

Precise Report Manager

User Guide

Version 10.1

IDERA

Precise Report Manager User Guide

Copyright © 2019 Precise Software Solutions, Inc. All rights reserved.

Document release version 1.0

Precise™, Precise Software™, the Precise™ Logo, Precise i³™, Precise Indepth™, Precise Insight™, Precise Savvy™, SmarTune™, Performance Warehouse™, Application Service Dashboard™, Precise for Storage Tiering™, Precise for Storage Tiering Plus Apps™, Precise for Database & Storage™, Precise for Applications™, Precise for Storage™, Precise Insight Inquire™, Performance Management Database™, and PMDB™ are trademarks or registered trademarks of Precise Software Solutions, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, decompilation and/or reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Precise Software Solutions, Inc. and its licensors, if any.

Certain third-party software may be distributed, embedded, or bundled with this product or recommended for use in connection with its installation and use. Such third-party software is separately licensed by its copyright holder. The list that includes the names of the copyright and license agreements can be found in the Release Notes document.

THE DOCUMENTATION IS PROVIDED "AS-IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PRECISE SOFTWARE SOLUTIONS, INC. SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Global Headquarters

Brookhollow Central III
2950 North Loop Freeway West, Suite 700
Houston, Texas 77092
Phone: +1-877-693-1886
Fax: +1-650-898-1666

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Precise product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Technical Support

For technical assistance, visit our customer portal at <http://www.idera.com/support/productsupport> where you can find an extensive knowledge base, product updates and our online community forums. You can also contact our Customer Support Team via our customer portal, or go to <http://www.idera.com/support/productsupport> for a list of our support access numbers in your country.

Contents

- INTRODUCING REPORT MANAGER8**
- About Report Manager 8**
- Components of Report Manager 8**
 - About the Precise FocalPoint 9
 - About the Report Manager FocalPoint 9
 - About the Performance Management Database..... 9
- How Report Manager works 9**
- What you can do with Report Manager..... 9**
- Where to get more information 9**
- PRECISE REPORT MANAGER BASICS..... 10**
- Introducing the Report Manager user interface 10**
- How the Report Manager tab is structured 10**
 - About the Precise bar 11
 - About the Tier selection bar 12
 - About the View control..... 12
 - About the Report set selector 12
 - About the Report set description..... 12
 - About the Reports table 12
 - About the Data view 13
 - About the Status view..... 14
 - About the Command buttons..... 14
- Tasks common to most tabs..... 14**
 - Exporting to the Precise Custom Portal..... 14
 - Adding, viewing, and deleting Favorites..... 15
 - Printing and saving 16
 - Sending an email message 16
- About role management in Report Manager..... 16**
- Launching Report Manager 17**

ABOUT REPORT MANAGER 23

- Viewing a report..... 23**
- Running a report 23**
- Setting report properties..... 19**
- Setting the schedule of a report generation 19**
- Creating a copy of a selected report 20**
- Creating customized reports..... 20**
 - Defining a query statement..... 20
 - Running and designing a customized report 22
 - Editing an existing customized report 22
- Report Manager settings 23**
 - Setting the General settings 23
 - Setting the Display settings 23
 - Setting the Multi settings 23
 - Setting the Mail settings 24

ABOUT THE REPORT MANAGER REPORTS..... 30

- Introducing Report Manager reports..... 30**
- About the Profile reports set..... 26**
 - About the standard profile report structure 27
 - About standard profile report properties..... 28
 - About the Application Profile report 29
 - About the Response time report 30
 - About the Schema Changes report..... 32
- About the Exception Reports set 33**
 - About the Standard Exception report structure..... 33
 - About Standard Exception report properties 34
 - About the Unused Objects report 36
- About the Availability reports set..... 37**
 - About the Instance-level Availability report structure 38
 - About the Entity-level Availability report structure..... 38
 - About the Availability Overview report structure 39
 - About the Cross-Tiers Availability report structure 39
 - About Availability report properties..... 40
- About the Load Balancing Reports set..... 41**
 - About the Cross-Tiers Load Balancing report structure..... 41
 - About the Tier Load Balancing report structure 42
 - About Load Balancing report properties 42
- About the Capacity Planning reports set 44**

- About the Capacity Planning report structure..... 44
- About Capacity Planning report properties..... 44
- Setting Capacity Planning report properties - time frame tab..... 45
- Setting Capacity Planning report properties - filters tab 45

- ADVANCED FEATURES THAT CAN HELP YOU CUSTOMIZE YOUR REPORTS 46**

- About customizing your reports 46**

- About defining properties for a customized report 46**
 - About creating a customized report..... 46
 - About updating Report Manager tables in Performance Management Database 47
 - Updating the .xml file with the report..... 49

- Filling in gaps in a customized report 49**
 - About creating a customized report..... 50
 - Modifying the .xml file with the report 50

- About importing and exporting customized reports..... 53**

- About copying customized reports..... 53**

- DATABASE MAPPING AND DEFINITION FILES..... 55**

- About database mapping and definition files..... 55**

- About the db.xml file 55**
 - About the Table tag..... 56
 - About the Join tag..... 56
 - About the Related tag 57

- About the def.xml file 57**

- About the tree.xml file 58**
 - About the Group tag 58
 - About the Object tag..... 59
 - About the Objref tag 60
 - About the Alt tag..... 60

- REPORT MANAGER REPORT SETS 71**

- Report Manager Report sets 71**

- Cross-Tiers Reports 63**

- Oracle..... 63**

- J2EE 64**

- Web 65**

- Microsoft .NET 65**

SQL Server	66
SAP	67
Sybase	68
Tuxedo	69
WebSphere MQ.....	69
Oracle Applications	70
OS	70
EMC Storage.....	71
Citrix	71
DB2	71
IMPORTING, EXPORTING AND COPYING YOUR CUSTOMIZED REPORTS.....	73
About importing and exporting customized reports.....	73
About copying customized reports.....	74
GLOSSARY.....	75

Introducing Report Manager

This section includes the following topics:

- [About Report Manager](#)
- [Components of Report Manager](#)
- [How Report Manager works](#)
- [What you can do with Report Manager](#)
- [Where to get more information](#)


About Report Manager

Today more than ever, the success of your company depends on the delivery of information, on time, on reliability, and on reasonable cost. Information is a bottom-line issue, affecting both overall revenue and costs. A performance slowdown, that is, the inability to deliver information on time, is likely to expose you to serious business risks.

Components of Report Manager

Report Manager employs a client/server architecture. The following figure illustrates the typical architecture of Report Manager.

Figure 1-1 Report Manager architecture



Report Manager consists of the following components:

- Precise FocalPoint that communicates with the client's browser.
- Report Manager FocalPoint that is installed on the system server (or a dedicated server).
- Performance Management Database (also known as PMDB).

About the Precise FocalPoint

Precise FocalPoint is the main FocalPoint of all Precise suite. It includes a Web Server that communicates with Internet browsers of Precise users. The exported report files (PDF and Microsoft Excel files) are stored in this Web Server.

About the Report Manager FocalPoint

Report Manager FocalPoint generates reports at predefined schedule and on demand. It consists of the following components:

GUI server	The GUI Server listens to requests from Precise users, such as viewing and generating reports. In case the request is to generate a report, it delivers the request to the Report Generator component.
Scheduler	The Scheduler reads the reports generation schedule. Based on this schedule, it requests the Report Generator component to generate reports.
Report generator	The Report Generator receives report generation requests from the GUI Server and Scheduler components. To generate a report, the Report Generator performs the following: <ul style="list-style-type: none"> • Prepares the appropriate SQL SELECT statements. • Executes the SQL SELECT statements on the Performance Management Database. • Creates the report based on the retrieved information. • Exports the report as a PDF file (for predefined reports) or as a Microsoft Excel file (for customized reports). • Transfers the file to the Precise FocalPoint.

About the Performance Management Database

The Report Manager FocalPoint retrieves the sampled data, which is collected by the various Precise and Insight products, from the Performance Management Database. In addition, it stores all reports configuration data in the Performance Management Database.

How Report Manager works

Report Manager collects and organizes historical information to focus the IT team onto infrastructure hot spots. Report Manager is your performance surge protection against dangerous resource consumption trends. You can use Report Manager to compare period-to-period performance versus a baseline and identify response time problems at-a-glance, before they affect your system performance.

What you can do with Report Manager

Report Manager is a reporting tool that queries the Performance Management Database and displays the results in reports that are generated on a scheduled basis or on demand.

The reports are designed to help you:

- Identify problematic conditions that may have an impact on performance before they become performance problems.
- Track long-term performance, volume trends and patterns, and then plan for future growth accordingly - capacity planning.
- View availability problems over long periods of time at different levels.
- Compare the performance of similar systems, such as different Oracle instances and server machines.
- Correlate between performance metrics of different products, such as between the number of EJB executions and the Lock Wait Time of an Oracle Program.
- Generate customized reports according to your demands, without the need to understand the physical structure of the Performance Management Database.

Where to get more information

More information on Precise, its products, technical notes, and so on, can be found in the Release Notes document for this version.

Precise Report Manager basics

This section includes the following topics:

- [Introducing the Report Manager user interface](#)
- [How the Report Manager tab is structured](#)
- [Tasks common to most tabs](#)
- [About role management in Report Manager](#)
- [Launching Report Manager](#)

Introducing the Report Manager user interface

Report Manager is a Web-enabled application. This means that you do not need to install any special software on your desktop to use Report Manager. Instead, the application opens in a browser.

Because Report Manager is Web-enabled, you can access it from almost any location that is connected to the World Wide Web as long as you can access your Report Manager server.

NOTE Although Report Manager opens in a browser window it is important to understand that it is an application rather than a Web site. To perform activities within Report Manager, use the Report Manager controls, not your browser controls. Do not open information in a New Window, and you should not use the Back or Refresh hot-keys. Use the Precise bar instead.

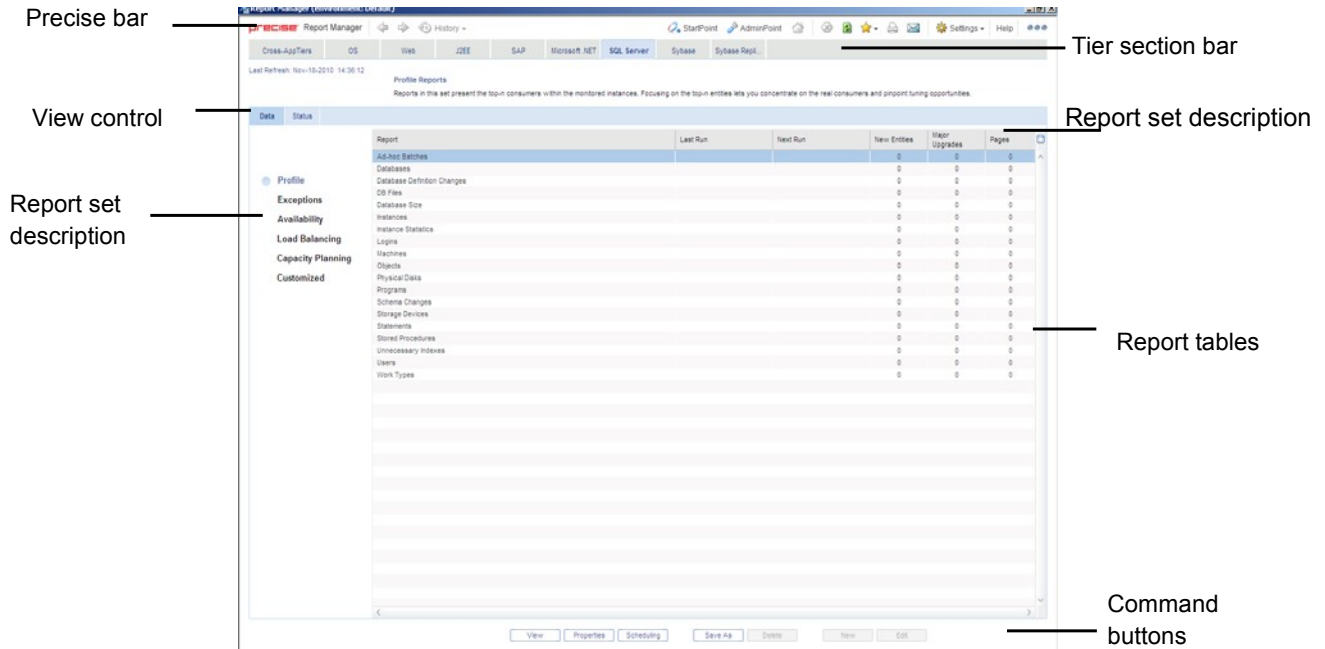
How the Report Manager tab is structured

Report Manager user interface consists of one main screen from which you can access the various Tier tabs. Report Manager main screen includes the following elements:

- [About the Precise bar](#)
- [About the Tier selection bar](#)
- [About the View control](#)
- [About the Report set selector](#)
- [About the Report set description](#)
- [About the Reports table](#)
- [About the Data view](#)
- [About the Status view](#)
- [About the Command buttons](#)

The following figure describes the Report Manager user interface elements.

Figure 2-1 How the Report Manager tab is structured









About the Precise bar

The Precise bar enables you to keep track of where you have been and provides various controls. The following table describes the function of each of the toolbar buttons

Table 2-1 Precise bar functions

Icon	Name	Description
	Backward	During a work session, keeps track of where you have navigated to. The Back button enables you to navigate between previously visited views. The Back control displays your previous view.
	Forward	Enables you to navigate to the next view. This button is only enabled if you clicked Back or if you chose a history option.
	History	The History control displays your navigation paths signifying the various views you have already visited. Choose a navigation path from the pull-down menu to display a previous view. The history does not record changes to the selected view or tab. For example, in the Current tab, if you have an instance selected as the entity in the Main area and you change the association area from Active applications to Applications and then click the Back button, you will not be taken back to see the Active Applications. Instead, you will be taken back to see the previous entity in the Main area, which may be in a different tab.
	StartPoint	Launches Precise StartPoint from your current Precise product. You can launch any Precise product from StartPoint.
	AdminPoint	Launches Precise AdminPoint.
	Home	Navigates to the highest level entity, usually the instance or Tier (all instances). The time frame settings remain the same.
	Stop	Stops a request for information from the server.
	Refresh	Updates the data currently displayed.

Icon	Name	Description
	Favorites	Enables you to add or remove favorites in your Favorites list.
	Print	Opens the Print Option dialog where you can select the area of the screen you want to print.
 Settings ▾	Settings	Opens the General Settings and Time Frame Settings dialog boxes.
 Actions ▾	Actions	Depending on the tab you are viewing, it lets you perform different actions, including editing an existing statement or saving a statement with a new name.
	Send	Opens a new email message in your email program with the link to the current application in context.
Help	Help	Opens the online help in context.
	Communications Indicators	Indicates whether or not the client is communicating with the server. Moving the cursor over this icon displays a ToolTip indicating the GUI server and port number.

About the Tier selection bar

Report Manager collects information and generates reports about the various Tiers in your application. The Tier Selection bar contains a button for each of these Tiers. When you select an Tier, Report Manager displays information relevant for that particular Tier in the tab. You can switch from one Tier to another at any time using these buttons. In addition, the Tier Selection bar contains the Cross-Tiers button that enables you to display high-level view of your application and generate integrated reports of more than one Tier.

About the View control

The View Control enables you to select the view of the Reports table. The following two views are available:

- Data view** Displays the reports list with statistical information about the report, such as the last run date of the report and number of pages in the report. See [About the Data view](#) on page 13.
- Status view** Displays the reports list with status information about the report, such as last run date and error description. See [About the Status view](#) on page 14.

About the Report set selector

Report Manager reports are grouped into six different sets as follows: Profile, Exceptions, Availability, Load Balancing, Capacity Planning, and Customized. When you select a report, Report Manager displays its assigned reports in the Reports Table.

In Cross-Tiers area, the Report Set selector is dimmed because the Cross-Tiers reports are assigned to a single report set.

About the Report set description

A short description of the report set, including few statistical information, which summarize the data available in the reports.

About the Reports table

The Reports Table, which is the main area of the user interface, contains a list of Report Manager reports that are assigned to the selected Tier and report set.

The Reports Table also displays statistical information according to the selected view: Data or Status.

About the Data view

In the Data view, the Reports Table shows a list of reports with statistical information displayed for each report. The statistical information includes columns that are common to all reports and additional columns that are specific to the selected report set.

NOTE The Customized reports set includes only the common columns: Report, Last Run, Next Run, and Pages.

The following table describes the information displayed in all reports.

Table 2-2 Information displayed in the Reports table

Column	Description
Report	Specifies the report name.
Last Run	Specifies the last time and date of generating the report.
Next Run	Specifies the next scheduled time and date for generating the report.
Pages	Specifies the number of pages in the report.

The following table describes the information that is specific to a specific report set.

Table 2-3 Information displayed according to specific report

Column	Description
Profile	<p>Reports in this set present the top <i>n</i> consumers within the monitored instances. Focusing on the top <i>n</i> entities lets you concentrate on the real consumers and pinpoint tuning opportunities. The following information is displayed in the Profile column:</p> <ul style="list-style-type: none"> • New Entities. Specifies the number of new entities in the Profile report. Entities are considered as new entities if they do not have a baseline or they are new in the top <i>n</i> entities list. • Major Upgrades. Specifies the number of entities that were upgraded by more than four places in the report.
Exceptions	<p>This set of reports is intended to help you keep your application tuned. Any substantial deviation from normal behavior is reported. The following information is displayed in the Exception column:</p> <ul style="list-style-type: none"> • Exceptions. Specifies the number of exceptions in the report. • New Entities. Specifies the number of new entities in the report. Entities are considered as new entities if they do not have a baseline. • Deviation (Avg). Specifies the average deviation from the baseline in the report. • Deviation (Max). Specifies the maximal deviation from the baseline in the report.
Availability	<p>This set of reports informs you about failures, unavailability times, and similar events. The following information is displayed in the Availability column:</p> <ul style="list-style-type: none"> • Availability (Avg). Specifies the average availability percentage of the entities in the report. • Availability (Min). Specifies the minimal availability percentage of the entities in the report.
Load Balancing	<p>Reports in this set are intended to help you keep your application balanced. The following information is displayed in the Load Balancing column:</p> <ul style="list-style-type: none"> • Machines. Specifies the total number of machines in the report. • Overloaded. Specifies the number of machines that were overloaded.
Capacity Planning	<p>These reports analyze the performance of the monitored instances to predict future trends and assist in planning. The following information is displayed in the Capacity Planning column:</p> <ul style="list-style-type: none"> • Increase. Specifies the additional consumption (%) that is predicted in the report. • Forecasted Weeks. Specifies the number of weeks included in the forecast.
Customized	<p>This set contains new user-defined reports and copies of reports from other sets. The customized reports include data that is not provided in the pre-defined reports.</p>

About the Status view

In the Status view, the Reports Table displays a list of reports along with their corresponding status information. The following table describes the information displayed in the Reports table in the Status view.

Table 2-4 Information displayed in the Reports table in the Status view

Column	Description
Report	Specifies the report name.
Last Run	Specifies when Report Manager started generating the last report.
Duration	Specifies how long it took Report Manager to generate the last report.
Error	Provides a description of the error that occurred while the report was last being generated.

Clicking Details or double-clicking on the selected report displays the status details of the selected report.

About the Command buttons

The Command buttons enable you to enable various functions for each report, such as opening the properties page or viewing its scheduling parameters. The following table describes the functions that can be carried out for each report.

Table 2-5 Command buttons

Column	Description
View	Displays the content of the selected report. See Viewing a report on page 23.
Properties	Opens the Properties dialog box of the selected report. See Setting report properties on page 19.
Scheduling	Opens the Scheduling dialog box, which enables you to select the schedule for running the selected report. See Setting the schedule of a report generation on page 19.
Save As	Opens the Save As dialog box, which enables to create a copy of the selected report. See Creating a copy of a selected report on page 20.
Delete	Deletes a selected report that is not original of Report Manager. For example, a new customized report or a Saved As report.
New	Creates a new customized report (this option is only available for the Customized reports set or the Cross-Tiers tab). See Creating customized reports on page 20.
Edit	This button enables to edit an existing customized report (available only for customized reports). See Editing an existing customized report on page 22.

Tasks common to most tabs

The following tasks are commonly performed in most tabs:

- [Exporting to the Precise Custom Portal](#)
- [Adding, viewing, and deleting Favorites](#)
- [Printing and saving](#)
- [Sending an email message](#)

Exporting to the Precise Custom Portal

The Export to the Precise Custom Portal Portlet feature enables you to export the view of the chosen table or graph and generate a portlet with that view in the Precise Custom Portal, so that it will provide you with another way of monitoring your application.

