



EDB Postgres™ Enterprise Manager Enterprise Features Guide Version 7.5

October 22, 2018

EDB Postgres Enterprise Manager Enterprise Features Guide
by EnterpriseDB® Corporation
Copyright © 2013 - 2018 EnterpriseDB Corporation. All rights reserved.

EnterpriseDB Corporation, 34 Crosby Drive Suite 201, Bedford, MA 01730, USA
T +1 781 357 3390 **F** +1 978 467 1307 **E** info@enterprisedb.com **www**.enterprisedb.com

EnterpriseDB, EDB Postgres, Postgres Plus, Postgres Enterprise Manager, and DynaTune are trademarks of EnterpriseDB Corporation. Other names may be trademarks of their respective owners. © 2018.

Table of Contents

1	Introduction	6
1.1	What's New	8
1.2	Typographical Conventions Used in this Guide	8
2	The PEM Query Tool	10
3	Package Deployment	12
3.1	Installing a New Package	15
3.1.1	Reviewing Scheduled Tasks	20
3.2	Upgrading an Installed Package	21
4	Performance Monitoring and Management	25
4.1	Using Dashboards to View Performance Information	25
4.2	Managing Custom Dashboards	29
4.2.1	Creating a Custom Dashboard	30
4.2.2	Creating an Ops Dashboard	35
4.3	Using the Manage Charts tab	36
4.3.1	Creating a Custom Chart	38
4.3.2	Importing a Capacity Manager Template	44
4.4	Customizing Probes	48
4.4.1	Creating a Custom Probe	50
4.4.2	Deleting a Probe	56
4.4.3	Copying a Probe	57
4.5	Alerting	58
4.5.1	Using the Alerts Dashboard	59
4.5.2	Using the Manage Alerts Tab	62
4.5.2.1	Creating a Custom Alert Template	63
4.5.2.2	Creating a New Alert	69
4.5.2.3	Modifying or Deleting an Alert	75
4.5.2.4	Copying an Alert	78
4.5.2.5	Audit Log Alerting	79
4.5.2.6	Creating an Email Group	80
4.5.3	Using PEM with Nagios	83
4.5.3.1	Enabling Nagios Notification for an Alert	85
4.5.3.2	Configuring Nagios-related behavior of the PEM Server	86

4.5.3.3	Creating the hosts.cfg and services.cfg File	87
4.5.3.4	Modifying the Nagios Configuration File	88
5	Capacity Manager	89
5.1	Capacity Manager Templates	94
6	Audit Manager	95
6.1	Setting the Advanced Server Instance Service ID	96
6.2	Setting the EDB Audit Configuration Probe	97
6.3	Configuring Audit Logging with the Audit Manager	98
6.4	Viewing the Log with the Audit Log Dashboard	106
7	Log Manager	108
7.1	Reviewing the Server Log Analysis Dashboard	119
8	Postgres Log Analysis Expert	121
8.1	Reviewing the Postgres Log Analysis Expert Report	127
9	SQL Profiling and Analysis	128
9.1	Creating a New SQL Trace	129
9.1.1	Creating a Trace	129
9.1.2	Opening an Existing Trace	133
9.1.3	Filtering a Trace	134
9.1.4	Deleting a Trace	136
9.1.5	Viewing Scheduled Traces	137
9.2	Using the Index Advisor	139
10	Tuning Wizard	140
11	Postgres Expert - Best Practice Enforcement	148
11.1	Using the Postgres Expert Wizard	148
11.2	Reviewing Postgres Expert Recommendations	152
12	Configuring Streaming Replication	154
12.1	Monitoring Streaming Replication and Failover Manager	169
12.1.1	Configuring High-Availability for PEM	171
13	Monitoring Failover Manager	173
13.1.1	Replacing a Master Node	175
14	Monitoring an xDB Replication Cluster	176
15	Reference	178
15.1	PEM Server Configuration Parameters - Reference	178
15.2	Capacity Manager Metrics - Reference	186

