
- report settings
 - changing, 549
- Report Settings folder, described, 520
- restoring
 - database, 97
- results frame options, 562
- retention days
 - data tables, 396, 397, 604
- retention time
 - defined, 458
 - for time-period tables, 458
 - setting, 458
- retention time, defined, 38
- rmon_collect, 165
- Row Counts
 - Create Row Counts Graphs, 298

S

- sar, 690
- saving
 - .ini files, 129
 - user preferences, 129
- server
 - backup
 - dumps and loads, 88
 - shutting down, 97
 - starting, 92
 - changing, 121
 - disconnecting from, 121
- services
 - changing settings, 565
 - for UNIX
 - starting, 105
 - for Windows
 - changing status of, 104
 - checking status of, 103
 - starting or stopping, 104
- Services Configuration folder, 516
- Services Diagnostics folder, 518
- setting defaults for, 458
- Settings page, described, 519
- setting up, proxy servers, 108

- shared memory, 675
 - Intimate Shared Memory feature, 676
 - locking, 676
- Site Appearance folder, 556
- site options
 - changing, 556
- SNMP Discovery
 - function of, 151
 - starting, 152
- SNMP profiles
 - creating, 283
 - deleting, 284
 - editing, 284
 - managing, 282
- SNMPv3
 - collecting data from, 64
- Solaris, 676
- sorting
 - database information, 394
- sorting information, 169
- sp_configure, 690
- SQL perfmon integration parameter (Sybase), 685
- SQL Server, backing up, 98
- srvbuild utility, 88
- SSL
 - digital certificates, 570
 - disabling, 573
 - importing, digital certificates, 573
 - keystore password, 572
- SSL Service page, described, 517
- starting
 - Catalog Manager, 463
 - Management Console, 118
 - Package Manager
 - command line for, 138
 - from Performance Insight, 137
 - from Windows desktop, 137
- summarization, table, 451
- swap file size, 675
- swapping, caution about, 690
- Sybase
 - add buffer pools to default data cache, 690
 - backing up SQL Server, 92
 - backup server, 88
 - script, 98
 - shutting down, 97
 - starting, 92
 - default data cache, 690
 - default transaction log, 694
 - devices, dumping, 93
 - parameter definitions, 684
 - RAM, 675
 - tuning
 - See* performance
- sybsystemprocs, 88
- System folder, 523
- System Manager, 575 to 583
 - adding systems, 576
 - changing settings, 570
 - deleting systems, 583
 - modifying system settings, 579
- System Manager page, described, 517

T

- table, 410
 - data, collectable, 440

- table alias names
 - finding, 395
- table data
 - exporting, 419
- Table Manager, 391 to ??
 - changing
 - table column size, 394
 - creating tables, 440
 - data tables
 - viewing by category, 603
 - viewing by type, 603
 - data types, 397, 604
 - finding
 - table alias names, 395
 - navigation, 393
 - panes, 392
 - refreshing the display, 394
 - sorting
 - table information, 394
 - starting, 392
 - summarization types, 451
 - time types, 397, 604
 - viewing
 - property tables, 398
 - table information, 394
- tables
 - column size, changing, 394
 - group member, 410
 - keymap, 410
 - property
 - binding, 697, 698
 - verify binding, 699
 - time-period, 458
- tempdb database
 - sizing guidelines, 381
- time lag
 - defined, 458
 - setting, 458
- time-out value, setting, 568
- time-period tables, 458
 - setting retention time, 458
- title, changing, 556
- tl_groups table, 107
- toolbar
 - Management Console, 124
- Tools menu
 - Management Console, 122
- total memory parameter (Sybase), 689
- traceflag parameters, 694
- transaction log
 - I/O size, 693
 - sizing guidelines, 381
 - Sybase default, 694
- Tree Nodes Per Level value, 309
- trend.log file, 81, 702
- trendtimer
 - stopping
 - on Windows, 105
- tuning
 - See* performance
- tuning PI, 673
- Type Discovery, 159
 - function of, 151

U

- undeploying
 - using the Wizard, 469
- uninstalling packages, 148
- URLs
 - deploying, 472

- user accounts, 525
 - creating, 527
 - deleting, 531
 - e-mailing, 532
 - locking, 530
 - modifying, 531
 - password validation rules, 525 to ??
 - sending e-mail to, 532
 - setting password expiration, 60, 530
- User Accounts folder, 515
- user connections, maximum number of, 688
- userimport utility, 525
- user preferences
 - saving, 129
- Users folder, 523
- Users folder, described, 516
- Using Copy Policy Manager, 623
- utilities
 - groupimport, 525
 - importing users and groups, 525
 - userimport, 525

V

- viewer.log file, 702
- viewing
 - data tables, 603
 - log information, 352
 - property tables, 398
- views
 - creating, 552
 - defined, 521, 552
 - deleting, 554
 - editing, 554
 - node groups, 253
- vmstat, 690

W

- Web Access Server
 - Administration link, 514, 515
 - Catalog link, 514, 522
 - changing link bar aliases, 559
 - changing link bar options, 556
 - changing navigation frame options, 561
 - changing results frame options, 562
 - changing settings, 565
 - changing title, 556
 - deploying reports, 563
 - user accounts, 525
- Web Access Server page, described, 517, 519
- website.log file, 702
- Windows
 - starting or stopping services, 104