Prerequisites

To be able to use this feature, you need to have the following rights in Precise:

- View permissions to all Tiers in the application

If you do not have sufficient rights, you will get an error message when trying to execute this feature.

Exporting the information

You can either export a table view or a graph view.

NOTE The name field has the following restrictions: maximum 100 characters.

To export a table view

1. Click the Column Chooser icon.
2. Select Export to the Precise Custom Portal Portlet.
3. Insert a name in the name field that clearly describes the table view.
4. Click **OK**.

To export a graph view

1. Right-click the graph.
2. Select Export to the Precise Custom Portal Portlet.
3. Insert a name in the name field that clearly describes the graph view.
4. Click **OK**.

Adding, viewing, and deleting Favorites

The Favorites feature enables you to save a specific location in Precise and to retrieve the same location later without having to navigate to it.

About the Favorites feature

The new Favorites feature includes the following option:

- **Relative Time Frame.** Saving relative time frame instead of static date. For example, saving the last 7 days will always display the last 7 days, depending on the day entered.
- **One click to specific location.** Once you open Precise by launching a saved Favorite item, you will not have to enter a login credential nor click the login button.
- **IE Favorites support.** Adding a new Favorite item in Precise will also add it to the IE Favorites menu.
- **Auto Complete.** The Favorites dialog includes a new combo box which supports Auto Complete.
- **Auto Naming.** The Favorites dialog generates item names based on the current location.

UI description

A Favorites menu has been added to the Precise bar in each product including StartPoint. An Add/Delete Favorites option under the Favorites menu allows you to save the current location or delete an existing one.

To add a new Favorite location

1. On the Add/Delete Favorites dialog box, enter the name of the new Favorites entry.
2. Click **Add**. The dialog box is closed and the new Favorite is added to the list.

To view a Favorites location

1. On the Precise bar, click **Favorites**.
2. Select the Favorites location you want to view.

To delete an existing Favorite location

1. On the Add/Delete Favorites dialog box, select the Favorite location to be deleted.
2. Click **Delete**. The dialog box closes and the selected Favorite is deleted from the list.

NOTE The favorite address is displayed in the Address field and cannot be edited.

Printing and saving

The new Print feature enables user-friendly printing and/or saving of most of the screens displayed in Precise products. Emphasis has been made on Activity, Dashboard, and Database layouts since they cover the majority of the screens the user may want to print or save.

About the Print feature

The new Print feature includes the following options:

- Capability to export a page from Precise to an .htm format (Save option). The saved file will include all the information contained the original file but all mouse options will be grayed-out except for the Print option.
- Capability to save the print page to the local disk.
- User-friendly and resizable print functionality.

NOTE The Save option is only possible within the same network of the installation. When used outside the network it may include broken links to images (instead of the image a frame with an X may be displayed).

Recommended print settings

For best fit and presentation we highly recommend you make the following changes to your print settings:

- Configure print options to landscape page layout.
- With the Internet Explorer open, go to **Tools>Internet Options>Advanced** tab, and in the Printing section verify that the Print background colors and images check box is marked.

To print wide tables with many columns, resize table columns to fit the landscape layout. Columns that overflow the landscape width will not be printed.

Additional information

- Note that when saving a file with a graph, the graph image is saved as a separate bmp file in the same folder. If you want to move the print file to a different folder you must also move the bmp file to that folder.
- New files are saved with a default name and path. To prevent inadvertently deleting a previously saved file, enter a different name and/or path for new files.

To print, save, or preview a currently displayed screen

1. Click **Print** on the Precise bar.
2. In the Print Option dialog, select the area(s) of the screen you want to print or save.
3. Click **OK**.
4. On the resizable preview screen, the dashed line on the right indicates the right margin. Change table column widths to ensure the information you want to print or save is within this margin.
5. Click the **Print** or **Save** icon at the top right-hand side of the screen, follow the displayed dialog box instructions and click **Print** or **OK**.

Sending an email message

You can send an email message to one or more recipients from the Precise toolbar. The default subject for the message will be "Link to a Precise application".

The email will include a link to the Precise product in the current context (time frame and selected entries). To send an email message

1. Click the email icon on the Precise toolbar. The default email program opens.
2. Fill in the required fields and click **Send**.

About role management in Report Manager

Use the role MONITOR.VIEW to view reports and MONITOR.FULL to manage reports.

Launching Report Manager

You can launch Report Manager user interface from the Precise StartPoint. The StartPoint is the opening screen, which provides a quick status overview of your application and start points for launching other Precise products.

You can also launch Report Manager for a specific Tier to display data in context with the specified Tier.

Report Manager displays information about the application selected on the StartPoint screen. If you use the Precise product suite to manage the performance of more than one application, verify that you select the required application within StartPoint before starting Report Manager.

NOTE You must have local administrator privileges on the server where the StartPoint is running.

To launch Report Manager from StartPoint, click the Application TPM drop-down menu at the top of StartPoint and select Report Manager. The Report Manager screen opens displaying the Cross-Tiers tab.

To launch Report Manager for a specific Tier (in context), from the StartPoint screen, right-click a specific Tier and from the pop-up menu select Report Manager. The Report Manager screen opens displaying the specified Tier tab.

About Report Manager

The section includes the following topics:

- [Viewing a report](#)
- [Running a report](#)
- [Setting report properties](#)
- [Setting the schedule of a report generation](#)
- [Creating a copy of a selected report](#)
- [Creating customized reports](#)
- [Report Manager settings](#)

Viewing a report

Report Manager displays generated reports in PDF format for pre-defined reports or in Microsoft Excel format for customized reports. You can view reports only in the Data view. The following procedure describes how to view reports that have already been generated at pre-scheduled times. See [Creating customized reports](#) on page 20 and [Running a report](#) on page 23.

To view a report

1. From the Tier Selection bar, select the Tier for which you want to view a report.
2. From the Report Set Selector, select the Report Set that contains the required report. The Reports Table displays the reports of the selected report set.
3. From the Reports Table, select the report you want to view and click **View** (or double-click the report). If the report is not yet generated or does not contain data (the Pages column indicates 0), Report Manager displays a message that allows you to rerun the report.

If the report is already generated and contains data (the Pages column indicates a number greater than 0) the report content is opened.

Running a report

Report Manager generates a report according to the settings of its properties. You can rerun a selected report with different property settings for that particular run. For example, you can rerun the report to include only specific counters and a specific Time Frame. See [Setting report properties](#) on page 19 and [Introducing Report Manager reports](#) on page 30.

To run a report

1. From a report view, click **Rerun**. The Properties for This Run dialog box is displayed. This dialog box contains similar tabs to the tabs of the report Properties dialog box (Description, General, Time Frame, and Filters).
2. In the Properties for This Run dialog box, set the properties according to your preferences.
These settings affect only the current run of the report and do not affect the default properties of the report.
3. Click **OK**. Report Manager generates the report and then displays it.

Setting report properties

Report Manager keeps the properties for each report in the Performance Management Database. You can view and modify these parameters in the Properties dialog box.

The Properties dialog box includes the following tabs:

Description	This tab contains the report description (read only).
General	This tab includes parameters - specific to the selected report - that are neither time frame parameters nor filters. For example, you can specify the counters to include in the report and the number of entities to display.
Time Frame	This tab contains parameters that define the time periods assigned to the various elements of the report.
Filters	This tab includes parameters - specific to the selected report - that enables you to filter the data to include in the report, such as the instance filter and entity filter. A filter can include a list of wildcards, separated by commas, where a wildcard is a string, which may contain asterisks. The format used to filter Report Manager reports by instance is: <code>InstanceName (ServerName)</code> . For example: <code>ORA92 (hp2)</code> .
Email	<p>This tab includes a list of addresses for emailing the selected report. Check the Email box to send copies of the selected report to the listed addresses.</p> <p>Type each address separately into the Address box and click Add to add the address to the Address list. The structure of the address is as follows: <i>employee name@company address</i>.</p> <p>Click Remove to remove an address from the Address list.</p>

To set report properties

1. From the Reports Table, select a report and click **Properties**. The Properties dialog box is displayed.
2. Set the properties in the General, Time Frame, Filters, and Email tabs, according to your preferences.
3. Click **OK**.

The new settings are used the next time the report is generated.

Setting the schedule of a report generation

Report Manager runs reports, by default, on a daily basis at 6:00 AM, following the Performance Management Database daily summary process. You can modify the default schedule of a selected report through the Scheduling dialog box.

The Scheduling dialog box includes the following options:

Once	Run only once in the selected date and time.
Hourly	Run every hour at the specified minute, starting at the specified date.
Daily	Run once a day at the specified time, starting at the specified date.
Weekly	Run once a week at the specified day of the week and time, starting at the specified date.
Monthly	Run once a month at the specified day of the month and time, starting at the specified date.
Never	Never run the report.

To set the schedule of a report generation

1. From the Reports Table, select the report and click **Scheduling**.
2. Set the schedule for the selected report and click **OK**.

Creating a copy of a selected report

Each report is assigned with a set of properties and schedule. If you want to have more than one report of the same set (Profile, Exception, and so on) and the same entity, but with different properties and schedule, you can create a copy of the report, and then modify the properties and schedule of the report copy. See [Setting report properties](#) on page 19 and [Setting the schedule of a report generation](#) on page 19.

To create a copy of a selected report

1. From the Reports Table, select the report and click **Save As**.
2. In the Report Name field, type a new name for the report copy.
3. From the Tier list, select the Tier for which you want to assign the report copy.
 You can assign the report copy either to the Tier of the original report or to the Cross-Tiers.
4. From the Set list, select the report set for which you want to include the report copy.
 If you have selected the Tier of the original report in the previous step, you can include the report copy either to the original report set or to the Customized set.
5. Click **OK**.

Creating customized reports

Report Manager enables advanced users who are familiar with the data logic of the Tier to create customized reports, which include data that is not provided in the pre-defined reports.

You can create a customized report only when the Customized set is selected, or when the Cross-Tiers tab is selected.

To create customized reports

1. On the Command buttons, click **New**.
2. Define a query.
3. Run and design a customized report.
 After creating a customized report, you can modify and redesign it. See [Editing an existing customized report](#) on page 22.




Defining a query statement

Each report can execute multiple queries that determine the customized report data. A query definition - whose result is an SQL SELECT statement - includes the following elements:

- Summary level** Specifies the granularity level of the data: time slice, hourly, daily, weekly, and monthly.
- Report fields** Specifies the fields selected from the Data Tree to be included in the report.
- Filter** Specifies the filtering conditions that determine the data in the report.

You can add, rename, or remove queries using the buttons located to the right of the Query Selection box, in the customized report data selection page.

Table 3-1 Query selection icons

Icon	Description
	Click to add a new query.
	Click to rename an existing query.
	Click to remove an existing query.

To define a query statement:

1. Set the summary level of your query. From the Summary Level box, select the view level for your report.
 The Performance Management Database stores data in different levels of granularity: Time slice, Hourly, Daily, Weekly, and Monthly. This enables you to present a detailed view and progressively higher-level view of the same data in your customized report.


2. Select the fields you want to view in your report. The Data Tree on the left pane represents the fields of the monitored Tiers stored in the Performance Management Database. Each data field is mapped to an expression in the database, for example, Timestamp is mapped to a time of an event, and Program is mapped to a program name. From the Data Tree, select the required field and click the arrow near the Report Fields area. The field is added to the Report Fields area. Repeat this step as for every field you want to view.
3. If you want to remove a field from the report, select the field and then click **Remove**.
4. Click the up and down arrows to modify the order in which the fields will be displayed in the resulting SQL SELECT statement, and in the resulting Microsoft Excel spreadsheet.
5. Use the Sort option to determine whether the rows will be displayed in ascending (A) or descending (D) order in the resulting Microsoft Excel spreadsheet. To sort, first select the box for the selected field. A check mark appears to show that this field was selected for sorting. Next click to set the order of the sort: ascending (A) or descending (D).

Setting filtering conditions

You can set filtering conditions to filter the data displayed in your customized report, through the Filter Table

The following figure shows how to set the filter conditions in your customized report.

Figure 3-1 Customized report data selection page - filtering conditions example



To set filtering conditions

1. From the Data Tree, select the field for which you want to create the condition, and then click (arrow) of the Filter Table. The selected field appears in the Field Selection box.
2. From the Operator box, select the required operator, such as **Like**, **=**, **<**, and **In**.
3. In the Value box, do one of the following:
 - a. Type a value by which the data of the field is filtered.
 - b. Click the List icon and choose a data field from the List of Values dialog box.
4. To remove a condition from the Selected Instances table, select it and then click **Remove**.
5. To update a condition, select it from the Selected Instances Table, modify the condition parameters, and then click **Update**.
6. To modify the order of the conditions, select the required condition from the Selected Instances Table and click **UpArrow** or **DownArrow** (located to the right of the Selected Instances Table).

Viewing and editing the SQL SELECT statement of a query

Report Manager creates an SQL SELECT statement for each query that you define. Actually, this statement is used to retrieve the required data from the Performance Management Database.

You can display the SQL SELECT statement of the query and modify it directly. This feature allows users who have experience with SQL to create complex queries.

After editing the SQL SELECT statement directly, no additional changes are allowed on the query through the Customized Report Data Selection page. For example, you cannot add new fields or remove existing ones from the query, because an SQL statement that is modified directly may include elements that collide with the Report Manager database mapping and SQL generation mechanism.

If you want to return to the original SQL statement to be able to update the query through the Customized Report Data Selection page, click **Revert** in the View SQL dialog box.

To view and edit the SQL SELECT statement of a query

1. After completing the query definition in the Customized Report Data Selection page, click **View SQL** (located at the top right). The View SQL dialog box is displayed.
2. To edit the SQL SELECT statement, click **Edit** and then modify the statement in the query text area.
3. Click **OK**.

Running and designing a customized report

Report Manager generates the customized report using the query that you define in the Customized Report Data Selection page, and then displays the report in a Microsoft Excel file. The Microsoft Excel file consists of the following worksheets:

- One worksheet for each query that you have defined. These worksheets are protected so that no manual changes are allowed (Report Manager overrides these worksheets for each execution). The name of each worksheet is composed of the query name appended to the prefix FS-.
- A worksheet named report-design, in which you can design the customized report. That is, creating charts, tables or any additional elements that describe the data included in the protected worksheets. In addition, you can create new worksheets with additional elements. Report Manager does not override these worksheets in each execution.

NOTE Verify that you do not name the new worksheet with the FS- prefix, because then it will disappear the next time that the report is run or edited.

To run and design a customized report

1. From the Customized Report Data Selection page, click **Execute**. Report Manager generates and then displays the report in a Microsoft Excel file. There are two worksheets in the Microsoft Excel file: the report design worksheet (labeled "report_design") and the protected worksheet (labeled by query name).
2. In the Microsoft Excel file, select the report-design worksheet and design the customized report. When you complete the design, save the Microsoft Excel file.
3. In Report Manager, in the Customized Report Data Selection page, type a name for your customized report in the Report name box, and then click **Save** (located at the top left).

Editing an existing customized report

Report Manager enables you to modify the queries of existing customized reports and to redesign the report-design worksheet of the Microsoft Excel file.

To edit an existing customized report

1. From the Reports Table, select the customized report and then click **Edit**. The Customized Report Data Selection page opens including all the queries defined for the report.
2. To redesign the customized report, click **Spreadsheet** (located at the top right). Report Manager opens the Microsoft Excel file of the last executed customized report.
3. To modify the selected data, change the existing queries or add new ones as required, and then click **Execute**.
Report Manager executes the queries and then displays the report in a Microsoft Excel file.
4. In the Microsoft Excel file, select the report-design worksheet and redesign the customized report. When you complete the design, save the Microsoft Excel file.
5. In Report Manager, in the Customized Report Data Selection page, click **Save**. Report Manager overrides the existing report.

Report Manager settings

Report Manager enables you to configure various settings through the `Settings` menu, located at the Precise bar.

Setting the General settings

The General Settings dialog enables you to set the following items:

- The subfolder in which Report Manager saves the PDF and Microsoft Excel files of executed reports. By default, the subfolder is not defined (Report Manager saves the PDF files in its server file system).
- An option that determines whether Report Manager generates links in the executed reports or not. These links enable users who read the reports to drill down from the reports to the exact location in the Precise or Insight products (depends on the selected entity location) for further investigation. By default, this option is enabled.

To set the General settings

1. From the Precise bar, click **Settings** and from the menu select **General Settings**. The General Settings dialog box is displayed. Check mark the desired settings.
2. To store PDF files in a specified subfolder in the Report Manager server, click the option box and type the path in the box, or open the Browse for Folder dialog box and then select the required subfolder.
3. In the Generate Launching Links option box, select **enable** (to enable) or **clear** (to disable) the links.

Setting the Display settings

The Display Settings dialog box enables you to set the reports that will be listed in the Reports Table when you launch Report Manager. Any settings you make in this dialog box affect only your display and do not affect the reports list when other users launch Report Manager. However, a user with administrator privileges can set the reports list also for other users.

To set the Display settings

1. From the Precise bar, click **Settings** and from the menu select **Display Settings**. The Display Settings dialog box is displayed. By default, all the pre-defined reports are selected.
2. From the Tier box, select the Tier for which you want to modify the reports list.
3. From the Listed to the User table, select the report that you do not want to display and then click **Left Arrow**. Repeat this step for all Tiers as required.

Setting the Multi settings

When you launch Report Manager, the Multi Settings dialog box enables you to set parameters and scheduling times for multiple Report Manager reports. You can select one of the Tiers installed in your application, Cross-Tiers, or ALL (Tiers).

There are various sets of Report Manager reports. With the Multi Settings dialog box you can select one or all sets of Report Manager reports. The sets of Report Manager reports you can select are as follows:

- All (sets)
- Profile
- Exceptions
- Availability
- Load Balancing
- Capacity Planning
- Customized

To set the Multi settings

1. From the Precise bar, click **Settings** and from the menu select **Multi Settings**. The Multi Settings dialog box is displayed.
2. From the Tier box, select the Tier (or all Tiers) you want listed in the Multi Settings reports table.
3. From the Set box, select the report set (or all sets) you want listed in the Multi Settings reports table.
4. To change report parameters and/or scheduling, Press **Ctrl+Shift** to select a range of reports.
5. Click **Properties** to set report properties. See [Setting report properties](#) on page 19.

6. Click **Schedule** to set the time schedule for the reports. See [Setting the schedule of a report generation](#) on page 19.
7. Click **Close** to close the Multi Settings dialog box.

Setting the Mail settings

The Mail Settings dialog box enables you to email Report Manager reports. The Mail Settings dialog box enables you to define the following:

- Enable/Disable the integration of Report Manager with an Email server.
- The name of the mail server.
- The name of the person the Report Manager messages are sent by.

To set the Mail settings

1. From the Precise bar, click **Settings** and from the menu select **Mail Settings**. The Mail Settings dialog box is displayed.
2. Select Integrate Report Manager with an Email server to enable the mail utility.
3. Insert the Email server name and originator of the mail messages.

About the Report Manager reports

The section includes the following topics:

- [Introducing Report Manager reports](#)
- [About the Profile reports set](#)
- [About the Exception Reports set](#)
- [About the Availability reports set](#)
- [About the Load Balancing Reports set](#)
- [About the Capacity Planning reports set](#)

Introducing Report Manager reports

Report Manager provides pre-defined reports divided into the following report sets:

Profile reports set	<p>These reports provide information about the top n consumers of your application and thus enables you to pinpoint tuning opportunities.</p> <p>NOTE: A new subset of the Profile reports set is the Geographical reports.</p> <p>These profile reports provide information about the top n consumers of your application, based on where they are located geographically; whether a state or city in the U.S.A., or a country shown on a two-dimensional world map. These graphic reports enable you to pinpoint tuning opportunities by geographic region.</p>
Exception reports set	<p>These reports provide information about any substantial consumption deviating from normal behavior. This information assists you with maintaining your application well tuned.</p>
Availability reports set	<p>These reports provide information about failures and downtime of the servers.</p>
Load balancing reports set	<p>These reports provide information about servers that are loaded more or less than others, so that you can maintain the load balance of your application.</p>
Capacity planning reports set	<p>These reports provide a forecast about future consumption of the current servers. This information is essential for planning your application evolution. For example, you can plan to acquire new hardware or to restrain the consumption.</p>

All the reports are arranged according to the hierarchy of the Tier as shown in the following general example:

Tier (level 1)	<p>Instance (level 2)</p> <ul style="list-style-type: none"> • Entity (level 3) • Entity (level 3) <p>Instance (level 2)</p> <ul style="list-style-type: none"> • Entity (level 3) • Entity (level 3)
----------------	---

Reports of a Tier-level, for example, provide a performance comparison between the various Tiers of your application. These reports can contain data of several Tiers, for example, Cross-Tiers Availability reports.