15.3	PEM Probes – Reference.....	189
15.4	PEM Pre-defined Alert Templates – Reference.....	194
15.4.1	Templates applicable on Agent	194
15.4.2	Templates applicable on Server.....	195
15.4.3	Templates applicable on Database	198
15.4.4	Templates applicable on Schema	201
15.4.5	Templates applicable on Table.....	202
15.4.6	Global Templates.....	202
16	Conclusion.....	203

1 Introduction

This guide will acquaint you with the tools and wizards that are built into the Postgres Enterprise Manager™ (PEM) web interface that make it easier for you to monitor and manage your system:

- **PEM Query Tool** - Chapter [2](#) introduces you to the PEM Query Tool. The Query Tool is a feature-rich Interactive Development Environment (IDE) that allows you to issue ad-hoc SQL queries against Postgres servers.
- **Package Deployment** – Chapter [3](#) introduces you to the PEM Package Deployment wizard.
- **Performance monitoring** – Chapter [4](#) provides an introduction to monitoring and troubleshooting performance on your database servers.
- **Capacity planning** – Chapter [5](#) discusses how to perform trend analysis and forecasting from data collected in Postgres Enterprise Manager.
- **Audit Manager** – Chapter [6](#) provides information about using Audit Manager to configure logging attributes, and using the Audit Log Analysis dashboard to filter and review the log files.
- **Log Manager** – Chapter [7](#) walks you through using Log Manager to modify server log configuration parameters, and using the Server Log Analysis dashboard to filter and review the server log.
- **Postgres Log Analysis Expert** – Chapter [8](#) introduces you to the Log Analysis expert. The Log Analysis Expert analyzes the log files of registered servers and produces a report that provides an analysis of your Postgres cluster's usage based on log file entries.
- **SQL Profiling** – Chapter [9](#) describes how to capture a SQL workload and analyze the output to tune and fix poorly running SQL.
- **Tuning Wizard** – Chapter [10](#) walks you through using the PEM Tuning Wizard. The PEM Tuning Wizard reviews your installation and recommends a set of configuration options that will help tune your installation.
- **Postgres Expert - Best Practices Enforcement** – Chapter [11](#) discusses how to use Postgres Expert to uncover configuration or design issues that need correcting on your database servers.

- **Streaming Replication Wizard** – Chapter [12](#) discusses using the Streaming Replication wizard to configure Postgres Streaming Replication on new or existing servers, and how to monitor the servers when the replication scenario is up and running.
- **Monitoring Failover Manager** - Chapter [13](#) provides an overview of PEM functionality that allows you to monitor EDB Failover Manager.
- **Monitoring an xDB Replication Cluster** - Chapter [14](#) provides an overview of PEM functionality that allows you to monitor EDB Failover Manager.

This guide is not a comprehensive resource; rather, it is meant to serve as an aid to help you evaluate the tool and bring you up to speed with the basics of how to use the product. For more detailed information about using PEM's functionality, please see the online help made available by the PEM client.

This document uses *Postgres* to mean either the PostgreSQL or EDB Postgres Advanced Server database.

1.1 What's New

The following features have been added to create Postgres Enterprise Manager 7.5:

- PEM 7.5 adds capabilities for monitoring and managing Advanced Server v11 and PostgreSQL v11 clusters.
- PEM 7.5 allows you to use PgBouncer as a connection pooler between the PEM database server and PEM agents.

For detailed information about supported features, please see the *EDB Postgres Enterprise Manager EDB Ark Management Guide* available from the EnterpriseDB website at:

<https://www.enterprisedb.com/resources/product-documentation>

1.2 Typographical Conventions Used in this Guide

Certain typographical conventions are used in this manual to clarify the meaning and usage of various commands, statements, programs, examples, etc. This section provides a summary of these conventions.