Reports of an Instance-level provide information about the complete instance that is the whole system, such as a Web Server, an Oracle Instance, or a Java Virtual Machine.

The structure of each report, in general, contains the following sections:

Opening Page	The opening page is similar for all reports. It contains the following information: <ul style="list-style-type: none"> • Report name • Report description • Summary about the report content • Report execution date-and-time • Report properties of the last execution
Summary	The Summary pages include comparison graphs (where applicable) or a Summary Heat Matrix (overview information), from which you can determine which instance you should start to investigate. Also available are comparison graphs displayed on a map of a geographic region, such as the United States, or on a world map.
Heat Matrix	The Heat Matrix is a table that summarizes the report information using a color code to indicate problematic entities or instances. The Heat Matrix differs for each type of report. The Heat Matrix Table is not displayed in reports where this information is not relevant (such as Availability and Exception reports).
Detailed Information	The detailed information, which relates to each reported entity, includes tables that list the events and Behavior Overtime graphs.

Report Manager generates all the reports and displays them in PDF format, which enables you to navigate within the report using bookmarks (Navigation pane) on the left and links.

There are two types of links: internal links and launching links. For example, the Heat Matrix contains internal links, which you can click to display its detailed information in the report. The launching links enable you to launch the relevant Precise or Insight product - for further investigation - in context with the information displayed in the report page (same entity, Time Frame, etc.).

NOTE You can save a copy of a report in your file system and display them later using Adobe® Acrobat Reader®.

Launching links are available only for reports where this option is enabled. To launch a Precise or Insight product, click the link on the upper right-hand corner of the window. See [Setting the General settings](#) on page 23.

About the Profile report set

The Profile reports provide information about the top n consumers of the application. The consumers can be either Entities such as programs and users, or Instances such as JVMs, SAP Systems, and Oracle Instances.

Using the Profile reports, you can focus on the hot spots, where performance improvement probably can be most effective for the overall behavior of your application.

For each Profile report, Report Manager calculates the top n consumers based on counters, such as In MS-SQL Times and Sessions. You can select the counters by which Report Manager calculates the top n consumers through the Standard profile report - general tab of the report properties dialog box. Each type of report may include a different list of counters.

For each Tier, Report Manager provides the following report types:

- A Profile report at the Instance-level.
- Profile reports at the Entity-level for the entities in the Tier.

The Instance-level report provides a comparison between the top n instances and displays over time information for each one of the top n instances.

The Entity-level report provides a comparison between the top n entities of each instance in the Tier-level. (This comparison applies in case you chose the Same entity option for the Consider Identically Named Entities in Different Instances property in the Standard profile report - general tab.)

The Profile report set includes reports that have similar structure and properties (standard reports). In addition, the Profile report set includes reports that have different structures and properties.

About the standard profile report structure

The standard Profile report structure consists of the following sections:

- [About the Profile report - opening page](#)
- [About the Profile report - summary heat matrix](#)
- [About the Profile report - entity heat matrix](#)
- [About the Profile report - entity mapping table](#)
- [About the Profile report - comparison graphs](#)
- [About the Profile report - overtime graphs](#)

The following information is displayed:

For each instance	Entity Heat Matrix and Entity Mapping Table
For each counter	Comparison graphs
For each entity	Overtime graphs

About the Profile report - opening page

The opening page of the Profile report is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of new entities among the top n consumers. A new entity means that the Performance Management Database did not yet calculate a baseline for the entity.
- Number of entities that normally are not among the top n consumers, and also are not new.

About the Profile report - summary heat matrix

The Summary Heat Matrix indicates, in a tabular format, the top instances with the most consuming entities for each of the reported counters. Each row of the table specifies an instance, and each column specifies a counter. The cell color-code and the left number, in each column, indicates the rank (order in the top n) of the instance for the specific counter. The right number specifies the normal rank, based on its baseline (for instances with no baseline, a dash sign is displayed instead of the right number).

From this page, you can determine which instance you should start to investigate. Normally, you would choose the instance with the highest value for a selected counter. However, the value of the selected counter for the entire instance does not always point to the significance of the top n consumers. Therefore, the instances are ranked by the combined counter value of the top n consumers, and not by the counter value of the entire instance. For example, in an application that runs two instances, HR and CRM, the top-10 users of the HR instance consume 800 hours out of 1,000 hours of the entire HR instance. The top-10 users of the CRM instance consume 500 hours out of 2,000 hours of the entire CRM instance. Even though the CRM instance consumes more than the HR instance, the CRM instance will be ranked the second, while the HR instance will be the first.

The Summary Heat Matrix applies only if the report meets all the following conditions:

- The report is an Entity-level Profile report.
- The report contains more than one instance.
- You chose the Different entity option for the Consider Identically Named Entities in Different Instances property. See [Setting Profile report properties - general tab](#) on page 28.

About the Profile report - entity heat matrix

The Entity Heat Matrix page indicates, in a tabular format, the main consumers of each reported counter. For Instance-level reports, a single Heat Matrix is displayed. For Entity-level reports, an Entity Heat Matrix is displayed for each instance.

Each row of the table specifies an entity, and each column specifies a counter. The cell color-code and the left number, in each column, indicates the rank (order in the top n) of the entity for the specific counter. The right number specifies the normal rank, based on its baseline (for entities with no baseline, a dash sign is displayed instead of the right number).

From this page, you should choose the entity with the highest rank. Each entity name links to the Entity Overtime Behavior graphs.

About the Profile report - entity mapping table

The Entity Mapping Table is displayed in reports whose entity names are too long to display in the detailed information pages. Report Manager generates the Entity Mapping Table only for the Statements reports, such as Oracle or SQL Server Profile Report - Statements and SQL Server Profile Report - Ad-hoc Batches.

In these reports, the detailed information pages display the statement IDs instead of the Statement Text (also a link to the Entity Mapping Table). See [About the Profile report - overtime graphs](#) on page 28.

NOTE You can specify the maximum number of characters in the Statement Text.

About the Profile report - comparison graphs

Report Manager generates a set of Comparison graphs for each of the reported counters. These graphs compare the top n entities of the counter. Instance-level reports have a single set of Comparison graphs for each counter, and Entity-level reports have a set of Comparison graphs for each instance.

Each graph includes a color coded legend. In addition, the legend below the graphs is common to all graphs in the page. This legend maps the number specified in the x-axis of the graph to the entity name (this number is the entity rank). You can click the entity name to display the Entity Overtime Behavior graphs.

NOTE Although an entity may be displayed in the Heat Matrix, it is omitted in the comparison charts of counters that are not part of their top n consumers.

About the Profile report - overtime graphs

Report Manager generates Behavior Overtime graphs for each of the reported entities. Unlike the comparison graphs, all reported entities have the same graph set, regardless of the entities consumption.

The Behavior Overtime graphs contain a launching link to the related Precise or Insight product. You can click the launching link for further investigation in context of the displayed entity. In addition, if the entity name is too long, such as in the Statements report, a link to the Statement Text (Mapping Table) is available.

About standard profile report properties

The Properties dialog box for standard Profile reports contains properties that are specific to the selected report. In general, the properties are similar for all the standard Profile reports. However, you may find some differences according to the relevancy of the property to the report.

The Properties dialog box contains three tabs in which you can set the properties:

- [Setting Profile report properties - general tab](#)
- [Setting Profile report properties - time frame tab](#)
- [Setting Profile report properties - filters tab](#)

In addition to these three tabs, the Description tab displays a textual description of the report structure and contents.

Setting Profile report properties - general tab

The General tab contains the following properties:

Calculate top n entities according to the following counters	This property includes an option box for each counter of the entity (note that the entity is a variable, such as statement, program, and so on). Select the counters, by which Report Manager calculates the top n consumers.
Consider identically-named entities in different instances to be	This property is available only for entity-level reports. Select the Same entities option when the same entities run on different instances, such as in Clustering Systems. Select the Different entities option to distinguish between the entities of the different instances.
Consider identically-named entities in different databases to be	This property is available only in reports that relate to the SQL Server Tier. It is similar to the previous one, except that it relates to SQL databases instead of instances.
Report top n entities	This property determines the number of top consumers to include in the report.
Display statement text up to n characters	This property is available for SQL Statement, Ad-hoc Batches, and Stored Procedures Entity reports. It specifies the maximum number of characters of the Statement Text in the Entity Mapping Table. The default value is 1000 characters. Note that "statement" is a variable.

To set Profile report properties - general tab

1. Select the counters for the entity.

2. If available, you can select an option for each of the drop-down menus.
3. Type the number of top consumers to be included in the report.
4. If available, type the maximum number of characters for the Statement Text in the Entity Mapping table.
5. Click **OK** to approve.

Setting Profile report properties - time frame tab

The Time Frame tab contains the following properties:

Detailed overtime graphs will show	Specifies the period for the Overtime graphs.
Calculate top consumers	Specifies the period by which Report Manager calculates the top entities. This property affects also the summary data in the comparison graphs. Set this period to the same or as part of the period specified in the Detailed Overtime graphs property.
Analyze the following hour groups only	Specifies the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

To set Profile report properties - time frame tab

1. Select the time period for the Overtime graphs.
2. Select the time period by which Report Manager will calculate the top entities.
3. Select the hour group to include in the report.
4. Click **OK** to approve.

Setting Profile report properties - filters tab

The Filters tab contains the following properties:

Analyze only the following instances	Format: wildcard1, wildcard2, ... where a wildcard is a string that may contain asterisks.
Analyze only the following databases	This property is available only for SQL server reports. Its format is the same as in the first property.
Analyze only the following entities	The format is the same as in the first property.
Ignore entities with total time less than <i>n</i> hours	This property is available only for Oracle and SQL Server reports. It filters entities that are considered not important for reporting.

To set Profile report properties - filters tab

1. In the appropriate fields, type the name of the instances, databases, statements, and etc. to be analyzed. (For instances you can click the browse button and select an instance from the list.)
2. If available, type the number of hours for filtering entities.
3. Click **OK** to approve.

About the Application Profile report

The Application Profile report is a Cross-Tiers report that provides information related to all the installed Tiers in the application. The Tiers are compared in terms of total time, average execution time, and the number of executions.

About the Application Profile report structure

The Application Profile report structure consists of the following sections:

- [About the Opening page](#)
- [About Tier comparison graphs](#)
- [About Cross-Tiers behavior overtime graphs](#)
- [About Tier-specific behavior overtime graphs](#)

About the Opening page

The summary section provides the number of installed Tiers in the application.

About Tier comparison graphs

The Tier Comparison graphs display summarized values and baselines of the Total Time, Average Time, and Executions for all the Tiers installed in the application.

About Cross-Tiers behavior overtime graphs

Report Manager generates a Cross-Tiers Behavior Overtime graph for each counter. The counters include the Total Time, Average Time, and Number of Executions. These graphs provide a behavior comparison over time between all installed Tiers.

About Tier-specific behavior overtime graphs

Report Manager generates a Behavior Overtime graph for each installed Tier and each counter within the Tier. The counters include the Total Time, Average Time, and Number of Executions. Each graph also contains the counter baseline.

Setting Application Profile report properties - time frame tab

The Properties dialog box of the Application Profile report contains one tab in which you can set the time frame properties, as follows:

Detailed Overtime graph will show This property determines the report period in the Overtime graphs.

Analyze the following hour groups only This property determines the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

In addition, there is a Description tab which displays a textual description of the report structure and contents.

To set Application Profile report properties - time frame tab

1. Select the report period for the Overtime graphs.
2. Select the hour groups.
3. Click **OK** to approve.

About the Response time report

Report Manager only generates the Response Time report for SAP Tier (client-side technologies). This report provides information about the response times of your system.

For the SAP Tier, the graph displays for each response time the percentage of SAP transactions accordingly that have a lower response time than a maximum value specified in the report properties - Response time report (general tab).

About the Response time report structure

The Response Time report structure consists of the following sections:

- [About the Response time report - opening page](#)
- [About the Response time distribution graph](#)

About the Response time report - opening page

The summary section provides the number of instances included in the report.

About the Response time distribution graph

Report Manager generates a Response Time Distribution graph for each instance running in the Tier. The graph displays a bar for each response time step. Each bar indicates the cumulative percentage and the actual number of executions (SAP transactions in the SAP Tier) that their response time is less or equal to the response time step.

Executions of which the response time is greater than the specified maximal value, are not included in the graph, and therefore the percentage of the total executions may not add up to 100 percent.

About Response time report properties

The Properties dialog box contains three tabs in which you can set the Response Time report properties:

- [Setting Response time report properties - general tab](#)
- [Setting Response time report properties - time frame](#)
- [Setting Response time report properties – filters tab](#)

In addition to these three tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Response time report properties - general tab

The General tab contains the following properties:

Start the graph with a minimal response time of Specifies the starting response time of the x-axis for which the cumulative response time bar will be calculated.

End the graph with a maximal response time of Specifies the ending response time of the x-axis for which the cumulative response time bar will be calculated.

Display response time in steps of Specifies the response time distribution of the x-axis.

To set Response time report properties - general tab

1. In the appropriate fields, type the number of seconds to start and end the response time.
2. Type the number of seconds for response time distribution.
3. Click **OK** to approve.

Setting Response time report properties - time frame tab

The Time Frame tab contains the following properties:

Detailed Overtime graph will show Specifies the report period for the Overtime graphs.

Analyze the following hour groups only Specifies the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

To set Response time report properties - time frame tab

1. Select the time period for the Overtime graphs.
2. Select the hour group to include in the report.
3. Click **OK** to approve.

Setting Response time report properties - filters tab

The Filters tab contains the following property:

Analyze only the following instances Format: wildcard1, wildcard2, ... where a wildcard is a string that may contain asterisks.

To set Response time report properties - filters tab

1. Type the name of the instances to be analyzed.
2. Click **OK** to approve.

About the Schema Changes report

Report Manager only generates the Schema Changes report for Oracle and SQL Server Tiers (database technologies). This report displays a chronological log of all changes in the database schema and definition in a tabular format. You can use this report to identify the relations between changes in the database schema and definition to trends in database performance.

To generate this report, Report Manager compares the status of the database at two points in time. Therefore, changes made after the first point in time of which the effect does not remain till the second point in time, will not be shown in the report. For example, a table that was created after the first point in time but was dropped before the second point in time will not appear in the report.

About the Schema Changes report structure

The Schema Changes report structure consists of the following sections:

- [About the Schema Changes report - opening page](#)
- [About Database Definition changes](#)
- [About Schema changes](#)

NOTE For the SQL Server Tier, the report only displays the schema changes. The database definition changes are displayed in a separate report named Database Definition Changes.

About the Schema Changes report - opening page

The summary section provides the total number of changes detected in the report period.

About Database Definition changes

For the Oracle Tier, the Database Definition Changes Table is part of the Schema Changes report, and the report displays a table for each instance.

For the SQL Server Tier, the Database Definition Changes Table is displayed in a separate report named Database Definition Changes, and the report displays a table for each instance and database.

About Schema changes

For the Oracle Tier, the report displays a table for each instance. For the SQL Server Tier, the report displays a table for each instance and database.

About Schema Changes report properties

The Properties dialog box contains two tabs in which you can set the Schema Changes report properties:

- [Setting Schema Changes report properties - time frame tab](#)
- [Setting Schema Changes report properties - filters tab](#)

In addition to these tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Schema Changes report properties - time frame tab

The Time Frame tab contains the following property:

Analyze information Specifies the report period.

To set Schema Changes report properties - time frame tab

1. Select the time period for the schema changes report.
2. Click **OK** to approve.

Setting Schema Changes report properties - filters tab

The Filters tab contains the following properties:

Analyze only the following instances Format: wildcard1, wildcard2, ... where a wildcard is a string that may contain asterisks.

Analyze only the following databases This property is available only for SQL Server reports. Its format is the same as in the first property.

Analyze only the following owners The format is the same as in the first property.

To set Schema Changes report properties - filters tab

1. In the appropriate fields, type the name of the instances and owners to be analyzed.
2. Click **OK** to approve.

About the Exception Reports set

The Exception reports provide information about any substantial consumption deviating from the normal behavior. This information assists you with maintaining your application well-tuned.

For each Tier, Report Manager provides the following Exception report types:

- Instance-level report type, which provides information about any deviation of the total instance's consumption.
- Entity-level report type, which provides information about any deviation of the entity's consumption. Most of the Exception reports are Entity-level reports.

The counters set may be different for each Tier. In addition, the counters may be different between the various reports in the same Tier. You can set the counters for each report.

For each report, you can set a threshold. When an entity counter consumption deviates from its baseline by more than the threshold, the exception is reported. You can specify up to two thresholds that will escalate the exception severity level.

You can define filters to eliminate negligible exceptions. Two filter types are available, an absolute filter, which defines a minimal total consumption of an entity over the period, and a relative filter, which defines a minimal consumption percentage of an entity from its instance's total consumption.

When Report Manager generates an Exception report, it also sends an alert to Alerts. Likewise, when working in Alerts, you can drill down to the relevant Exception report from the Alerts user interface. For more information, see the *Precise Alerts User's Guide*.

The Exception reports set includes reports that have similar structure and properties (standard reports) described in the following topics:

- [About the Standard Exception report structure](#)
- [About Standard Exception report properties](#)

In addition, the Exception reports set includes a report named Unused objects report that has different structure and properties.

About the Standard Exception report structure

The standard Exception report consists of the following sections:

- [About the Exception report – opening page](#)
- [About the Exception report - summary heat matrix](#)
- [About the Exception report - entity heat matrix](#)
- [About the Exception report - new entities table](#)
- [About the Exception report - entity mapping table](#)
- [About the Exception report - summary table and overtime graphs](#)

An instance for a standard Exception report displays the following tables:

- Entity Heat matrix
- New Entities table
- Entity Mapping table

The following information is displayed for each entity:

- Exception Summary Table
- Overtime Behavior Graphs

About the Exception report - opening page

The opening page of the Exception report is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of entities that met the exception criteria, with their mean deviation and maximal deviation from the baseline.
- Total number of exceptions and their distribution among the various counters.
- Number of new entities within the Tier. A new entity means that the Performance Management Database did not yet calculate a baseline for the entity.

About the Exception report - summary heat matrix

The Summary Heat Matrix indicates, in a tabular format, the top instances with the most severe exceptions for each of the reported counters. It summarizes the number of deviating entities that are detected and their severity level.

In Entity-level reports, each row in the Summary Heat Matrix represents a different instance. Each column represents a different counter. The cell color indicates the average deviation severity of all the reported entities of the instance. The number in the cell indicates the number of entities having exceptions. The instance name is a link to the Related Entity Heat Matrix.

In Instance-level reports, the Summary Heat Matrix displays a single row that indicates the average deviation severity and exceptions number in the entire Tier. (This also applies if you chose the Same entity option for the Consider Identically Named Entities in Different Instances property in the Standard exception report - general tab.)

About the Exception report - entity heat matrix

The Entity Heat Matrix page indicates, in a tabular format, the most problematic entities that have the most significant deviations. Each row in the Entity Heat Matrix represents a different entity. Each column represents a different counter. The cell color indicates the severity of the counter deviation from its normal behavior (baseline). The entity name is a link to the Entity Overtime Behavior graphs.

About the Exception report - new entities table

The New Entities Table lists new entities in the current instance (or Tier in Instance-level reports), which met the report filtering criteria of the minimal consumption amount. Entities are considered new if the Performance Management Database did not calculate a baseline for the entity yet.

For each new entity the table displays its total time value and the percentage that its total time occupies from the entire instance consumption (or Tier consumption in case of Instance-level reports).

About the Exception report - entity mapping table

The Entity Mapping Table is displayed in reports whose entity names are too long to display in the detailed information pages. Report Manager only generates the Entity Mapping Table for Statements reports, such as the Oracle or SQL Server Exception Report - Statements and the Oracle or SQL Server Exception Report - Ad-hoc Batches.

In these reports, the detailed information pages display the statement IDs instead of the Statement Text (also a link to the Entity Mapping Table). See [About the Exception report - summary table and overtime graphs](#) on page 34.

About the Exception report - summary table and overtime graphs

Report Manager generates an Exception Summary Table and Behavior Overtime graphs for each of the reported entities. Each row in the Exception Summary Table includes a deviated counter, with its most recent exception details, including: the Exception Time, the Actual Counter Value, its Baseline Value, and the Percent of Change.

Following the Exception Summary Table, a set of Overtime graphs is displayed, one for each counter. Each graph displays the Actual Counter Behavior, the Counter Baseline, and the Exception Threshold.

This page contains a launching link to the related Precise or Insight product. You can click the launching link for further investigation in context of the displayed entity. In addition, if the entity name is too long, such as in the Statements report, a link to the Statement Text (Mapping Table) is available.

About Standard Exception report properties

The Properties dialog box for Standard Exception reports contains properties that are specific to the selected report. In general, the properties are similar for all the Standard Exception reports. However, you may find some differences according to the relevancy of the property to the report.

The Properties dialog box contains three tabs in which you can set the properties:

- [Setting Exception report properties - general tab](#)
- [Setting Exception report properties - time frame tab](#)
- [Setting Exception report properties - filters tab](#)

In addition to these three tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Exception report properties - general tab

The General tab contains the following properties:

Report the following exceptions	This property includes an option box for each counter of the entity. Select the counters, by which Report Manager detects the exceptions. The counters and their default selections are different for each Tier and each report within the same Tier.
Report an exception when baseline is exceeded by more than <i>n</i> percent	This property determines the percentage by which Report Manager reports the first level exception. One color-code represents this exception in the Heat Matrix.
Then, escalate exception when baseline is exceeded by more than <i>n</i> percent	This property determines the percentage by which Report Manager reports the second level exception. This property is optional, that is, if you set it, a second color-code represents this exception in the Heat Matrix.
Then, escalate exception when baseline is exceeded by more than <i>n</i> percent	This property is enabled only if the previous one is set. It determines the percentage by which Report Manager reports the third level exception. This property is optional, that is, if you set it, a third color-code represents this exception in the Heat Matrix.
Consider identically-named entities in different instances to be	This property is available only for entity-level reports. Select the Same entities option when the same entities run on different instances, such as in Clustering Systems. Select the Different entities option to distinct between the entities of the different instances.
Consider identically-named entities in different databases to be	This property is available only in reports that relate to the SQL Server Tier. It is similar to the previous one, except that it relates to SQL databases instead of instances.
Display Statement Text up to <i>n</i> characters	This property is available only for SQL statement entity reports. It specifies the maximum number of characters of the Statement Text in the entity mapping table. The default value is 1000 characters.

To set Exception report properties - general tab

1. Check the desired exceptions.
2. Type the desired percent number by which Report Manager reports the first level exception.
3. Check and type the percent number if you want Report Manager reports for second level and third level exceptions.
4. If available, you can select an option for each of the drop-down menus.
5. If available, type the maximum number of characters for the Statement Text in the Entity Mapping table.
6. Click **OK** to approve.

Setting Exception report properties - time frame tab

The Time Frame tab contains the following properties:

Detailed Overtime graphs will show	Specifies the period for the Overtime graphs.
Detect exceptions	Specifies the period by which Report Manager detects exceptions. Set this period to the same or as part of the period specified in the detailed Overtime Graphs property.
Analyze the following hour groups only	Specifies the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

To set Exception report properties - time frame tab

1. Select the time period for the Overtime graphs.

2. Select the time period by which Report Manager detects exceptions.
3. Select the hour group to include in the report.
4. Click **OK** to approve.

Setting Exception report properties - filters tab

The Filters tab contains the following properties:

Analyze only the following instances	Format: wildcard1, wildcard2, ... where a wildcard is a string that may contain asterisks.
Analyze only the following databases	This property is only available for SQL Server reports. Its format is the same as in the first property.
Analyze only the following entities	The format is the same as in the first property.
Ignore entities with total time less than <i>n</i> hours	This property filters exceptions that are not considered important for reporting.
Ignore entities with total time less than <i>n</i> percent of their instance	This property provides an additional filtering condition (to the previous one). It uses relative values to filter exceptions that are not considered important for reports that the previous filter did not eliminate. For example, using the previous property, one hour per week may be sufficient for some applications but for others may not.

To set Exception report properties - filters tab

1. In the appropriate fields, type the names of the instances, databases, statements, and etc. to be analyzed. (For instances you can click the browse button and select an instance from the list.)
2. If available, type the number of hours for filtering exceptions.
3. If desired, type the percent number for an additional filter to exceptions.
4. Click **OK** to approve.

About the Unused Objects report

This report is applicable for database technologies, such as Oracle and SQL Server. The report lists the objects that were not used during the specified time frame of the report. You can use this report to identify cases such as forgotten tables and redundant indices. These cases are frequent in data warehouse applications, which may contain tables and clusters that are not accessed by statements. By dropping these unused objects, you can free disk space without affecting the database performance.

About the Unused Objects report structure

The unused objects report structure consists of the following sections:

- [About the Unused Objects report - opening page](#)
- [About the Unused Objects table](#)

About the Unused Objects report - opening page

The summary section provides the number of unused objects in the report.

About the Unused Objects table

Report Manager generates tables of unused objects as follows:

- For the Oracle Tier, one table for each instance and schema owner.
- For the SQL Server Tier, one table for each instance, database, and schema owner.

The table lists all the unused objects detected during the report period, their owner, and their type.

About the Unused Objects report properties

The Properties dialog box contains two tabs in which you can set the Unused Objects report properties:

- [Setting Unused Objects report properties - time frame tab](#)
- [Setting Unused Objects report properties - filters tab](#)

In addition to these tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Unused Objects report properties - time frame tab

The Time Frame tab contains the following property:

Analyze information Specifies the report period by which Report Manager detects unused objects.

To set Unused Objects report properties - time frame tab

1. Select the time period by which Report Manager detects unused objects.
2. Click **OK** to approve.

Setting Unused Objects report properties - filters tab

The Filters tab contains the following properties:

Number of blocks greater than *n* This property specifies the minimum number of blocks of objects that Report Manager reports. This eliminates reporting of small objects.

Consider only the following object types This property determines the type of objects to report so that you can concentrate only on a specific object type: Table, Index, or Cluster.

Analyze only the following instances Format: wildcard1,wildcard2 ... where a wildcard is a string that may contain asterisks.

Analyze only the following databases This property is only available for SQL Server reports. Its format is the same as in the previous property.

Analyze only the following owners The format is the same as in the previous property.

To set Unused Objects report properties - filters tab

1. Type the minimum number of blocks for objects that Report Manager reports.
2. Select the type of objects to report.
3. In the appropriate fields, type the name of the instances, databases, and owners to be analyzed.
4. Click **OK** to approve.

About the Availability reports set

The Availability reports provide information about failure and downtime of Tiers in your application. For most Tiers, Report Manager tracks only the Instance availability. However, for some Tiers, such as SAP and SQL Server, Report Manager also tracks the availability of other entities.

Report Manager provides the following Availability report types:

- Instance-level Availability report type for all Tiers except OS. This report type provides information about the availability of all the instances in the Tier.
- Entity-level Availability report type for entities other than Instance-level, whose availability is being tracked.
- Availability Overview report type for Tiers that include Entity-level reports. The Overview report provides a summary of the Tier availability, by combining the availability information of the Instance-level and Entity-level in one table.
- Cross-Tiers Availability report that summarizes the availability information of all Tiers in one report.

The availability percentage calculation takes into account planned down time, which can be configured through Precise AdminPoint. For more information, see the *Precise Administration Guide*.

When Report Manager generates an Availability report, it checks if an instance availability falls below its availability threshold, unless there is no information (regarding availability) for the Tier. For example, a newly installed Tier may not have data currently available regarding availability; in this case a bar labeled Unknown appears on the graph. If an instance availability falls below its availability threshold, Report Manager sends an alert to Alerts. Likewise, when working in Alerts, you can drill down to the relevant Availability report from the Alerts user interface. For more information, see the *Alerts User's Guide*.

About the Instance-level Availability report structure

The Instance-level Availability report structure consists of the following sections:

- [About the Instance-level Availability report - opening page](#)
- [About the Tier Availability table](#)
- [About the Instance-level Availability Overtime graphs](#)

About the Instance-level Availability report - opening page

The opening page of the Instance-level Availability report is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of reported instances, and their average and minimal availability within the report period.

About the Tier Availability table

The Tier Availability Table displays the availability details of each instance of the Tier. These details include the Average and Minimal Availability Percentage, and the Availability Threshold. The Availability Threshold can be configured through Precise AdminPoint. If the Instance Availability is below the threshold, Report Manager indicates its data in red. If there is no data currently available regarding availability, Report Manager indicates no data as Unknown and in blue.

By default, Report Manager generates Overtime graphs only for instances whose availability is below their threshold. Therefore, only for these instances, the names in the Instance column of the table are links, while the others are not links. You can change this default setting so that all the names will be links. See [Setting Availability report properties - filters tab](#) on page 40.

About the Instance-level Availability Overtime graphs

Report Manager generates, by default, Overtime graphs for each instance that its availability is below the threshold. The Availability Overtime graphs include the following:

- An Availability graph that displays the availability percentage of the instance along with the threshold. For a Tier where no data is currently available, the bar appears as Unknown in the graph.
- An Unavailability graph that displays the duration of the instance unavailability.

To generate Overtime graphs of instances with availability above the threshold, modify the report properties and rerun it. See [Setting Availability report properties - filters tab](#) on page 40.

The Availability Overtime graphs contain a launching link to the related Precise or Insight product. You can click the launching link for further investigation in context of the displayed instance.

About the Entity-level Availability report structure

The Entity-level Availability report structure consists of the following sections:

- [About the Entity-level Availability report - opening page](#)
- [About the Entity-level Availability report - summary table](#)

The following information is displayed for each instance:

- [About the Entities Availability table](#)
- [About the Entities Availability Overtime graphs](#)

About the Entity-level Availability report - opening page

The opening page of the Entity-level Availability report is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of Reported Entities
- Their Average and Minimal Availability within the report period

About the Entity-level Availability report - summary table

This Summary Table compares all the instances of the Tier by the availability of their entities. The Summary table contains a row for each reported instance that displays:

- Its Total Amount of Entities
- The Number of Entities with Availability Problems
- The Average Entities Availability within the report period

About the Entities Availability table

The Entities Availability Table displays the availability details of each entity in the current instance. These details include the Average and Minimal Availability Percentage, and the Availability Threshold. The Availability Threshold can be configured through Precise AdminPoint. If the entity availability is below the threshold, Report Manager indicates its data in red.

By default, Report Manager generates Overtime graphs only for entities whose availability is below their threshold. Therefore, only for these entities, the names in the System column of the table are links, while the others are not links. You can change this default setting so that all the names will be links. See [Setting Availability report properties - filters tab](#) on page 40.

About the Entities Availability Overtime graphs

Report Manager generates, by default, Overtime graphs for each entity that its availability is below the threshold. The Availability Overtime graphs include the following:

- An Availability graph that displays the availability percentage of the entity along with the threshold.
- An Unavailability graph that displays the duration of the entity unavailability.

To generate Overtime graphs of entities with availability above the threshold, modify the report properties and rerun it. See [Setting Availability report properties - filters tab](#) on page 40.

The Availability Overtime graphs contain a launching link to the related Precise or Insight product. You can click the launching link for further investigation in context of the displayed entity.

About the Availability Overview report structure

The Availability Overview report structure consists of the following sections:

- [About the Availability Overview report - opening page](#)
- [About the Availability Overview table](#)

About the Availability Overview report - opening page

The opening page of the Availability Overview report is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- For each reported entity type, Number of Entities and their Average and Minimal Availability within report period.

About the Availability Overview table

The Availability Overview Table lists for each instance of the Tier the availability information for each level, that is, instance-level and any entity for which availability information is being collected.

The table contains a row for each reported instance that displays the Instance Availability Percentage, the Total Amount of Entities for each entity type and their Average Availability Percentage.

About the Cross-Tiers Availability report structure

The Cross-Tiers Availability report structure consists of the following sections:

- [About the Cross-Tiers Availability report - opening page](#)
- [About the Application Availability table](#)
- [About Availability Overtime graphs](#)

About the Cross-Tiers Availability report - opening page

The opening page of the Cross-Tiers Availability report is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of reported Tiers and their Average and Minimal Availability within the report period.

About the Application Availability table

The Application Availability table displays the availability details of each Tier in the application. These details include the Average and Minimal Availability Percentage, and the Availability Threshold. The Availability Threshold can be configured through Precise AdminPoint. If the instance availability is below the threshold, Report Manager indicates its data in red.

By default, Report Manager only generates Overtime graphs for Tiers whose availability is below their threshold. Therefore, only for these Tiers, the names in the Tier column of the table are links, while the others are not links. You can change this default setting so that all the names will be links. See [Setting Availability report properties - filters tab](#) on page 40.

About Availability Overtime graphs

Report Manager generates, by default, Overtime graphs for each Tier of which its availability is below the threshold. The Availability Overtime graphs include the following:

- An Availability graph that displays the availability percentage of the Tier along with the threshold.
- An Unavailability graph that displays the duration of the Tier unavailability.

To generate Overtime graphs of Tiers with availability above the threshold, modify the report properties and rerun it. See [Setting Availability report properties - filters tab](#) on page 40.

About Availability report properties

The Properties dialog box for the Availability reports contains two tabs in which you can set the properties:

- [Setting Availability report properties - time frame tab](#)
- [Setting Availability report properties - filters tab](#)

In addition to these tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Availability report properties - time frame tab

The Time Frame tab contains the following properties:

Report on events	Specifies the period by which Report Manager calculates the availability percentage and determines the events to display in the Overtime graphs.
Analyze the following hour groups only	Specifies the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

To set Availability report properties - time frame tab

1. Select the time period by which Report Manager calculates the availability percentage and determines the events to display in the Overtime graphs.
2. Select the hour groups to include in the report.
3. Click **OK** to approve.

Setting Availability report properties - filters tab

The Filters tab contains the following properties:

Analyze only the following instances	Format: wildcard1,wildcard2,... where a wildcard is a string that may contain asterisks.
Analyze only the following entities	The format is the same as in the first property.
Generate Overtime graphs only for entities which did not comply with their availability threshold	Clear this option box to display Availability Overtime graphs also for entities that their availability is above the threshold.

To set Availability report properties - filters tab

1. In the appropriate fields, type the names of the instances and databases to be analyzed.
2. Clear this option box to display Availability Overtime graphs.
3. Click **OK** to approve.

About the Load Balancing Reports set

Load Balancing reports assist you to maintain the application balanced, by providing information about servers that are loaded more or less than others.

Report Manager provides two types of Load Balancing reports:

Cross-Tiers Load Balancing Report	This report type compares all the servers installed in the application, and displays a high level balancing view of the entire application.
Tier Load Balancing Report	This report type assists you only to balance the machines in which instances of the Tier are currently running.

Report Manager compares the servers' load using the operating system counters. Therefore, these reports are only available when Insight's OS Tier is installed.

For more information, see the *Precise Installation Guide*. The operating system counters include:

- CPU usage (%)
- CPU load - average number of processes waiting for CPU
- Paging

Report Manager compares the servers also by the Tier's Total Time, Volume, and Average Time. To overcome the differences among the servers' power (such as, CPU and memory), the different servers are assigned with weights. By default, a UNIX server is weighted 7 and a PC is weighted 4. The Total and Average Times are multiplied by the weight. The weights of the servers are stored in the Performance Management Database and can be updated using appropriate SQL UPDATE statements.

In the report properties, you can set the deviation percentage from the average, by which Report Manager determines if the server is loaded. If the server is overloaded or below the average load, Report Manager sends an event to Alerts.

Report Manager generates the Tier Load Balancing report for all the servers that run instances of the specified Tier (Report Manager does not generate the Tier Load Balancing report for Tiers that contain a single server).

About the Cross-Tiers Load Balancing report structure

The Cross-Tiers Balancing report structure consists of the following sections:

- [About the Cross-Tiers Load Balancing report - opening page](#)
- [About Cross-Tiers Balancing - heat matrix](#)

About the Cross-Tiers Load Balancing report - opening page

The opening page of the Cross-Tiers Load Balancing reports is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of servers that are overloaded, and number of servers whose load is significantly under the average load
 (the number of servers that have one of the OS counters significantly above or under the average value).
- For each OS counter, the number of servers overloaded and underloaded related to the counter's average value.

About Cross-Tiers Balancing - heat matrix

The Cross-Tiers Balancing Heat Matrix summarizes the load of each server in the application. Each row in the Balancing Heat Matrix represents a different server, and each column represents one of the balanced OS counters: CPU Usage, CPU Load, and Paging.

The color of the cell indicates the balancing state of the selected OS counter for the current server. Red indicates an overloaded server and orange indicates that the server load is under average.

The Heat Matrix also introduces technologies indicating where instances are installed on each server.

About the Tier Load Balancing report structure

The Tier Load Balancing report structure consists of the following sections:

- [About the Tier Load Balancing report - opening page](#)
- [About OS Summary graphs](#)
- [About OS Overtime graphs](#)
- [About Tier Overtime graphs](#)
- [About Tier Consumption Summary graphs](#)

About the Tier Load Balancing report - opening page

The opening page of the Tier Load Balancing reports is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- Number of servers that are overloaded and number of servers whose load is significantly under the average load (the number of servers that have one of the OS counters significantly above or under the average value).
- For each OS counter, the number of servers overloaded and underloaded related to the counter's average value.

About OS Summary graphs

Report Manager generates an OS Summary graph for each balanced OS counter, such as CPU Usage and CPU Load and Paging. This graph displays the Counter Summarized Value versus the Average Load, for each reported machine.

About OS Overtime graphs

Report Manager generates an OS Overtime graph for each balanced OS counter, such as CPU Usage, CPU Load and Paging. This graph displays an Overtime Behavior graph of the Top n and Bottom m Servers compared to the Counter Average Value.

About Tier Overtime graphs

This graph displays a Summary of the Overtime Behavior of the reported Tier. It displays the Actual Values of the counters for the Top n and Bottom m Servers.

Report Manager generates a Tier Overtime graph for each of the following counters:

- Total Time
- Average Time
- Volume

About Tier Consumption Summary graphs

This graph displays a Summary of the Counter Values for each installed Tier on the Top n and Bottom m Servers. Report Manager generates a Tier Consumption Summary graph for each of the following counters:

- Total Time
- Average Time
- Volume

About Load Balancing report properties

The Properties dialog box of the Load Balancing report contains three tabs in which you can set the properties:

- [Setting Load Balancing report properties - general tab](#)
- [Setting Load Balancing report properties - time frame tab](#)
- [Setting Load Balancing report properties - filters tab](#)

In addition to these three tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Load Balancing report properties - general tab

The General tab contains the following properties:

Report an exception when a counter exceeds average by more than <i>n</i> percent	This property determines the threshold percentage that Report Manager considers the server consumption as overloaded.
Report an exception when a counter falls below the average by more than <i>n</i> percent	This property determines the threshold percentage that Report Manager considers the server consumption as under the average consumption. This property is optional. If you disable this property, Report Manager does not report on server whose consumption is less than the average
Report top <i>n</i> servers	This property determines the number of servers that are most overloaded.
Report bottom <i>m</i> servers	This property determines the number of servers that are most underutilized.

To set Load Balancing report properties - general tab

1. Type the desired percent number threshold that Report Manager will use to report servers whose consumption is overloaded.
2. If desired, check the option and type a percent number if you want Report Manager to report on servers whose consumption is less than the average.
3. Type the desired number of servers to report that are most overloaded.
4. Type the desired number of servers to report that are most underutilized.
5. Click **OK** to approve.

Setting Load Balancing report properties - time frame tab

The Time Frame tab contains the following properties:

Detailed Overtime graph will show	Specifies the period for the Overtime graphs. (This property is not available for the Cross-Tiers report.)
Detect exceptions	Specifies the period by which Report Manager detects exceptions. Set this period to the same or as part of the period specified in the Detailed overtime graph will show property.
Summary graphs depict	This property determines the period, which is used for a Summary graph calculation. This period should be included within detailed Overtime graph period. This property is not presented in Cross-Tiers report.
Analyze the following hour groups only	Specifies the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

To set Load Balancing report properties - time frame tab

1. Select the time period for Overtime graphs.
2. Select the time period by which Report Manager detects exceptions.
3. Select the time period to be used for a Summary graph calculation.
4. Select the hour groups to include in the report.
5. Click **OK** to approve.

Setting Load Balancing report properties - filters tab

The Filters tab contains the following properties:

Analyze only the following instances	Format: wildcard1,wildcard2,... where a wildcard is a string that may contain asterisks. (This property is not available for the Cross-Tiers report.)
Analyze only the following servers	The format is the same as in the first property.

To set Load Balancing report properties - filters tab

1. In the appropriate fields, type the names of the instances and servers to be analyzed.
2. Click **OK** to approve.

About the Capacity Planning reports set

The Capacity Planning report provides a forecast about future consumption of the current instances. This information is essential for planning your application evolution. For example, you can plan to acquire new hardware or to restrain the consumption (If you choose to restrain the consumption, you can use the Profile reports set to detect the top consumers of the specific resource).