In the following descriptions a *term* refers to any word or group of words that are language keywords, user-supplied values, literals, etc. A term's exact meaning depends upon the context in which it is used.

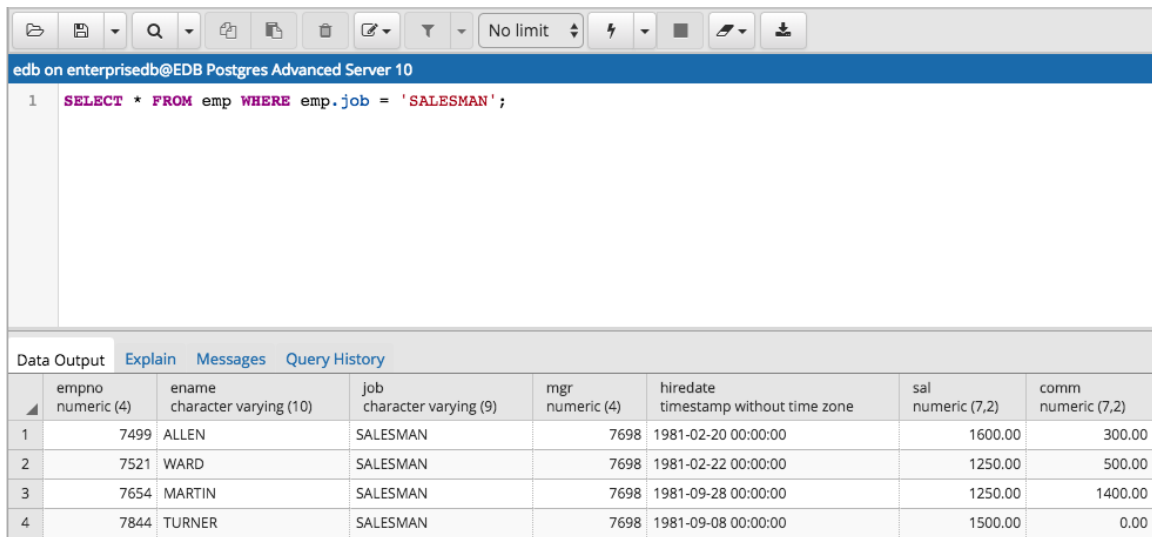
- *Italic font* introduces a new term, typically, in the sentence that defines it for the first time.
- Fixed-width (mono-spaced) font is used for terms that must be given literally such as SQL commands, specific table and column names used in the examples, programming language keywords, etc. For example, `SELECT * FROM emp;`
- *Italic fixed-width font* is used for terms for which the user must substitute values in actual usage. For example, `DELETE FROM table_name;`
- A vertical pipe | denotes a choice between the terms on either side of the pipe. A vertical pipe is used to separate two or more alternative terms within square brackets (optional choices) or braces (one mandatory choice).
- Square brackets [] denote that one or none of the enclosed term(s) may be substituted. For example, [a | b], means choose one of “a” or “b” or neither of the two.
- Braces { } denote that exactly one of the enclosed alternatives must be specified. For example, { a | b }, means exactly one of “a” or “b” must be specified.

- Ellipses ... denote that the proceeding term may be repeated. For example, [a | b] ... means that you may have the sequence, “b a a b a”.

2 The PEM Query Tool

PEM contains a feature-rich Interactive Development Environment (IDE) that allows you to issue ad-hoc SQL queries against Postgres servers. To open the `Query Tool` SQL IDE from within PEM, simply highlight the name of the database you want to query in the tree control, and select `Query tool` from the `Tools` menu.

The `Query Tool` dialog provides an interface that allows you to manually enter SQL queries, graphically execute and interpret SQL statements, `EXPLAIN` queries and much more (see Figure 2.0).



edb on enterprisedb@EDB Postgres Advanced Server 10							
1 <code>SELECT * FROM emp WHERE emp.job = 'SALESMAN';</code>							
Data Output Explain Messages Query History							
	empno numeric (4)	ename character varying (10)	job character varying (9)	mgr numeric (4)	hiredate timestamp without time zone	sal numeric (7,2)	comm numeric (7,2)
1	7499	ALLEN	SALESMAN	7698	1981-02-20 00:00:00	1600.00	300.00
2	7521	WARD	SALESMAN	7698	1981-02-22 00:00:00	1250.00	500.00
3	7654	MARTIN	SALESMAN	7698	1981-09-28 00:00:00	1250.00	1400.00
4	7844	TURNER	SALESMAN	7698	1981-09-08 00:00:00	1500.00	0.00

Figure 2.0 - The PEM Query Tool.

The upper panel of the `Query Tool` contains the SQL Editor. You can use the panel to manually enter a query, or read the query from a file. If you are manually entering a SQL query, the edit entry window also contains auto completion code and formatting features that help you write queries.

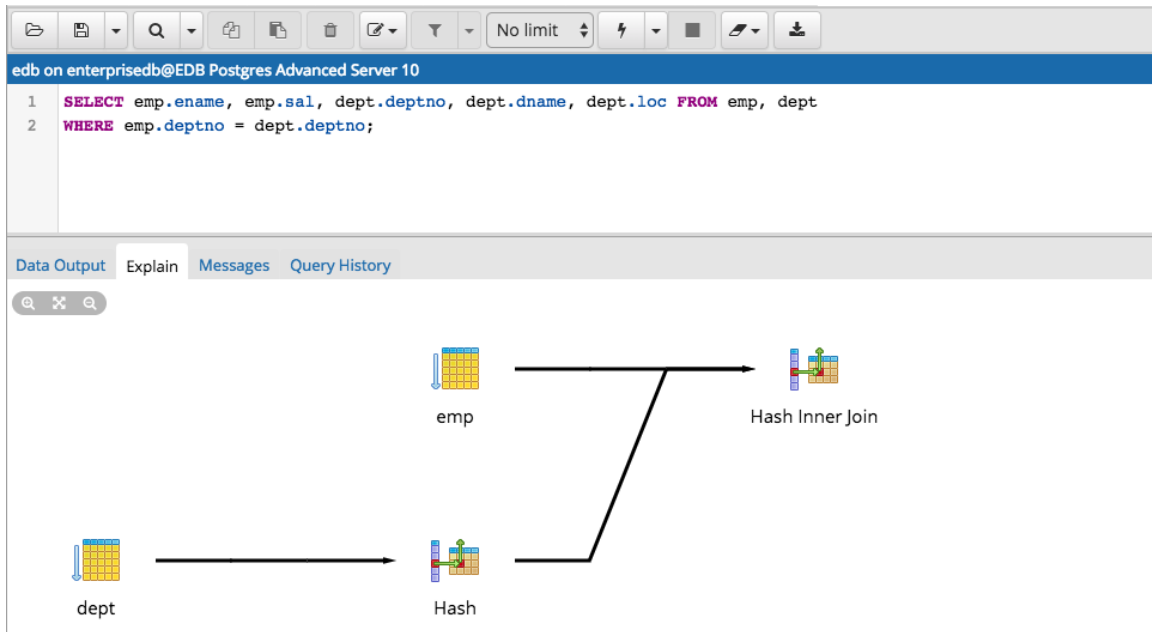


Figure 2.1 - The Query Tool Graphical explain.

After executing a query, you can view the result set or an `EXPLAIN` plan in the lower panel of the Query Tool (see Figure 2.1). As with all PEM features, detailed online help text is available with the click of a button.

3 Package Deployment

The Package Deployment wizard walks you through the process of scheduling the installation of new packages or upgrades of existing packages. Please note: the PEM server must have internet access to deploy packages for deployment.

Before invoking the Package Deployment wizard, you must modify the PEM agent configuration file and restart the agent; first on the server, and then on each system where packages will be deployed.