Capacity Planning reports take into consideration only the current active systems, and do not assume any new programs and enhancements, which may be applied in the future.

When Report Manager generates a Capacity Planning report, it checks if the expected increase in a forecasted counter exceeds the baseline by more than the increase threshold. In this case, it sends an alert to Alerts. Likewise, when working in Alerts, you can drill down to the relevant Capacity Planning report from the Alerts user interface. For more information, see the *Precise Alerts User's Guide*.

About the Capacity Planning report structure

A Capacity Planning report structure consists of the following sections:

- [About the Capacity Planning report - opening page](#)
- [About the Capacity Planning report - heat matrix](#)
- [About Behavior and Forecast overtime graphs](#)

About the Capacity Planning report - opening page

The opening page of the Tier Capacity Planning reports is similar to all the pre-defined reports except that the summary section provides information related to the specific report as follows:

- For each forecasted counter, the average increase of the counter consumption at the end of the forecast period.

About the Capacity Planning report - heat matrix

The Heat Matrix page indicates, in a tabular format, the most significant consumption increase of the forecasted counters. Each row in the Heat Matrix represents a different instance. Each column represents a different forecasted counter. The cell color indicates the increase level of the counter related to its increase thresholds. You can set the increase thresholds through the Capacity planning report - general tab of the report Properties dialog box. The instance name is a link to the instance's Overtime Behavior graphs.

About Behavior and Forecast overtime graphs

Report Manager generates - for each instance - a Behavior and Forecast Overtime graph for each predicted counter. Each Overtime graph presents the past behavior, baseline, and forecast of the counter.

The Behavior and Forecast Overtime graphs page contain a launching link to the related Precise or Insight product. You can click the launching link for further investigation in context of the displayed instance.

About Capacity Planning report properties

The Properties dialog box of the Capacity Planning report contains three tabs in which you can set the properties:

- [Setting Capacity Planning report properties - general tab](#)
- [Setting Capacity Planning report properties - time frame tab](#)
- [Setting Capacity Planning report properties - filters tab](#)

In addition to these three tabs, there is a Description tab which displays a textual description of the report structure and contents.

Setting Capacity Planning report properties - general tab

The General tab contains the following properties:

Forecast the following counters	This property includes an option box for each counter of the instance. Select the counters, by which Report Manager calculates the forecast.
Escalate the increase when current consumption is exceeded by more than <i>n</i> percent	This property determines the threshold percentage that Report Manager considers as the first level increase in instance's consumption (appears yellow in the Heat Matrix Table).

Then, escalate increase when the current consumption is exceeded by more than n percent

This property determines the threshold percentage that Report Manager considers as the second level increase in instance's consumption (appears orange in the Heat Matrix Table). This property is optional. If you do not select it and specify the percentage, only one color will be presented in the Heat Matrix Table.

Then, escalate increase when the current consumption is exceeded by more than n percent

This property determines the threshold percentage that Report Manager considers as the third level increase in instance's consumption (appears red in the Heat Matrix Table). This property is optional. If you do not select it and specify the percentage, only two colors will be presented in the Heat Matrix Table.

To set Capacity Planning report properties - general tab

1. Check the desired number of counters by which Report Manager will calculate the forecast.
2. Type the threshold percent number that Report Manager will consider as the first level increase in instance consumption.
3. If desired, check this option and type the threshold percent number that Report Manager will consider as the second level increase in instance consumption.
4. If desired, check this option and type the threshold percent number that Report Manager will consider as the third level increase in instance consumption.
5. Type the percent number to define the probability that the forecast is correct.
6. Click **OK** to approve.

Setting Capacity Planning report properties - time frame tab

The Time Frame tab contains the following properties:

Detailed Overtime graph will show

Specifies the period for the Overtime graphs.

Base the forecast on information from

Specifies the history period by which Report Manager calculates the forecast. Set this period to the same or as part of the period specified in the Detailed Overtime graph will show property.

Extrapolate over the next

Specifies the forecast period. To get accurate forecasting results, it is recommended to set the forecast period at most half of the history period.

Analyze the following hour groups only

Specifies the hour groups (for example, All, Night, Day, and Weekend) to include in the report.

To set Capacity Planning report properties - time frame tab

1. Select the time period for Overtime graphs.
2. Select the history time period by which Report Manager calculates the forecast.
3. Type the number of weeks for the forecast period.
4. Select the hour groups to analyze.
5. Click **OK** to approve.

Setting Capacity Planning report properties - filters tab

The Filters tab contains the following property:

Analyze only the following instances

Format: wildcard1, wildcard2,... where a wildcard is a string that may contain asterisks.

To set Capacity Planning report properties - filters tab

1. Type the name of the instances to be analyzed.
2. Click **OK** to approve.

Advanced features that can help you customize your reports

This appendix includes the following topics:

- [About customizing your reports](#)
- [About defining properties for a customized report](#)
- [Filling in gaps in a customized report](#)
- [About importing and exporting customized reports](#)
- [About copying customized reports](#)

About customizing your reports

Similar to predefined reports (Profile, Exceptions, and so on), customized reports can have property definitions set through the Properties dialog box. However, this is currently not supported by the Report Manager User Interface and must be implemented manually.

Retrieving over time information from the Performance Management Database may result in gaps (missing information). These are specific timestamps for which no data is available. To fill in these gaps, Report Manager FocalPoint supports holes filling. However, the current version of the Report Manager User Interface does not support the holes filling definition.

To implement these features, a few manual changes must be made to the tables and XML files of Report Manager FocalPoint. These changes should only be carried out by a person who is experienced with SQL and familiar with XML.

NOTE These two features are supported by Report Manager FocalPoint, but are not supported by the Report Manager User Interface.

About defining properties for a customized report

Defining properties for a Customized report enables the user to rerun a report with its preferred property values. This procedure enables you to add new properties to the General, Filtering, and Time Frame tabs in the Properties dialog box.

After you complete this procedure, do not edit the Customized report by clicking Edit, because Report Manager will override the newly defined properties. If you have to edit the report, you will have to carry out this procedure again to restore these properties.

Defining new properties for a Customized report includes the following stages:

- [About creating a customized report](#)
- [About updating Report Manager tables in Performance Management Database](#)
- [Updating the .xml file with the report](#)

About creating a customized report

Set a filter condition that includes a fixed valid value for each new property, (the value should be distinguishable, so that you can easily recognize it later when you replace it with the new property name in the generated report's .xml file).

When you complete the filter condition settings for all the new properties, execute the customized report and then save it. Report Manager creates two files, the report's Microsoft Excel file and .xml file. See [Creating customized reports](#) on page 20.

About updating Report Manager tables in Performance Management Database

The following changes should be made to the Report Manager tables of the Performance Management Database, in order that the new properties you defined for your customized report can be supported:

- [About adding a new property to the Report Parameters table](#)
- [About locating the report identifier](#)
- [About associating a new property to the customized report](#)

NOTE Check with your Precise administrator about the location of the Performance Management Database and user permission details.

About adding a new property to the Report Parameters table

Open the Report Parameters Table, named PS_FSRA_REPORT_PARAMETERS, using a database tool, such as SQL*Plus or Microsoft SQL Server Enterprise Manager (depending on your Performance Management Database type).

Add the new property to the Report Parameters Table as described in the following table. Verify that the property identifier (FSRA_ID) and name (FSRA_NAME) are unique in this table.

The following table describes the properties that can be added to the Report Parameters table.

Table A-1 Additional properties that can be added to the Report Parameters table

Column	Description
FSRA_ID	Specifies the property identifier. Set a unique and consecutive number, starting from 100000 (enables you to easily distinguish between the new properties you added and the existing pre-defined properties).
FSRA_NAME	Specifies a unique property name.
FSRA_IS_DISPLAYABLE	Specifies whether or not to display the property through the Properties dialog box. Valid values are: <ul style="list-style-type: none"> • T - True • F - False Set this value to 'T'.
FSRA_IS_CONVERTIBLE	Specifies whether or not to convert the % (wildcard sign) to * in the Properties dialog box. The valid values include: <ul style="list-style-type: none"> • T - True • F - False If the property value contains wildcards, set this value to 'T'.
FSRA_DEFAULT_VALUE	Specifies the default value of the property. Set this value to NULL, since you will later define the default value when you associate the new property to the customized report.
FSRA_PROMPT_TEXT	Specifies the text that is displayed with the property box in the Properties dialog box. Enter a description.
FSRA_TAB_ID	Specifies the tab ID - of the report Properties dialog box - in which this property is displayed. Set to one of the following valid values according to your property: <ul style="list-style-type: none"> • G - General • F - Filters • T - Time Frame
FSRA_INDEX	Specifies the property index in the tab, indicating its order.

If you use a database tool such as Microsoft SQL Server Enterprise Manager, you can use the following INSERT statement as an example, to add a new parameter to the PS_FSRA_REPORT_PARAMETERS table.

```
INSERT INTO PS_FSRA_REPORT_PARAMETERS
(FSRA_ID, FSRA_NAME, FSRA_IS_DISPLAYABLE,
FSRA_IS_CONVERTIBLE, FSRA_DEFAULT_VALUE,
```

```
FSRA_PROMPT_TEXT,FSRA_TAB_ID,FSRA_INDEX)
VALUES (100000,'ProgramFilter','T','T',NULL,'
Analyze only the following programs:', 'F', 10)
```

About locating the report identifier

Open the Reports Table (named PS_FSRP_REPORTS) and search for the Report Identifier of the new report that you created in the FSRP_ID column. You can search for it according to the FSRP_NAME column, which contains the report name you specified when saving the report, or according to the FSRP_GENERATOR_NAME column that contains the string 'excel'.

About associating a new property to the customized report

To associate the new property to the Customized report, set the property values in the Parameter Values Table (named PS_FSPV_PARAMETER_VALUES).

The following table describes the parameters values which can be customized.

Table A-2 Parameter Values

Column	Description
FSPV_FSRA_ID	Specifies the property identifier. Set the number as you specified in the FSRA_ID column of the Report Parameters Table. See About adding a new property to the Report Parameters table on page 47.
FSPV_FSPG_ID	Specifies the group identifier. Set to NULL.
FSPV_FSRP_ID	Specifies the report identifier. Set the number that you found in the Reports Table. See About locating the report identifier on page 48.
FSPV_CTX_ID	Specifies the context identifier. To find the context identifier of the new property, run the following query: SELECT FSPV_CTX_ID FROM PS_FSPV_PARAMETER_VALUES WHERE (FSPV_FSRP_ID = report_id) AND (FSPV_FSRA_ID = parameter_id) Where <code>report_id</code> is the report identifier that you found in the Reports Table and <code>parameter_id</code> is the report parameter that you added in the Report Parameters Table. See About adding a new property to the Report Parameters table on page 47. The query result includes the required value for FSPV_CTX_ID. Set the resulted value accordingly.
FSPV_VALUE	Specifies the default value of the new property.
FSPV_PROMPT_TEXT	Specifies the text that is displayed with the property box in the Properties dialog box. You can set to NULL to apply the text that you already set in the FSRA_PROMPT_TEXT column of the Report Parameters Table, or you enter a new text to override it. See About adding a new property to the Report Parameters table on page 47.

If you use a database tool such as Microsoft SQL Server Enterprise Manager, you can use the following example to add a new parameter to the PS_FSPV_PARAMETER_VALUES Table:

```
INSERT INTO PS_FSPV_PARAMETER_VALUES
(FSPV_FSRA_ID,FSPV_FSPG_ID,FSPV_FSRP_ID,
FSPV_CTX_ID,FSPV_VALUE,FSPV_PROMPT_TEXT) VALUES
(100000, NULL, 100086, 10, 'Scheduler', NULL)
```


Updating the .xml file with the report

When you create a Customized report, Report Manager creates two files: a Microsoft Excel file that contains the report details, and a .xml file that contains the report definition. This procedure provides step-by-step instructions for updating the definitions of the query tag in the .xml file. See [About creating a customized report](#) on page 50 and [About adding a new property to the Report Parameters table](#) on page 47. The following example shows the query tag before modification.

```
<query updated_manually="false">SELECT SQPN_STRING_VALUE "SQProgram"
FROM (SELECT SQSS_PROGRAM_ID FROM PW_SQSS_SESSIONS_STATS_D,
PS_INAP_APP_TIER, PW_SQPN_PROGRAM_NAMES_N, PS_INCE_INSTANCE,
PS_INEN_ENVIRONMENT
WHERE SQSS_PWII_INSTANCE_ID=INCE_ID AND INCE_INAP_ID=INAP_ID AND
INAP_INEN_ID=INEN_ID AND SQSS_PROGRAM_ID = SQPN_ID AND SQPN_ID
<>0 AND SQPN_STRING_VALUE LIKE 'My Program' AND INEN_ID = 10
GROUP BY SQPN_STRING_VALUE, SQSS_PROGRAM_ID)intab,
PW_SQPN_PROGRAM_NAMES_N WHERE SQSS_PROGRAM_ID = SQPN_ID AND SQPN_ID
<>0</query>
```

The following example shows the query tag after modification:

```
<query updated_manually="true">SELECT SQPN_STRING_VALUE "SQProgram"
FROM (SELECT SQSS_PROGRAM_ID FROM PW_SQSS_SESSIONS_STATS_D,
PS_INAP_APP_TIER, PW_SQPN_PROGRAM_NAMES_N, PS_INCE_INSTANCE,
PS_INEN_ENVIRONMENT
WHERE SQSS_PWII_INSTANCE_ID=INCE_ID AND INCE_INAP_ID=INAP_ID AND
INAP_INEN_ID=INEN_ID AND SQSS_PROGRAM_ID = SQPN_ID AND SQPN_ID
<>0 AND SQPN_STRING_VALUE LIKE '$P{ProgramFilter}' AND INEN_ID = 10
GROUP BY SQPN_STRING_VALUE, SQSS_PROGRAM_ID)intab,
PW_SQPN_PROGRAM_NAMES_N WHERE SQSS_PROGRAM_ID = SQPN_ID AND SQPN_ID
<>0</query>
```

To update the report's .xml file

1. Open the .xml file of your Customized report from the following folder in the server where Report Manager FocalPoint is installed:

```
<precise_root>/products/foresight/etc/reports/udr/storage
```

The file name format consists of the Context Identifier (FSPV_CTX_ID) concatenated with the Report Identifier (FSRP_ID). For example, in the 10_100086.xml Customized report file, the context identifier is 10 and the report identifier is 100086 (you can also check the file's modification date to identify your report's .xml file).

2. In the .xml file, search for the fixed value that you have specified in the filter conditions when creating the report.

The fixed value can be found in both the filter and query tags. Update the query tag as follows:

- Change the value of the **updated_manually** attribute to `true`.
- Replace the fixed value of the filter condition with the string '\$P{parameter_name}', where parameter_name is the new property name (FSRA_NAME) that you defined in the Report Parameters Table.

3. Save the .xml file.
4. Click **Scheduling** or **View** followed by **Rerun** to rerun the report.

Do not rerun the report by clicking **Edit**.

Filling in gaps in a customized report

Retrieving information collected over time from the Performance Management Database may result in missing information. These are specific timestamps for which no data is available. For example, an Oracle program may not include data for a specific timestamp, if the program was not invoked during the interval that started with this timestamp. In this case, you may want, for example, to fill in the number of executions counter with a zero value.

To fill in gaps in a customized report

1. Create a customized report
2. Modify the report's .xml file

After filling in gaps in a customized report, don't click `Edit` to modify the report. If after filling in gaps in a customized report you mistakenly click `Edit`, Report Manager will override your modifications. If you must edit the Customized report, fill in the gaps again so as to restore the properties you modified.

About creating a customized report

The selected data fields in the report must contain one or more identifiers and a single timestamp data field, which must be a full timestamp data field and not a partial time data field, such as Date and Month.

You should also sort the identifiers followed by the timestamp. The timestamp must be sorted in ascending order, while the identifiers can be sorted in ascending or descending order.

After completing report settings, execute the customized report and then save it. Report Manager creates two files, the report's Microsoft Excel file and `.xml` file. See [Creating customized reports](#) on page 20.

Modifying the `.xml` file with the report

When you create a Customized report, Report Manager creates two files: a Microsoft Excel file that contains the report details, and a `.xml` file that contains the report definition.

The following example shows the report's `.xml` file before modification. The bold text indicates the value that should be replaced.

```
<report>
<subreport name="default_query" sum_level="Daily" id="47">
<selection>
<object id="4000" name="Instance" />
<object id="4003" name="User" />
<object id="21" name="Timestamp" />
<object id="4181" name="Sessions" />
</selection>
<order>
  <object id="4000" name="Instance" order="asc" />
  <object id="4003" name="User" order="asc" />
  <object id="21" name="Timestamp" order="asc" /> </order>
</filter />
<query firstLoad="true" updated_manually="false">SELECT intab.c2
"SQInstance", SQUN_STRING_VALUE "SQUser", intab.c3 "Timestamp",
intab.c0 "SQSessions" FROM (SELECT SQSS_TIMESTAMP c3, SQSS_USER_ID ,
SUM(SQSS_NUM_OF_ENDED_SESSIONS_SUM) c0, INCE_NAME c2
FROM PW_SQSS_SESSIONS_STATS_D, PS_INAP_APP_TIER, PS_INCE_INSTANCE,
PS_INEN_ENVIRONMENT WHERE SQSS_PWII_INSTANCE_ID=INCE_ID AND
INCE_INAP_ID=INAP_ID AND INAP_INEN_ID=INEN_ID AND
SQSS_PWII_INSTANCE_ID=INCE_ID AND INCE_INTE_CODE='SQ' AND
INCE_DELETED='F' AND INEN_ID = 10 GROUP BY INCE_NAME,
SQSS_USER_ID, SQSS_TIMESTAMP) intab, PW_SQUN_USER_NAMES_N
WHERE SQSS_USER_ID = SQUN_ID AND SQUN_ID <>0 ORDER BY
"SQInstance" asc, "SQUser" asc, "Timestamp" asc
</query>
</subreport>
</report>
```

The following example shows the report's `.xml` file after modification:

```
<report>
<subreport name="default_query" sum_level="Daily" id="47">
<selection>
<object id="4000" name="Instance" />
<object id="4003" name="User" />
<object id="21" name="Timestamp" />
<object id="4181" name="Sessions" />
</selection>
<order>
<object id="4000" name="Instance" order="asc" />
<object id="4003" name="User" order="asc" />
<object id="21" name="Timestamp" order="asc" />
</order>
</filter />
<query firstLoad="true" updated_manually="true">SELECT intab.c2
"SQInstance", SQUN_STRING_VALUE "SQUser", intab.c3 "Timestamp",
intab.c0 "SQSessions"
```

```
FROM (SELECT SQSS_TIMESTAMP c3, SQSS_USER_ID,
SUM(SQSS_NUM_OF_ENDED_SESSIONS_SUM)
c0, INCE_NAME c2 FROM PW_SQSS_SESSIONS_STATS_D, PS_INAP_APP_TIER,
PS_INCE_INSTANCE, PS_INEN_ENVIRONMENT WHERE
SQSS_PWII_INSTANCE_ID=INCE_ID AND
INCE_INAP_ID=INAP_ID AND INAP_INEN_ID=INEN_ID AND
SQSS_PWII_INSTANCE_ID=INCE_ID AND INCE_INTE_CODE='SQ' AND
INCE_DELETED='F' AND INEN_ID = 10 GROUP BY INCE_NAME, SQSS_USER_ID,
SQSS_TIMESTAMP)intab, PW_SQUN_USER_NAMES_N
WHERE SQSS_USER_ID = SQUN_ID AND SQUN_ID
<>0 ORDER BY "SQInstance" asc, "SQUser" asc, "Timestamp" asc
</query> <fill_holes fill="true">
  <object id="4181" name="Sessions" fill="const" const="666" />
  <timeframe start="2004-05-20 00:15:23.11"
end="2004-06-10 22:35:23.11" />
</fill_holes>
</subreport>
</report>
```

This procedure provides step-by-step instructions for updating the report's .xml file, including updating the query tag and setting the fill_holes tag.

To modify the report's .xml file

1. Open the .xml file of your Customized report from the following folder in the server, where Report Manager FocalPoint is installed:

```
<precise_root>/products/foresight/etc/reports/udr/storage
```

The file name format consists of the Context Identifier (FSPV_CTX_ID) concatenated with the Report Identifier (FSRP_ID).

For example, in the 10_100086.xml file, the context identifier is 10 and the report identifier is 100086. (You can also check the file's modification date to identify your report's .xml file.)

2. In the .xml file, do the following:
 - Search for the query tag.
 - Change the value of the updated_manually attribute to True.
 - Add the fill_holes tag after the query tag at the same level.
3. Set the fill attribute of the fill_holes tag to True to enable the data completion feature.
4. Add the object tag under the fill_holes tag level and set its attributes.

For example, <object id="4000" name="Instance" fill="const" const="666"/>

id Specifies the object identifier.
Set to one of the objects ID defined in the selection tag at the beginning of the report's .xml file.

name Specifies the data field name.
This attribute is optional. You may set this attribute to add clarity to the report.

fill Specifies the method of data completion.

Possible values include:

- **const**
The data is completed with a constant value. In this case, you must also set the const attribute.
- **prev**
The data is completed with the value of the previous timestamp.
- **next**
The data is completed with the value of the next timestamp.

const Specifies the constant value that completes the missing data.
 Set the `const` attribute if you set the `fill` attribute to `const`.
 By default, Report Manager fills in counters with zeros, and strings with nulls. You can set the `const` attribute, together with the `prev` or `next` fill types, to specify a different value other than the default value (zero or null), so that Report Manager also completes the margins, where there is no previous or next value.

5. Add the `timeframe` tag under the `fill_holes` tag level and set its attributes.

The `timeframe` tag defines the time frame for which the data completion applies. You can set the time frame using one of the following methods:

- Absolute date range.

Set the start and end attributes. For example, `<timeframe start="2004-05-20 00:15:23.11" end="2004-06-10 22:35:23.11" />`.

- Relative date time.

Set the `lastNType`, `lastNStart` and `lastNLen` attributes.

start Specifies the start date and time.
 The format is: `yyyy-mm-dd hh24:mi:ss.ms`
 You must apply the entire format, except the `ms` (milliseconds), which is optional. For example:
`<timeframe start="2004-05-15 18:30:15.00" end="2004-05-20 23:59:59.59"/>`

end Specifies the end date and time. See the format of the start attribute.

lastNType Specifies the time units for the `lastNStart` and `lastNLen` attributes.

Possible values include:

- H - Hours
- D - Days
- W - Weeks
- M - Months

Set a value that is not lower than the report summary level resolution (that is, the report summary level must be at least as the `lastNType` value).

For example: `lastNType="D"` can be used with a report summary level of H (hours).

lastNStart Specifies the start of the time frame counting back from the current time (current time is always 0).

For example, if `lastNType="D"`, then a `lastNStart="0"` signifies that the start time of the time frame is today.

If `lastNStart="10"` and `lastNLen = "5"`, this signifies that the start time of the time frame was 15 days ago.

lastNLen Specifies the length of the time frame counting back from the `lastNStart` time.
 The following example defines a start time of 15 days ago and end time of 10 days ago.

`<timeframe lastNType="D" lastNStart="10" lastNLen="5"/>`

6. Save the `.xml` file.

7. Rerun the report by doing one of the following:

- Click **Scheduling**.
- Click **View** and then **Rerun**.

Don't click **Edit** to rerun the report

About importing and exporting customized reports

The following section describes how to export a customized report from one Precise system and import it to another Precise system.

NOTE Before you begin, find the source Tier ID and the destination Tier ID in the PS_INAP_APP_TIER or PS_FSEC_ENV_CONTENTS tables.

To export/import a customized report

1. In the source Precise system, create the `ForesightExportRequest.xml` file in the Precise root folder. For example:

```
<root name="Export">
  <report-name>In Sybase Breakdown</report-name>
  <apptier>1201</apptier>
</root>
```

NOTE The report-name value is case-sensitive.

2. Run the following command:


```
infra\bin\psin_http_requestor -dp -x fs-export-import FS < ForesightExportRequest.xml > out.xml
```
3. Check the `out.xml` file in the `<precise_root>`, to view the status.
4. Copy the exported report's `.zip` file from the source Precise root folder to the destination Precise installation in:

```
<precise_root>\distribution
```

NOTE The `.zip` file name consists of `<technology code>_<report name>` (`SY_In Sybase Breakdown.zip`). For cross-Tiers the Tech code is `ALL` (`ALL_CA_Rep.zip`).

5. In the destination Precise system, create the `ForesightImportRequest.xml` file in the Precise root folder. For example:

```
<root name="Import">
  <report-name>SY_In Sybase Breakdown.zip</report-name>
  <apptier>1174</apptier>
</root>
```

6. Run the following command:


```
infra\bin\psin_http_requestor -dp -x fs-export-import FS < ForesightImportRequest.xml > out.xml
```
7. Check the `out.xml` file in the Precise root folder, to view the status.

About copying customized reports

The following section describes how to copy customized reports from one Tier to another with the CLI utility. To copy a customized report

1. Run the following commands from the Precise root folder: UNIX

```
infra/bin/psin_cli.sh -i3-user <user> {-i3-clear-password <i3_clear_password> |
-i3-encrypted-password <i3_encrypted_password>}
-action fs-copy-foresight-custom-report
-parametersfile <parameters-file.xml>
```

Windows

```
infra\bin\psin_cli.bat -i3-user <user> {-i3-clear-password <i3_clear_password> |
-i3-encrypted-password <i3_encrypted_password>}
-action fs-copy-foresight-custom-report
-parametersfile <parameters-file.xml>
```

Example:

```
infra\bin\psin_cli.bat -i3-user admin -i3-clear-password admin -action fs-copy-foresight-custom-report -
parametersfile my_parameters_file.xml
```

2. Create the `parameters-file.xml` file in the Precise root folder and verify that it contains the following:

```
<parameters>
<parameter name="source-environment" value="Default"/>
<parameter name="source-apptier" value="SQL Server"/>
<parameter name="source-report" value="My_report"/>
<parameter name="destination-environment" value="Default"/>
<parameter name="destination-apptier" value="SQL Server"/>
</parameters>
```

and verify that the information in the following notes is applied.

NOTE The parameter values are case-sensitive.

NOTE The "source-report" name must not contain blanks.

NOTE The new report name will have a suffix with a serial number (for example: "My_report_1").

NOTE To copy a Cross-Tier report, you should use the "Cross-Tier" string for both "source-apptier" and "destination-apptier" parameters.

Database mapping and definition files

This appendix includes the following topics:

- [About database mapping and definition files](#)
- [About the db.xml file](#)
- [About the def.xml file](#)
- [About the tree.xml file](#)

About database mapping and definition files

An understanding of database mapping and definition files is essential if you want to add internal proprietary information such as, user name and telephone number, to the reports generated by Report Manager.

NOTE Entering invalid strings into these files can cause the customized reporting mechanism to fail. Any modifications made to the xml files falls under the sole responsibility of the customer. Before modifying these files verify that you are experienced with the SQL language, familiar with XML, and that you have fully read and understood this section.

The database mapping and definitions for customized reports are stored in the following files:

db.xml	Describes the database mapping.
def.xml	Defines dependency conditions.
tree.xml	Defines the hierarchical structure of the Data Tree.

These files are found located in: `<precise_root>/products/foresight/etc/reports/udr`

About the db.xml file

The `db.xml` file defines all Database Tables and Join conditions between tables, which are used by the customized reporting mechanism.

The `db.xml` file uses the following tags to define the database mapping:

- [About the Table tag](#)
- [About the Join tag](#)
- [About the Related tag](#)

About the Table tag

The Table tag defines the table and its attributes, as described in the following table. The Table tag is a top level element in the db.xml file.

Example:

```
<table name="PS_INEN_ENVIRONMENT" type="lookup"/>
<table name="PW_SPTS_TRAN_STAT_$sumlvl" type="statistics"
summary_levels="T,H,D,W,M" >
</table>
```

The following table describes the table tag attributes.

Table B-1 Table tag attributes

Attribute	Description
name	Specifies the table name as defined in the database schema. For a statistics table type, you can use the \$sumlvl variable. In this case, you must set also the summary_levels attribute. When executing the customized report, the \$sumlvl variable uses the value that you set in the summary_levels attribute.
type	Defines the table type, as follows: <ul style="list-style-type: none"> • Lookup - a lookup table maps between entity identifiers and entity names. • Statistics - a statistics table that includes statistics performance information and is summarized into various time intervals (such as, hourly, daily, and weekly). • History - history table that includes statistics performance information, but is not summarized into various time intervals.
summary_levels	This attribute is only relevant for a statistics table type. It lists the available summary levels for this table. These values include: <ul style="list-style-type: none"> • T – Time slice • H – Hourly • D – Daily • W – Weekly • M – Monthly
time_field	Defines the timestamp column name. Use this attribute only if the column name is different than the standard xxxx_timestamp column name, where xxxx indicates the name of the table (abbreviated).
noteable-name	A flag that enables or disables Join conditions. Set this flag to true for tables that contain Join conditions written in Java code. In this case, replace the table-name with the table name whose Join conditions you want to disable. For example, noPS_INSM_SERVER_MACHINE="true", will disable the Join conditions in table PS_INSM_SERVER_MACHINE.
alias	Defines an alias name for a table that is used frequently. For example, you can use the same lookup table again for different purposes.

About the Join tag

The Join tag is a second level element in the db.xml file, and is part of the table definition. The Join tag defines Join conditions between the current table and other specified tables. The Join tag defines a one-way connection from the current table to another table, but not in the opposite direction. Usually, lookup tables do not include a Join definition. Instead, they will be referenced inside a Join definition of Statistics Tables. You should not define a Join between two Statistics Tables, unless it is a One-to-one Join. The Join predicates are included in the tag contents.

Example:

```
<table name="PS_INCE_INSTANCE" type="lookup">
  <join toTable="PS_INEN_ENVIRONMENT"
    otherTables="PS_INAP_APP_TIER">
    INEN_ID = INAP_INEN_ID AND INCE_INAP_ID = INAP_ID
  </join>
</table>
```


The following table describes the Join tag attributes.

Table B-2 Join tag attributes

Attribute	Description
toTable	Specifies the table that is joined with the current table.
otherTables	Specifies the name of a connection table. This attribute is used if the Join is performed through another table.
normalization_field	<p>Defines a Join condition for a Lookup Table, whose name ends with <code>_N</code>. You can use the <code>normalization_field</code> attribute to specify the name of the normalization column. In this case, the select will include a 'group by' with the <code>normalization_field</code> column and will be wrapped by another select, which joins with the Lookup Table specified in the <code>toTable</code> attribute. The reasons for applying this method are:</p> <ul style="list-style-type: none"> Performance - the Join with the Lookup Table is done only in the wrapping select and not in the inner select which retrieves more rows. Database limitations - in Oracle it is impossible to apply any option on a column of a long data type, or to use it in the 'group by'.

About the Related tag

The Related tag is a second level element in the `db.xml` file, and is part of the table definition. This tag is only relevant for lookup tables. It defines the statistics table used for retrieving entity identifiers that depend on a specified condition (see the `depend` attribute).

A table can include more than one Related tag. Report Manager uses the first related table whose dependency condition is met. It is possible to define the last line without a `depend` attribute, so that if no previous dependency was met, this table is always selected. If you want an error message to be displayed when no dependency conditions are met, do not specify a related table with no dependency.

The Performance Management Database includes Snapshot Tables (for details about these tables, see the Performance Management Database documentation). To achieve optimal performance when using Snapshot Tables, it is recommended to define the smallest Snapshot Tables first, followed by the larger ones.

For example:

```
<table name="PW_SPLD_LOCALE_DIM" type="lookup">
<related table="PW_SPTS_TRAN_STAT_LS_$sumlvl" depend="loc srv snapshot"/>
<related table="PW_SPTS_TRAN_STAT_AOLS_$sumlvl" depend="org app loc srv snapshot"/>
</table>
```

The following table describes the Related tag attributes.

Table B-3 Related tag attributes

Attribute	Description
table	Defines the name of the Statistics Table.
depend	Specifies the dependency condition identifier. This identifier is defined in the <code>def.xml</code> file. If the dependency condition is met, then the table specified in the <code>table</code> attribute is used to retrieve the list of entity identifiers.

About the def.xml file

The `def.xml` file is a configuration file that defines dependency conditions using a single tag called `dependency`. These definitions are used by the `depend` attribute of the `db.xml` file.

A special subset of these definitions are dependency conditions for Snapshot Tables. A Snapshot Table is used if all the selected data fields in the query exist in this snapshot.

The `dependency` tag contains a single attribute, called `name`. The value of this attribute is a character string that identifies the dependency. This character string must be unique in the `def.xml` file.

The dependency tag format contains dependency conditions, as follows:

```
<![CDATA[unary-operator{field-id,...}] logical-operator ...>
```

where:

- The `field-id` specified in the brackets is a numeric identifier of a data field in the Data Tree as defined in the `tree.xml` file.

You can specify one or more data field identifiers inside the brackets according to the specified unary-operator.

- The `unary-operator` defines the relation between the data field identifiers specified in brackets and the SELECT list in the query.

The value can be one of the following:

- `a` - all. All the data field identifiers must exist in the SELECT list.
- `na` - not all. All of the data field identifiers must not exist in the SELECT list.
- `o` - one. At least one data field identifier must exist in the SELECT list.
- `no` - not one. At least one data field identifier must not exist in the SELECT list.
- The logical-operator between dependency conditions can be one of the following:
 - `|` – OR
 - `&` – AND

Example:

```
<dependency name="form appl user snapshot and user"><![CDATA[na{2003,4} & a{2002}]]></dependency>  
<dependency name="oa user"><![CDATA[a{2002}]]></dependency>
```

About the tree.xml file

The `tree.xml` file defines the hierarchical structure of the Data Tree. Each data field that is a leaf node, includes a unique numeric identifier. The data field identifiers are used in the `def.xml` file. For this reason you should not modify the existing data field identifiers, but you can add additional data field identifiers to the Data Tree. If you add new data field identifiers to the Data Tree, you must also add new data fields to the `def.xml` file.

The `tree.xml` file uses the following tags:

- Group
- Object
- Objref
- Alt

About the Group tag

This tag defines a group that is a non-leaf node. A group can contain data fields (leaf nodes) and other groups. Example:

```
<group name="General" desc="0">  
<group isApptier="true" tech="SP" name="SP" desc="17" />  
</group>  
<group isApptier="true" tech="ST" installCode="SQ-SQCOL" name="ST" desc="1076" />
```

The following table describes the Group tag attributes.

Table B-4 Group tag attributes

Attribute	Description
name	Defines a string representing the group name. Report Manager translates this string to a text defined in the <code>udrTxtResource.xml</code> resource file, and then displays the translated text in the Data Tree.
desc	Defines a unique numeric identifier that represents the description of the group. Report Manager translates this number into text defined in the <code>udrTxtResource.xml</code> resource file. Before adding a new identifier to the file, you must verify that this number is not already in use. (The numeric identifier at the end of the file list is not necessarily the highest one.)
isApptier	Defines whether the group is an Tier group or not.
tech	Defines a string that represents the technology of the Tier group. Report Manager translates this string into text defined in the <code>udrTxtResource.xml</code> resource file.
installCode	Defines a string that represents the installation code. Its format is a combination of the product code and init code. For example, <code>IS-FP</code> , where <code>IS</code> is the product code and <code>FP</code> is the init code. Report Manager validates the installation code using the Infrastructure Installation Table named <code>PS_INLL_INSTALLATION</code> , to determine whether or not to display this group in the specific Precise application. (The Infrastructure Installation Table contains information regarding the application's system installations.)

About the Object tag

This tag defines a data field that is displayed as a leaf node in the Data Tree. It can be a column from a table or a formula.

Example:

```
<object id="1" type="id" data_type="string" name="Environment" desc="1">
<object id="1000" type="instance_id" name="SPSystem" desc="20">
<object id="21" type="time" data_type="time" xls_format="m/d/yy h:mn"
name="Timestamp" desc="7">
```

The following table describes the Object tag attributes.

Table B-5 Object tag attributes

Attribute	Description
id	Defines a numeric identifier of the data field.
type	Defines a string that represents the type of object. The value can be one of the following: <ul style="list-style-type: none"> id - identifier data field time - time data field instance_id - instance data field cnt - counter data field (in this case, also defined as an Aggregate function)
data_type	Describes the type of each data field. Default values include: <ul style="list-style-type: none"> id and instance_id - string cnt - double time - time <p>It is possible to override the default type of the data field type by defining the <code>data_type</code> attribute. Possible values include: string, time, int, and float.</p>

Attribute	Description
xls_format	<p>The format of each type of data field has a default value used by Microsoft Excel. By default, Microsoft Excel uses no special format for strings.</p> <p>For numeric data types, Microsoft Excel uses the floating decimal point format.</p> <p>For date/time data types, Microsoft Excel uses the numeric format. You can override the default format by defining the xls_format attribute.</p> <p>For example, you can override the format of the date/time data type with one of the following formats: m/d/yy, d-mmm-yy, d-mmm, mmm-yy, h:mm AM/PM, h:mm:ss AM/PM, h:mm, h:mm:ss, and m/d/yy h:mm.</p>
name	<p>Defines a string that represents the data field name.</p> <p>Report Manager translates this string into a text defined in the <code>udrTxtResource.xml</code> resource file, and then displays the translated text in the Data Tree.</p>
desc	<p>Defines a unique numeric identifier that represents the description of the data field.</p> <p>Report Manager translates this number into a text defined in the <code>udrTxtResource.xml</code> resource file.</p> <p>Before adding a new identifier to the file, you must verify that this number is not already in use. (The numeric identifier at the end of the file list is not necessarily the highest one.)</p>
installApp	<p>Defines a string that represents a Tier code.</p> <p>A specific data field may depend on a specific Tier, so if the Tier is not installed in the current application, the data field will not be displayed in the Data Tree.</p> <p>The Tier installation is verified against the Infrastructure Tier Table named <code>PS_INAP_APP_TIER</code>.</p>
installCode	<p>Defines a string that represents the installation code.</p> <p>Its format is a combination of the product code and init code.</p> <p>For example, <code>IS-FP</code>, where <code>IS</code> is the product code and <code>FP</code> is the init code.</p> <p>Report Manager validates the installation code using the Infrastructure Installation Table named <code>PS_INLL_INSTALLATION</code>, to determine whether or not to display this group in the specific Precise application. (The Infrastructure Installation Table contains information regarding the application's system installations.)</p>
includeQ	<p>Defines a new query used if the object returns no data (in this case, the data field is not displayed).</p> <p>For example:</p> <pre><object id="4026" type="id" name="SQSiebelView" desc="446" includeQ="SELECT INLI_INCE_ID FROM PS_INLI_LITE_INSTANCE, PS_FSEC_ENV_CONTENTS WHERE INLI_INTE_CODE = 'SQ' AND INLI_INCE_ID = FSEC_INSTANCE_ID AND INLI_DELETED = 'F' AND FSEC_ENV_ID = ?"></pre>

About the Objref tag

This tag defines a reference to another Object that is already defined in the Data Tree definition. Use the Objref tag if the data field you want to add to the Data Tree has the same attributes as an existing object (instead of repeating the entire definition of the object).

For example, the Instance ID object (data field) may appear under various groups of a Tier.

The following example creates a reference to an object 3 (Server) definition:

```
<objref id="3" />
```

About the Alt tag

The meaning of a data field can depend on its context, that is, on the other fields that are selected with it. The Alt tag defines an alternative selection of expression for its data field, based on the other fields selected. When the data field is selected, Report Manager scans the alternatives in the same order as they were listed in the file, and chooses the first alternative whose dependency is met.

Example:

```
<object id="5050" type="cnt" name="ORInOracleTimeSum" desc="825">
  <alt select="SUM(ORSS_IN_ORACLE_TIME_SUM)"
    table="PW_ORSS_STATEMENTS_STATS_$sumlvl" depend="or stmt no sess" />
  <alt select="SUM(ORSA_IN_ORACLE_TIME_SUM)"
    table="PW_ORSA_STMT_APPL_STATS_$sumlvl" depend="or stmt and sess" />
  <altselect="SUM(ORAS_IN_ORACLE_TIME_SUM)"
    table="PW_ORAS_APPLICATION_STATS_$sumlvl" />
</object>
```

In this example, if the user selects the In Oracle time data field:

- If the `or stmt no sess` dependency is met (defined in the `def.xml` file), then Report Manager selects the `SUM(ORSS_IN_ORACLE_TIME_SUM)` expression from the `PW_ORSS_STATEMENTS_STATS_$sumlvl` Table.
- If the first dependency is not met, and the `or stmt and sess` dependency is met (defined in the `def.xml` file), then Report Manager selects the `SUM(ORSA_IN_ORACLE_TIME_SUM)` expression from the `PW_ORSA_STMT_APPL_STATS_$sumlvl` Table.
- If the second dependency is not met, Report Manager selects the `SUM(ORAS_IN_ORACLE_TIME_SUM)` expression from the `PW_ORAS_APPLICATION_STATS_$sumlvl` Table.