- On Linux, modify the `agent.cfg` file, setting the `allow_package_management` property to `true`. The configuration file is located in: `/opt/PEM/agent/etc`
- On Windows, use the Registry Editor to modify the registry entry for the agent, setting the value of the `AllowPackageManagement` property to `true`. The entry is located in:
`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\EnterpriseDB\PEM\agent`

After modifying the agent configuration properties, you must restart the PEM agent. On a Linux host, you can use the `service` command:

```
service pemagent restart
```

Or on a Windows host, use the Services dialog to restart the PEM agent service:

```
Postgres Enterprise Manager - pemAgent
```

After enabling package management and restarting the agents, you should also confirm that agent-level probes are enabled on the host of the PEM server, and on any system on which a package will be deployed. To access the `Manage Probes` tab, highlight the name of the PEM agent in the PEM client tree control, and select `Manage Probes...` from the `Management` menu. The following probes must be enabled:

- the `Package Catalog` probe on the PEM server host.
- the `Installed Packages` probe on any system on which you wish to install packages.

Then, to open the Package Deployment wizard, select `Package Deployment...` from the `Management` menu. The Package Deployment wizard `Welcome...` dialog opens (see Figure 3.1).

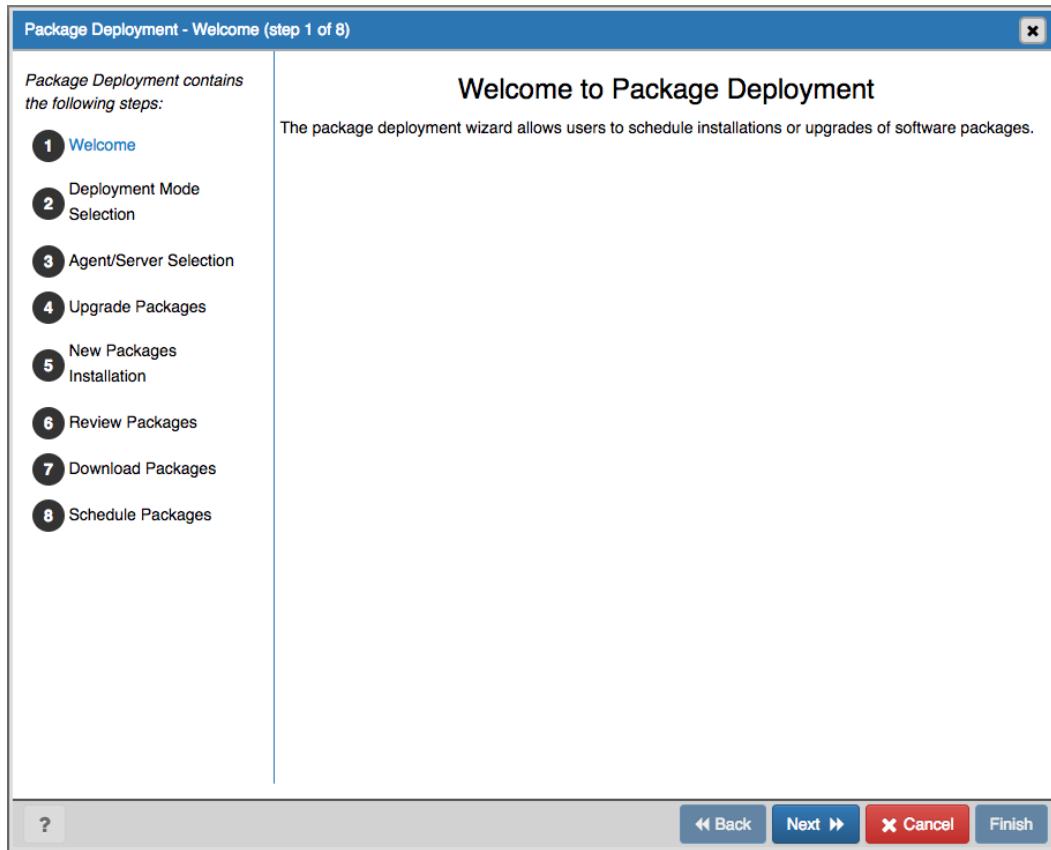


Figure 3.1 - The Package Deployment Welcome dialog.