The following table describes the Alt tag attributes.

Table B-6 Alt tag attributes

Attribute	Description
select	Defines a table column name or an expression. For a counter (<code>type="cnt"</code>), an expression with an aggregated function is required. The expression can include the <code>\$timeFld</code> string, which Report Manager replaces with the timestamp column of the table used in the query. In this case, the <code>table</code> attribute must be empty.
table	Specifies the name of the table as defined in the Performance Management Database.
depend	Specifies the name of the dependency condition as defined in the <code>def.xml</code> file. If the specified condition is met, Report Manager uses the alternative SELECT.
otherTables	Defines a name of a connection table. If you need to join one table to other tables for this data field only, you can specify its conditions in the <code>alt</code> tag and it is therefore not necessary to define this Join in the <code>db.xml</code> file.

Report Manager report sets

This appendix includes the following topics:

- [Report Manager Report sets](#)
- [Cross-Tiers Reports](#)
- [Oracle](#)
- [J2EE](#)
- [Web](#)
- [Microsoft .NET](#)
- [SQL Server](#)
- [SAP](#)
- [Sybase](#)
- [Tuxedo](#)
- [WebSphere MQ](#)
- [Oracle Applications](#)
- [OS](#)
- [EMC Storage](#)
- [Citrix](#)
- [DB2](#)

Report Manager Report sets

Report Manager is a reporting tool that queries the Performance Management Database and displays the results in reports that are generated on a scheduled basis or on demand.

Report Manager provides pre-defined reports divided as follows:

- Availability
- Capacity planning
- Exceptions
- Load Balancing
- Profile

Reports of an Tier-level, for example, provide a performance comparison between the various Tiers of your application.

Reports of an Instance-level provide information about the complete instance, that is the whole system, such as Web Server, an Oracle Instance, or a (J2EE) Java Virtual Machine.

Cross-Tiers Reports

Cross-Tier reports contain data for several Tiers, for example, Cross-Tiers Availability report. Application reports are also available for Cross-Tiers reports only.

The Report Manager Cross-Tiers Reports are as follows:

- Availability
- Load Balancing
- Application

Oracle

The following table lists report sets for Oracle.

Table C-1 Report sets for Oracle

Report set	Report name
Availability	Instances
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Actions
Exceptions	Applications (Oracle Apps)
Exceptions	File Systems
Exceptions	Files
Exceptions	Forms (Oracle Apps)
Exceptions	Functions (SAP R/3)
Exceptions	Host Users
Exceptions	Instances
Exceptions	Logical Volume
Exceptions	Machines
Exceptions	Modules
Exceptions	Objects
Exceptions	Physical Devices
Exceptions	Programs
Exceptions	Rarely Used and Unused Objects
Exceptions	Requests (Oracle Apps)
Exceptions	Screens (SAP R/3)
Exceptions	Statements
Exceptions	Storage Devices
Exceptions	Transactions (SAP R/3)
Exceptions	Users
Profile	Actions
Profile	Applications (Oracle Apps)
Profile	Database Size
Profile	File Systems
Profile	Files

Report set	Report name
Profile	Forms (Oracle Apps)
Profile	Functions (SAP R/3)
Profile	Host Users
Profile	Instance Statistics
Profile	Instances
Profile	Logical Volume
Profile	Machines
Profile	Modules
Profile	Objects
Profile	Physical Devices
Profile	Programs
Profile	RAC Databases
Profile	Requests (Oracle Apps)
Profile	Schema Changes
Profile	Screens (SAP R/3)
Profile	Statements
Profile	Storage Devices
Profile	Transactions (SAP R/3)
Profile	Users

J2EE

The following table lists report sets for J2EE.

Table C-2 Report sets for J2EE

Report set	Report name
Availability	JVMs
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Entry Points
Exceptions	JVMs
Profile	Entry Points
Profile	JVMs

Web

The following table lists report sets for Web.

Table C-3 Report sets for Web

Report set	Report name
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Instances
Exceptions	Locations
Exceptions	Transactions
Exceptions	Users
Profile	Instances
Profile	Locations
Profile	Transactions
Profile	Users

Microsoft .NET

The following table lists report sets for Microsoft .NET.

Table C-4 Report sets for Microsoft .NET

Report set	Report name
Availability	CLRs
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Entry Points
Exceptions	CLRs
Profile	Entry Points
Profile	CLRs

SQL Server

The following table lists report sets for the SQL Server.

Table C-5 Report sets for SQL Server

Report set	Report name
Availability	Databases
Availability	Instances
Availability	Overview
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Ad-hoc Batches
Exceptions	Applications (COM+)
Exceptions	Client Machines (PeopleSoft)
Exceptions	Components (COM+)
Exceptions	Databases
Exceptions	DB Files
Exceptions	Instances
Exceptions	Logins
Exceptions	Machines
Exceptions	Methods (COM+)
Exceptions	Objects
Exceptions	Physical Disks
Exceptions	Programs
Exceptions	Rarely Used and Unused Objects
Exceptions	Siebel Views
Exceptions	Statements
Exceptions	Storage Devices
Exceptions	Stored Procedures
Exceptions	Transactions (SAP)
Exceptions	User IDs (PeopleSoft)
Exceptions	Users
Exceptions	Work Types
Profile	Ad hoc Batches
Profile	Applications (COM+)
Profile	Client Machines (PeopleSoft)
Profile	Components (COM+)
Profile	Database Definition Changes
Profile	Database Size
Profile	Databases
Profile	DB Files
Profile	Instance Statistics

Report set	Report name
Profile	Instances
Profile	Logins
Profile	Machines
Profile	Methods (COM+)
Profile	Objects
Profile	Physical Disks
Profile	Programs
Profile	Schema Changes
Profile	Siebel Views
Profile	Statements
Profile	Storage Devices
Profile	Stored Procedures
Profile	Transactions (SAP)
Profile	Unnecessary Indexes
Profile	User IDs (PeopleSoft)
Profile	Users
Profile	Work Types

SAP

The following table lists report sets for SAP Systems.

Table C-6 Report sets for SAP Systems

Report set	Report name
Availability	Application Servers
Availability	Locales
Availability	Organizations
Availability	Overview
Availability	Systems
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Application Servers
Exceptions	Applications
Exceptions	Locales
Exceptions	Organizations
Exceptions	Systems
Exceptions	Transactions
Exceptions	Users
Profile	Application Servers
Profile	Applications
Profile	Locales
Profile	Organizations

Report set	Report name
Profile	Response Time
Profile	Systems
Profile	Transactions
Profile	Users

Sybase

The following table lists report sets for Sybase.

Table C-7 Report sets for Sybase

Report set	Report name
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Ad hoc Batches
Exceptions	Database Devices
Exceptions	Databases
Exceptions	Instances
Exceptions	Logins
Exceptions	Machines
Exceptions	Programs
Exceptions	Statements
Exceptions	Stored Procedures
Exceptions	Users
Profile	Ad hoc Batches
Profile	Database Devices
Profile	Database Size
Profile	Databases
Profile	Instances
Profile	Logins
Profile	Machines
Profile	Programs
Profile	Statements
Profile	Stored Procedures
Profile	Users

Tuxedo

The following table lists report sets for Tuxedo.

Table C-8 Report sets for Tuxedo

Report set	Report name
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Client IPs
Exceptions	Instances
Exceptions	Server Processes
Exceptions	Services
Exceptions	Users
Profile	Client IPs
Profile	Instances
Profile	Server Processes
Profile	Services
Profile	Users

WebSphere MQ

The following table lists report sets for WebSphere MQ

Table C-9 Report sets for WebSphere MQ

Report set	Report name
Availability	Instances
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Instances
Exceptions	Queues
Exceptions	Reader Applications
Exceptions	Users
Exceptions	Writer Applications
Profile	Instances
Profile	Queues
Profile	Reader Applications
Profile	Users
Profile	Writer Applications

Oracle Applications

The following table lists report sets for Oracle Applications.

Table C-10 Report sets for Oracle Applications

Report set	Report name
Availability	Instances
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Applications
Exceptions	Client IPs
Exceptions	Forms
Exceptions	Instances
Exceptions	Locations
Exceptions	Users
Profile	Applications
Profile	Client IPs
Profile	Forms
Profile	Instances
Profile	Locations
Profile	Users

OS

The following table lists report sets for OS (Operating System).

Table C-11 Report sets for OS

Report set	Report name
Load Balancing	OS Counters
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Commands
Exceptions	Programs
Exceptions	Servers
Exceptions	Users
Profile	Commands
Profile	Programs
Profile	Servers
Profile	Users

EMC Storage

The following table lists report sets for EMC Storage.

Table C-12 Report sets for EMC Storage

Report set	Report name
Capacity Planning	Long Term
Capacity Planning	Short Term
Exceptions	Devices
Exceptions	Front End Directors
Exceptions	Units
Profile	Devices
Profile	Front End Directors
Profile	Units

Citrix

The following table lists report sets for Citrix Servers.

Table C-13 Report sets for Citrix Servers

Report set	Report name
Exceptions	Citrix Locations
Exceptions	Citrix Users
Exceptions	Client IPs
Profile	Published Applications
Profile	Citrix Locations
Profile	Citrix Users
Profile	Client IPs
Profile	Published Applications

DB2

The following table lists report sets for DB2 Servers.

Table C-14 Report sets for DB2 Servers

Report set	Report name
Profile	Database Partitions
Profile	Programs
Profile	Statements
Profile	Machines
Profile	Modules
Profile	Actions
Profile	OS Users
Profile	DB Users
Profile	Work Types
Exception	Database Partitions
Exception	Programs

Report set	Report name
Exception	Statements
Exception	Machines
Exception	Modules
Exception	Actions
Exception	OS Users
Exception	DB Users
Exception	Work Types

Importing, exporting and copying your customized reports

This appendix includes the following topics:

- [About importing and exporting customized reports](#)
- [About copying customized reports](#)

About importing and exporting customized reports

The following section describes how to export a customized report from one Precise system and import it to another Precise system.

NOTE Before you begin, find the source Tier ID and the destination Tier ID in the PS_INAP_APP_TIER or PS_FSEC_ENV_CONTENTS tables.

To export/import a customized report

1. In the source Precise system, create the `ForesightExportRequest.xml` file in the Precise root folder. For example:

```
<root name="Export">
  <report-name>In Sybase Breakdown</report-name>
  <apptier>1201</apptier>
</root>
```

NOTE The report-name value is case-sensitive.

2. Run the following command:

```
infra\bin\psin_http_requestor -dp -x fs-export-import FS
< ForesightExportRequest.xml > out.xml
```

3. Check the `out.xml` file in the `<precise_root>` directory to view the status.
4. Copy the exported report `.zip` file from the source Precise root directory to the destination Precise installation in:

```
<i3 root>\distribution
```

NOTE The `.zip` file name consists of `<technology code>_<report name>` (SY_In Sybase Breakdown.zip). For cross-Tiers the Tech code is ALL (ALL_CA_Rep.zip).

5. In the destination Precise system, create the `ForesightImportRequest.xml` file in the Precise root directory. For example:

```
<root name="Import">
  <report-name>SY_In Sybase Breakdown.zip</report-name>
  <apptier>1174</apptier>
</root>
```

6. Run the following command:

```
infra\bin\psin_http_requestor -dp -x fs-export-
import FS < ForesightImportRequest.xml > out.xml
```

7. Check the `out.xml` file in the Precise root folder to view the status.

About copying customized reports

The following section describes how to copy customized reports from one Tier to another with the CLI utility.

To copy a customized report

1. Run the following commands from the Precise root folder: UNIX

```
./infra/bin/psin_cli.sh -i3-user <user> {-i3-clear-password
<i3 clear password> |
-i3-encrypted-password <i3 encrypted password>}
-action fs-copy-foresight-custom-report
-parametersfile <parameters-file.xml>
```

Windows

```
<i3_root>\infra\bin\psin_cli.bat -i3-user <user> {-i3-clear-
password <i3 clear password> | -i3-encrypted-password <i3
encrypted password>}
-action fs-copy-foresight-custom-report
-parametersfile <parameters-file.xml>
```

Example:

```
<i3_root>\infra\bin\psin_cli.bat -i3-user admin -i3-clear-password admin
-action fs-copy-foresight-custom-report -
parametersfile my_parameters_file.xml
```

2. Create the `parameters-file.xml` file in the Precise root folder and make sure it contains the following:

```
<parameters>
  <parameter name="source-environment" value="Default"/>
  <parameter name="source-apptier" value="SQL Server"/>
  <parameter name="source-report" value="My_report"/>
  <parameter name="destination-environment" value="Default"/>
  <parameter name="destination-apptier" value="SQL Server"/>
</parameters>
```

and make sure to apply the information in the following notes.

NOTE The parameter values are case-sensitive.

NOTE The "source-report" name must not contain blanks.

NOTE The new report name will have a suffix with a serial number (for example: "My_report_1").

NOTE To copy a Cross-Tier report, use the "Cross-AppTier" string for both the "source-apptier" and "destination-apptier" parameters.

Glossary

abandonment rate	In Web, a counter that keeps track of the percentage of users that abandon the loading of the Web page before it completes downloading.
action	An operation that Alerts FocalPoint automatically performs when detecting a warning or critical status for a specific metric. According to the defined action, Alerts FocalPoint opens a message box, sends an email or SNMP trap, or executes a program.
action item (manual)	Some installation steps cannot be executed automatically by Precise Installer. The administrator is required to execute them manually. Action Items are presented (when necessary) at end of installation or product update.
AdminPoint	The central administration console of Precise that facilitates the maintenance, configuration, and management of all installed Precise components, such as monitoring the status of all Precise agents and PMDB processes, getting license information, starting and stopping the agents, getting log data on agents and events, changing PMDB settings, and installing patches. See also PMDB and agent .
advice	In Oracle, an algorithm that is designed to recommend on gathering statistics, creating new indexes, change existing indexes, and add or delete hints to make Oracle's Optimizer choose a better access path and make the statement perform better. Can be activated from any DML (Data Manipulation Language) statement.
agent	A program that runs on a server machine to collect, process, or communicate performance information. The Precise installation consists of multiple agents.
alert	The state of an Alerts metric that has reached a near-critical or critical status. An alert is issued by Alerts, triggering an action and informing of a problem that has occurred or is likely to occur within the area sampled by the specific metric. See also action .
Alerts	These product provide alerts to problematic conditions before they turn into performance problems, based on predefined metrics and thresholds. Alerts can automatically perform an action, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program.
Alerts FocalPoint	An agent that receives data from the InformPoint agents, stores it, and performs any action that has been user-defined for that specific alert, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program. See also InformPoint .
alert type	In Alerts, the status of all metrics belonging to a metric group or a monitored instance, indicating the current performance level through colors. See also agent and metric .
Application	The highest Precise logical group. It may contain multiple AppTiers of various technologies that serve an application together. For example, a Payroll Production environment may contain all Web servers, application servers, transaction managers, databases, and servers that serve this application. Alternatively, it may contain any set of instances that form together an administrative group. Since Precise version 8.5, an instance can belong to more than one environment. Let's say a single Oracle database serves two different applications plus the DBA wants to associate this database with a group of other databases under his responsibility regardless of served applications. In this case the Oracle database will be associated to three different Precise environments. See also AppTier and instance
application server metric	In J2EE, a metrics that is provided by the application server or by customer code. This can include metrics published by the Java Management Extension (JMX) APIs or vendor-specific APIs, such as IBM's Performance Management Interface (PMI).
AppTier	The abbreviation for an application tier in a Precise environment. An AppTier contains one or more instances of the same technology and purpose. Application tiers do not necessarily correspond to distinct physical servers: in many cases, the tiers are logical, with one server running multiple AppTiers or one AppTier spans a cluster of servers. A Precise environment

may contain multiple AppTiers on the same technology. For example, you may group J2EE instances (JVMs) into a J2EE Presentation AppTier and a J2EE Business Logic AppTier. Segmenting application service time into the contribution of individual application tiers is helpful in identifying the source of performance problems. Analyzing the performance and behavior of each tier separately is crucial for isolating the root causes of performance problems.

cabinet	In Oracle, SQL Server, DB2, and Sybase, the highest logical level in the SQL workspace hierarchy. A cabinet contains folders and, within folders, statements.
call path	A subset of a stack trace including only those methods that have been instrumented.
client cache	In Web, a counter that keeps track of the percentage of requests taken from the client cache (http status 304).
client-side collection	See Web client browser-side .
cluster	A group of servers or instances that are configured to be treated as a single entity. In Oracle, all Oracle instances of the same Oracle database (RAC configuration) are treated as a cluster and presented in the Precise for Oracle UI as a Database. In J2EE, a group of JVMs that serve the same application can be defined as Cluster (even if they are not defined as a cluster in the application server configuration). All JVMs of the same cluster share the same configuration.
collapsed access plan	In SQL Server and Sybase, the access plan of a unique group of statements or batches that belong to the same collapsed statement or batch but have different access plans. This can differ due to the constants in the text of the original (not collapsed) statements. See also explained statement and collapsed statement .
collapsed statement	In Oracle, SQL Server, Sybase, and DB2, a statement whose constants are replaced with parameters. Each collapsed statement can have several access plans, according to the occurrences of its statements.
Collector	<p>The program that runs on the monitored server to collect performance information.</p> <p>Some technologies allow a single Collector (agent) to serve multiple instances running on the same server. Other technologies require a dedicated Collector per monitored instance.</p> <p>In most cases, it must be installed on the monitored server. Collectors for SAP, SQL Server, and Sybase may reside on a remote server. See also agent.</p>
collector slice size	In J2EE, the smallest unit of data presentation and storage. It can be configured as 30 seconds or five minutes. If SmartLink is configured, relevant Web and Microsoft .NET Collectors will report to SmartLink using the same collector slice size. Regardless of the defined collector slice size, the most granular slice size presented on both the Precise for Web UI and the Precise for Microsoft .NET UI will be 15 minute intervals.
completion rate	In Web, a counter that keeps track of the percentage of users that completed the loading of a Web page.
count	The number of occurrences observed during a measurement interval.
CPU time	The average amount of time consumed by the operating system actively processing instructions on behalf of a running activity.
critical status	In Alerts, the status represented by a red bullet indicating that the value of the sampled metric has exceeded the near-critical and critical threshold values. See also metric .
Cross-AppTiers	In Insight, Report Manager, and Alerts, a perspective that provides high-level information about the performance of all AppTiers in an environment, including operating system data. It provides a holistic view of the entire environment and helps understand how the AppTiers interact. See also AppTier .
current data	In SmartLink, the end-to-end activity, detected by Precise over the last five minutes, including activities that are still in progress. In all databases products, a workspace that presents Current Sessions and the SQL statements that were executed including statements that are still in progress.
custom	Represents invocations of a type not specifically defined. See also work time .
customized metric	In Alerts, a user-defined metric that measures site-specific parameters.
drilldown	Within Precise products and Insight SmartLink, the filtering of analyzed data by clicking a specific entity. Then, additional information about the selected entity is presented, plus a breakdown of its activity by another entity. For example, drilldown on a transaction and get an overtime graph of the selected transaction, plus a list of top users that executed the selected transaction.