Click **Next** to continue.

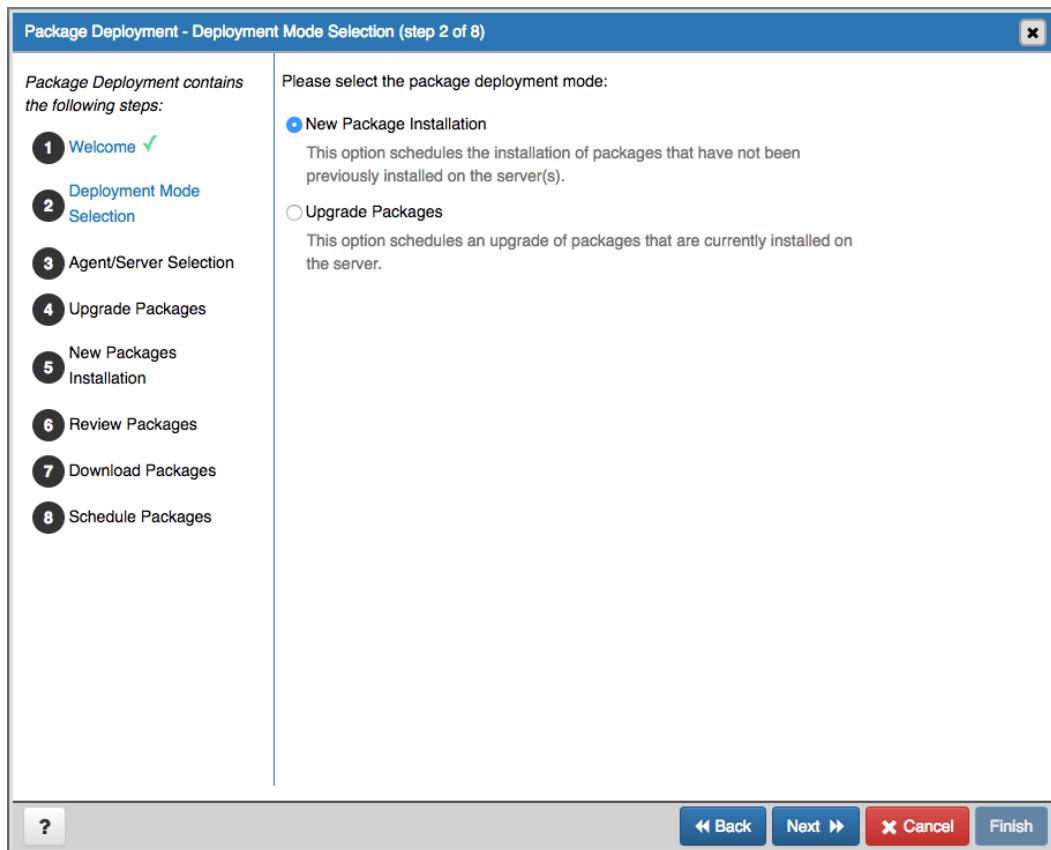


Figure 3.2 - The deployment mode selection dialog.

Use the radio buttons on the `Deployment Mode Selection` dialog (see Figure 3.2) to specify the type of deployment that you are scheduling:

- Select the `New Package Installation` radio button to schedule the installation of a package that has not been previously installed on the server. This is the default.
- Select the `Upgrade Packages` radio button to schedule an upgrade of packages that are currently installed on the server.

When you've made a selection, click `Next` to continue.

3.1 Installing a New Package

If you select **New Package Installation** on the **Deployment Mode Selection** dialog, the **Package Deployment wizard** opens the **Agent/Server Selection** dialog (see Figure 3.3), allowing you to specify the agents and servers on which the new applications will be installed.

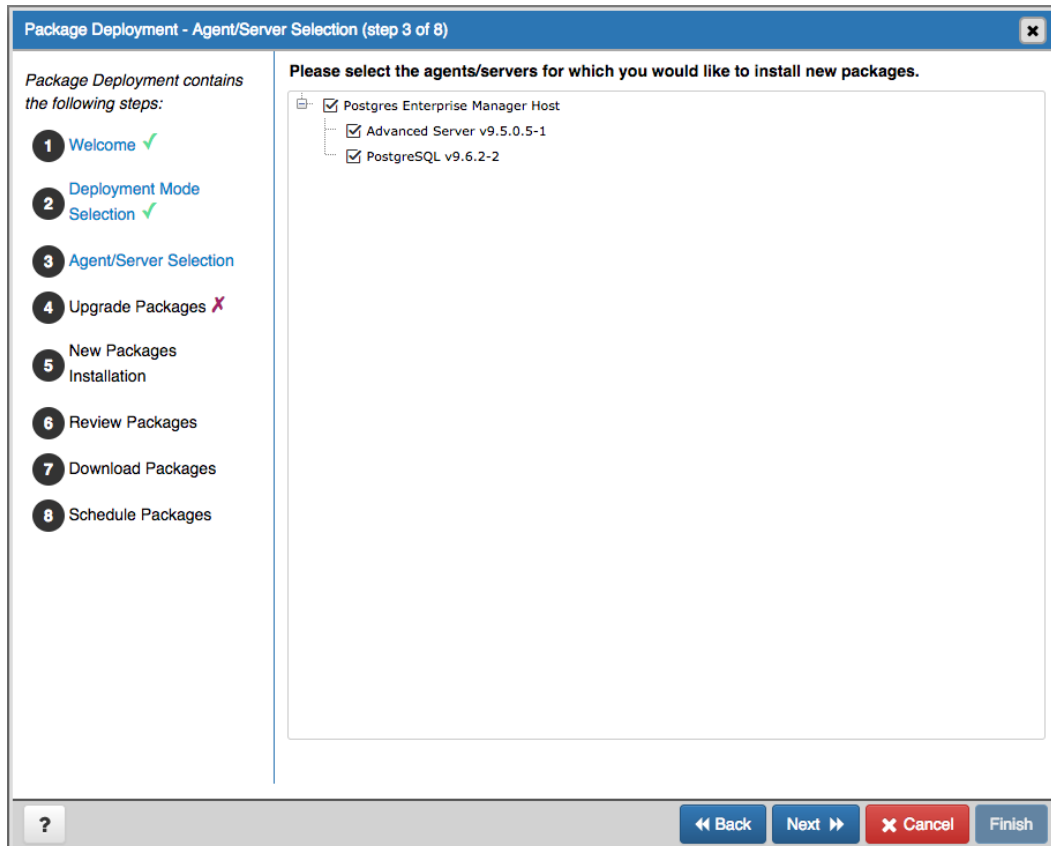


Figure 3.3 - Specify the target Agents and Servers.

Expand the tree control, and check the box next to each server on which you wish to install a new package and click **Next**.

The `New Packages Installation` dialog opens (see Figure 3.4).