entry point	Usually a top-level HTTP or EJB request. An entry point can originate when a user clicks in a browser or an E-commerce server invokes a remote EJB. J2EE makes a distinction between service request invocations and other invocations. The first HTTP or EJB invocation within a call path is designated as an entry point.
ERP Extension	In Oracle, DB2, and SQL Server, an extension to the Collector that provides detailed information on the activities and resource consumption of packaged application components. It correlates the database information and the packaged application information and lets you see users, transactions, reports, and other elements of ERP applications, such as Oracle Applications, SAP, PeopleSoft, and Siebel.
event	In Alerts, the occurrence of a sampling or progress. A sampling occurrence occurs every time a metric samples. A progress occurrence occurs when a metric's progress status is changed or when the investigated status reaches the end of the given investigation time. In AdminPoint, all occurrences reported by Precise agents, including informational events, warnings, and errors related to one of the agents. All events are shown in the Events view. See also metric and progress .
executions	In Oracle, SQL Server, DB2, and Sybase, the number of times a SQL statement was executed during the selected time frame. In SmartLink and SAP, the number of times a transaction was executed during the selected time frame.
explained statement	In Oracle, SQL Server, DB2, and Sybase, a statement whose access path (chosen by the RDBMS Optimizer) is clarified and translated into a visual display. Explained results include information on the objects referenced by the statement and the operations performed on these objects.
extended collection	In Oracle, a function that proactively specifies a future period during which Oracle activity data is collected and organized for subsequent analysis. Extended collections are an easy means to view collected information, assess application resource consumption, and identify bottlenecks that are inhibiting application performance and end-user productivity.
Federated Precise	Federated Precise is a version of Precise (starting with version 8.5) that can manage multiple Precise installations within unified StartPoint and AdminPoint screens, displaying and managing all environments, instances, and installations.
findings	A ranked list of top performance problems in the selected context. They are presented in Oracle, SQL Server, .NET, and J2EE. Each finding provides: a short explanation of the problem, detailed background information, and most important: one or more links to further investigate the problem and possibly get more granular findings on the selected context.
first byte time	In Web, a counter that keeps track of the time that it takes from the moment a new Web page is called until the first byte arrives back from the Web server.
FocalPoint	An agent that communicates with the Listeners installed on the monitored servers, receives data from the Collectors, periodically processes and stores this data in the PMDB, and serves UI requests.
folder	In SQL Server, Oracle, DB2, and Sybase, the intermediate logical level in the SQL workspace hierarchy. Folders are grouped into cabinets and contain SQL statements. See also cabinet .
Framework Installer	The application that facilitates the installation of Precise framework components. It can be invoked from the installation DVD to install a new Precise deployment. It can also be invoked from an existing Precise deployment to install an additional framework node and attach it to the originating Precise deployment. See also Framework node .
Framework node	A set of FocalPoint agents that are installed together and manages a set of monitored instances. The performance data of these instances will be loaded into a dedicated PMDB. A single Precise deployment may contain multiple framework nodes (using a separate PMDB for each node). An environment cannot span over multiple framework nodes.
garbage collection	An automatic process in the Java runtime environment that periodically reclaims memory used by objects that are no longer referenced. The process can impact an application's performance while memory is being reclaimed. Java programmers may initiate garbage collection explicitly.
grouper	In Web, the identifier that is used to group other identifiers, such as sites or URLs.
hint	In Oracle, an instruction directed at the Oracle Optimizer that includes considerations for an execution plan. The Oracle Optimizer will build an execution plan based on the hint, ignoring its own set of considerations.

hour group	A unit that reflects the type and level of activity within the system at different times. By defining the times of the day that are peak and off-peak, or day and night, the performance analysis can be focused on those particular times of the day. If, for example, most performance problems occur within nighttime and weekend batches, it can be useful to focus only on them.
InformPoint	In Alerts, an agent that retrieves performance data from all installed Precise products, analyzes it, and sends an alert to the Alerts FocalPoint when the predefined thresholds are exceeded. See also agent .
Insight	The Precise product family that facilitates the process of monitoring and correlating system performance. It consists of Insight. See also Report Manager and FocalPoint .
Insight FocalPoint	An agent that receives performance information from Insight Savvies, which monitor the environment. Insight FocalPoint then correlates, processes, and stores this information in a centralized location. The Insight performance history is stored in the PMDB. See also Savvy and PMDB .
instance	A monitored object of a specific technology. The following list specifies what constitutes an instance for the various supported technologies: J2EE - a Java Virtual Machine (JVM - a logical name set by the user), Microsoft .NET - a Common Language Runtime (CLR - a logical name set by the user), Oracle Applications - an Oracle Applications Form server, Oracle - an Oracle instance, SAP - a SAP system, SQL Server - an SQL instance, Sybase - a Sybase instance, Tuxedo - a Tuxedo domain, Web - a Web server, WebSphere MQ - an IBM WebSphere Queue Manager, DB2 - a DB2 non-partitioned database or a DB2 database partition. During installation, the instance is associated with one AppTier and environment. An Instance can be moved to a different environment or associated with multiple environments without re-installation nor losing historical data. See also Application .
instance statistics	See performance counter .
instance/database changes	In SQL Server and Oracle, the part of the Collect Schema Changes process that captures instance definition changes and database option changes and saves them in the PMDB.
instrumentation	The process of inserting fault-tolerant recording hooks in Java byte code, .NET MSIL, HTML pages, or other monitored components, resulting in the capture of performance metrics. In J2EE, a mechanism that enables collecting performance information when an application is executed. The process involves inserting special fault-tolerant recording hooks into application class objects. In WEB, the insertion of recording hooks into HTML pages can either be in memory (Dynamic Instrumentation) or file-based (Static Instrumentation).
instrumentation context	See invocation context .
Internal invocation	The process of invoking a request from an HTTP request (Servlet or JSP) or an EJB. J2EE displays an internal invocation like any other invocation.
invocation	An execution of a J2EE entity (a Servlet, EJB, SQL Statement, method, and so on). When mentioned in plural (Invocations), means the amount of times that the entity was executed.
invocation context	The context within which a method is invoked. For example, if Method A is invoked both from Method B and from Method C, there will be two different invocation contexts for the performance metrics collected for Method A, one for when it is invoked from Method B and one for when it is invoked from Method C. See also call path and instrumentation context .
JRE (Java Runtime Environment)	As the runtime part of the Java software, the combination of the components that enable the execution of a Java program: a Java virtual machine, the core class libraries, and the files that form the Java platform.
JSP (Java Server Page)	An HTML page with special tags for Java scripting. An application server processes the tags and generates a Servlet.
JVM (Java Virtual Machine)	An instance of a JRE that executes Java programs. A server-side Java application server is itself a Java program that runs inside a JVM. Servlets, JSPs and EJBs are Java programs (applications) that run within the application server's JVM. J2EE monitors the JVM running the application server and the server-side Java applications within the application server.
JVM Heap Memory	The amount of real computer memory that is allocated to the JVM for executing Java programs.

key metric	In Alerts, a parameter that monitors a very important performance aspect. The status of the Oracle instance (up/down), for example, is crucial for system performance. If the instance is down, this is the first problem that needs to be solved. Marking a metric as key metric ensures that a critical alert raised for this metric receives top priority by the person that is responsible to handle alerts at any time. It can also help determine which alert to handle first in case of multiple alerts. In Alerts, key metrics are always displayed at the top of a metric table when it is sorted by alerts so that they get immediate attention.
Listener	The agent that facilitates the communication between the various Precise agents across different servers must be installed on every server where Collectors or FocalPoint agents are installed. The Listener allows communication with all other agents installed on the monitored server, while only the Precise Listener port is known by other servers.
login name	In SQL Server and Sybase, the session identifier that represents the credential used to connect to the database. When an ERP extension is installed, the user name of the packaged application's client overrides the login name. For example, when SAP extension is installed, the SAP user name overrides the login name.
machine	A session identifier. A machine as sampled, for example, by the SQL Server Collector is the identification of the machine where the client process executes. Machine is also sampled by Oracle, DB2 and Sybase. In Insight terminology Machine is called Client machine.
Main framework node	A main framework node is the single point for login and also serves as the Precise FocalPoint for the entire deployment.
major collection count	Number of estimated major garbage collection events that occurred during the last J2EE collector slice. A major garbage collection can stop the application while JVM heap memory is being reclaimed. See also garbage collection .
major garbage collection time	Percentage of time spent by the JVM executing major garbage collection events during the displayed interval. See also garbage collection .
MBean	A Java object that represents a manageable resource. In J2EE, MBeans, or Managed Beans, are used for application server metrics.
memory logger interval	The interval at which the J2EE Collector gathers JVM heap memory data. All snapshots of the memory logger's data collected according to the memory logger interval are summarized with counters in the current aggregation interval. The memory logger interval's time span is typically a small fraction of the aggregation interval.
metric	In Alerts, a query that helps measure performance in the environment. Three types of metrics are available: System metrics relate to the internal resources, operations, and objects of the monitored infrastructure; application metrics reflect the way the applications perform; user-defined metrics can be customized to specifically relate to a site. When a metric's value exceeds one of the defined thresholds, its status changes to near-critical (yellow bullet) or critical (red bullet). See also key metric .
metric set	In Alerts, a unit that groups metrics that measure related performance aspects. The following metric sets exist: <ul style="list-style-type: none"> • Status: includes metrics that alert to functional problems related to the instance. Performance: includes metrics that alert to performance problems related to the instance. • Load: includes metrics that alert to instance-related load problems that may later cause errors or crashes. Service: includes metrics that alert to instance-related SLA breaches. The metrics in this set are sampled by Insight. • Performance Trending: includes metrics that alert to potential future performance problems. The metrics in this set are sampled by Report Manager. • Load Trending: includes metrics that alert to potential future load problems. The metrics in this set are sampled by Report Manager. • Customized: includes user-defined metrics. • Precise Status: includes metrics that alert to the near-critical or critical status of the installed Precise environment.
module	A session identifier. In Oracle it contains the value of MODULE column in V\$SESSION table. In DB2 it contains the command name of non-SQL statements or the package name for SQL statements.

near-critical status	In Alerts, a status indicating that the value of the sampled metric has exceeded the defined near-critical threshold. A near-critical status is indicated by a yellow bullet.
network time	In Web, a counter that keeps track of the time spent on network activity from the server side perspective. This includes the time to read the request from the network and the network time to send the response back to the client.
node	See Framework node . NOTE: when creating a new installation from DVD, it would be called "Framework." In the UI screens (columns in tables, choosing a system for a new environment etc), we would call them "Nodes." For example: choose a node for the new environment. When adding a new system within AdminPoint, it would be called a "Framework node".
packaged application	An application that is created and/or maintained by a third party and is not custom-built to one's specific needs. The following packaged applications have special treatment by Precise: SAP, Oracle Applications, PeopleSoft, and Siebel. They are harder to modify because the application code is either not available or hard to understand. See also ERP Extension .
page size	In Web, a counter that keeps track of the amount of data loaded from the server to display the page.
page views	In Web, a counter that keeps track of the number of Web pages viewed at a specific Web site during a selected time period.
parent metric	In Alerts, a joining of several child metrics. Each time a parent metric samples, it gathers data from a set of child metrics and presents it as a single metric query. The individual child metric values are displayed on the Thresholds tab of the Properties dialog box in Alerts. See also submetric .
performance counter	In SQL Server, a Windows performance counter as reported by the operating system. In Oracle, an instance level statistic as gathered from some Oracle V\$ views. In Sybase, a sysmonitors counter collected by Sybase. In Web, an operating system or a Web server performance counter.
PMDB	The Precise data warehouse of performance and availability data. It can be hosted on an Oracle or SQL Server database.
portal server	An application server for Web-based applications that commonly provide personalization, single sign-on, and content aggregation from different sources and that host the presentation layer of Information Systems. Aggregation is the action of integrating content from different sources within a Web page. A portal may have sophisticated personalization features to provide customized content to users. Portal pages may have a different set of portlets creating content for different users.
portlet	A Java-based Web component, managed by a portlet container, that processes requests and generates dynamic content. Portlets are used by portals as pluggable user interface components that provide a presentation layer to Information Systems.
Precise deployment	An independent Precise system. It contains and manages various agent types and provides centralized monitoring and administration. A Precise deployment may contain multiple framework nodes. One of them is defined as the main framework node and it manages all other nodes of the Precise deployment.
Precise FocalPoint	See Precise FocalPoint agent .
Precise FocalPoint agent	An agent that manages all agents in a single Precise deployment. Additional product FocalPoint agents manage specific technologies and resources.
Precise Proxy agent	When multiple Framework Nodes are managed by a single Precise deployment, the Precise FocalPoint agent manages all of the agents of the main Framework node, while every other framework node has a Precise Proxy agent to manage all its agents.
Precise root	This is the term used in a path for the Precise installation directory. The terms <i><i3_root></i> or <i>i3 root</i> can appear in text too.

program	A session identifier in Oracle, DB2, SQL Server, Sybase, Insight, and Alerts. A program as sampled, for example, by the SQL Server Collector, is the name of the executable that connects to the database. Applications that do not set the application name have N/A as program. When a SAP extension is installed, the SAP transaction overrides the program, and it may change during the application's lifetime. When a Siebel extension is installed, the Siebel views override the program, and it may change during the application's lifetime. In Insight, program is the name of an executable as recognized by the operating system. If an executable is invoked from a script (a batch or a shell), the script is displayed as the command entity. In Alerts, program is part of the customized metrics definition and is the name of the executable or stored procedure executed in the database that runs when the metric is sampled. In action definitions, program is the name of the executable that will run if the metric exceeds its threshold.
progress	In Alerts, the management state of a metric for which an alert has been raised. The following statuses exist: <ul style="list-style-type: none"> • Open: An alert is raised. • Investigated: The alert is taken care of. • Closed: The problem has been solved.
Proxy FocalPoint	See Precise Proxy agent .
recommended index	In SQL Server, a function that uses the Microsoft Index Tuning Wizard to recommend on adding indexes or statistics for a selected statement, batch, or table. Based on the results of this function, the Optimizer will choose a better access plan and make the respective statement or batch perform better. For statements and batches, recommendations are based on the content of the statement or batch. For tables, recommendations are based on all the statements that are stored in the PMDB, are executed during the selected time period, and have an average duration time greater than the value defined in the registry (where 0 is the default). See also advice .
related SQL	In Oracle, a generated statement that uses alternative syntax to access the database in different ways and returns the same output as the original statement.
relative frequency	Number of contributor invocations per entry point. For example, if a service request calls three methods each time it is invoked, the Relative Frequency for the method is three invocations per service request execution. Similarly, if a service request calls one method every other time it is invoked (half of the time), the Relative Frequency for the method is 0.5 invocations per service request execution.
rendering time	In Web, a counter that keeps track of the time that it takes for a Web page to be loaded from the moment the first byte arrives until the Web page is fully loaded or the user interrupts or abandons the loading process.
report	A collection of queries, programming code, and layout settings that Report Manager executes to generate graphical results like tables and charts.
Report Manager	Uses historical information to identify problematic conditions, track long-term performance, volume trends and patterns, view availability problems over long periods of time and on different levels, compare the performance of similar systems, correlate between performance metrics of different products, assist in capacity planning, and generate demand-driven, user-defined reports.
Report Manager FocalPoint	An agent that examines the PMDB tables to produce scheduled performance reports. See also PMDB .
report parameter	In Report Manager, a keyword used in a report. Its value is set during the report execution. A parameter's value can be updated either permanently or for the current execution only.
report property	In Report Manager, the attributes that define a specific report, consisting of report parameters and scheduling information. See also report parameter .
request	In Web, a counter that keeps track of the number of HTTP requests sent for a viewed entity.
request error	In Web, a counter that keeps track of the percentage of requests completed with an HTTP error.
sampling	In Alerts, the process during which a metric queries a Precise product for a specific instance, retrieves values and calculates the metric's alert level.
sampling base	In Alerts, the start time of a scheduled sampling process (by default Sunday, 00:00 AM).

sampling period	In specific Alerts metrics, the time frame during which statistical data is returned from other Precise products. Such a metric is for example Oracle's General Behavior metric, which returns database behavior for a certain period of time.
sampling rate	In Alerts, the frequency of a metric's regular sampling schedule. The sampling rate is measured in minutes.
Savvy	An Insight agent that collects AppTier-specific performance data.
scalability	A system's ability to withstand load. For example, positive scalability means that the system continues to function properly even when it is called upon to service a larger number of users.
schema changes	In Oracle and SQL Server, a process that captures schema changes and saves them in the PMDB. In SQL Server, instance configuration parameters and database option changes are also captured.
Script installer	A component of Precise Agent Installer that enables adding a small JavaScript script, also known as a Web browser-side agent, to the Web pages of your Web site.
server	Also called a host machine or server machine. The combination of a computer and associated operating system software that is accessed by multiple clients to perform requested functions and typically returns results to the clients.
server side collection	In Web, the collecting of performance data from the Web server instance through the Web server agent.
service + network time	In Web, a counter that keeps track of the service and network time, including the total amount of time the request took to reach the server.
Service Level Agreement	See SLA (Service Level Agreement) .
service time	In Web and SmartLink, the time elapsed from when a request is received by the server to the time a response is sent back to the computer that made that request. The service time is measured on the server side. In J2EE, the time it takes an invocation to complete execution. In other words, service time is the average length of time between the start time and end time of a Java method execution. For example, the SQL service time is the time it takes the JDBC method executing the SQL statement to be completed. The service time includes CPU and wait time. The service times are reported in the interval in which they complete execution. Though a contributor's average service time may be very small, the contributor may be called many times. As a result, a contributor's overall contribution to performance may be large even though its average service time is low.
servlet	A class that is loaded only once and for which the application server uses multithreading to process requests. The servlet generates an HTML page that is sent back to the Web browser.
size	In Web, a counter that keeps track of the bytes sent and received.
SLA (Service Level Agreement)	A formal definition of an information system's performance goals. Within Precise, an SLA consists of clauses corresponding to various system activities. Once a system's SLA is defined, its SLA compliance can be analyzed, and breaches can be isolated to identify their causes.
SmartLink	In Insight, a component that provides transactive, correlated information across AppTiers from the user's perspective. Insight SmartLink functions in PeopleSoft and Web applications environments.
SmarTune	In Oracle, SQL Server, and .NET, an algorithm that analyzes performance metrics, identifies and ranks potential problem areas, and provides advice for correcting the problems.
socket	An I/O abstraction layer that enables processes to communicate with each other, regardless of whether they are on the same server. Sockets are bi-directional FIFO pipes that look like regular file I/O to the developer with the abstraction layer handling all of the low-level communication details.
StartPoint	The opening page of Precise. It provides a quick overview of the environment status and links to launch any of the Precise products.
Statistics workspace	In Oracle, DB2, SQL Server, Sybase, J2EE, and Web this workspace presents instance level statistics such as hit ratios, utilization, JMX metrics, logging, and more.
submetric	Also called child metric. In Alerts, the subquery of a parent metric. A child metric gathers its own data and combines it with the data gathered by other child metrics to form the result of the parent metric. Each child metric has its own thresholds and may be enabled or disabled individually. See also parent metric .

summary table	<p>A container that stores the information collected by the Precise agents and loads it into the PMDB.</p> <p>Summary tables store the same data at different levels of granularity: time slice, hourly, daily, weekly, and monthly. By storing data in multiple summary tables, Precise can present a detailed view and progressively higher-level views of the same data. Summary tables are particularly useful for data aging. A data purging policy can be implemented for each summary table so that detailed data is retained for short-term historical analyses while more summarized data is used for long-term analyses and trending.</p>
technology	<p>A technology identifies the monitored object. For example, Oracle, SQL Server, Sybase, and DB2 are different technologies, while all Web servers (such as: Apache, IIS, and WebSphere) are defined as a single Web technology. A single monitored object can be monitored by two different technology's Collectors. For example, WebLogic server can be monitored by both a Web Collector and a J2EE Collector.</p>
throughput	<p>The number of average completions per second that are observed during an interval.</p>
time slice	<p>A unit used to break up long sessions into smaller time periods. The length of a time slice is preset and cannot be changed. It represents the maximum time that passes before the data collected can be displayed. For example, if the length of a time slice is 15 minutes, the collection is only updated at 15-minute intervals.</p> <p>The length of a time slice is different for each technology: In Oracle, SQL Server, Sybase, DB2, Web, .NET, MQ, Tuxedo, Network, and OS, a time slice is 15 minutes. In J2EE and SmartLink, a time slice is either 30 seconds or 5 minutes. In Precise for SAP, a time slice is 5 minutes. See also collector slice size.</p>
URI (Uniform Resource Identifier)	<p>The relative path to a resource after the location (network node) is found.</p>
URL mapping	<p>In Web, a function that defines rules that map URLs (Uniform Resource Locators) with dynamic parameters originating from a specific domain to a format that identifies the Web pages and prevents them from having different URLs.</p>
Web client browser side	<p>In Web, the gathering of performance data from the browser running on the desktop of the Web application's user, through the static or dynamic instrumentation of Web pages.</p>
Web server	<p>A program that receives client requests for Web pages, retrieves the static pages and/or issues a request for dynamic page creation to an application server, and sends the pages to the client. In Web, a computer that delivers (serves up) Web pages. Every Web server has an IP address and possibly a domain name.</p>
work time	<p>The time spent in the specific entity excluding time spent by called (other) entities. In J2EE, work time of a method is calculated as the service time of that method minus the service time of the methods it invoked.</p>
work type	<p>In Web, a counter that keeps track of the type of the Web instance, such as Web, PeopleSoft, SAP, or Siebel.</p>
workspace	<p>A display unit in Precise products. All workspaces display data from different perspectives. For example, in Oracle, SQL Server, and Sybase, the Current workspace shows information on the sessions currently active in an application, and the Objects workspace displays information on Oracle or SQL Server database objects that can be used to understand relationships and associations between database schema objects.</p>