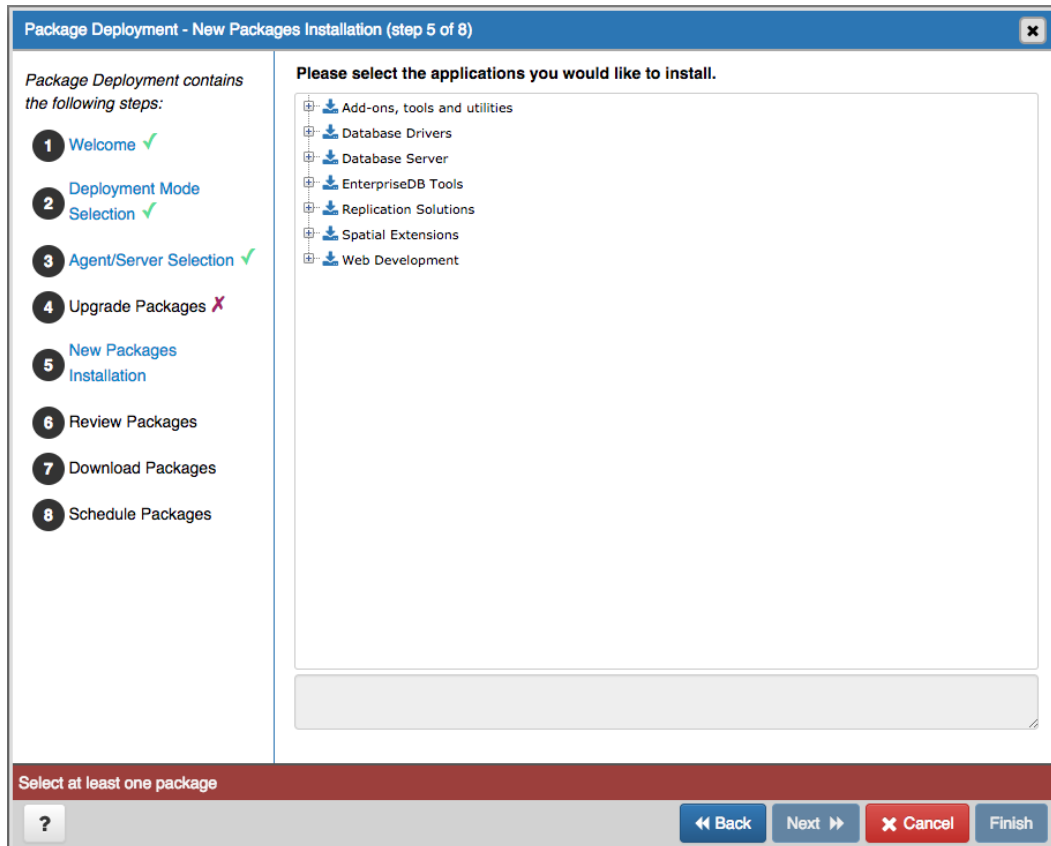


Figure 3.4 - Select which applications are to be installed.

Expand the tree control to review a list of applications that are available for installation. Check the box next to an application name to mark the application for installation. Note that the Package Deployment wizard will automatically check the boxes next to any supporting applications required by the applications you select.

When you've selected all of the packages you wish to add, click `Next` to continue.

Package Deployment - Review Packages (step 6 of 8)

Package Deployment contains the following steps:

- 1 Welcome ✓
- 2 Deployment Mode Selection ✓
- 3 Agent/Server Selection ✓
- 4 Upgrade Packages ✗
- 5 New Packages Installation ✓
- 6 Review Packages
- 7 Download Packages
- 8 Schedule Packages

Review your selections and provide installer options for each selected package.

Postgres Enterprise Manager Host

- ✓ Advanced Server v9.5.6.11-1
- ✓ PostGIS v2.1.8-2

Note: Options shown in bold are required.

Option Name	Option Value
Installation Directory	\$PATH
HOST	localhost
User Name	enterprisedb
Password	*****
Port	\$PORT
Language selection	en
Extract Only	No
Unattended Mode UI	minimal
Debug information level of verbosity	2

Installation Directory is required for package "Advanced Server v9.5.6.11-1" of agent "Postgres Enterprise Manager Host"

? << Back Next >> ✗ Cancel Finish

Figure 3.5 - Specify installation options.

Review the list of packages that will be installed, and (if prompted) provide any options requested (see Figure 3.5). Click **Next** to continue.

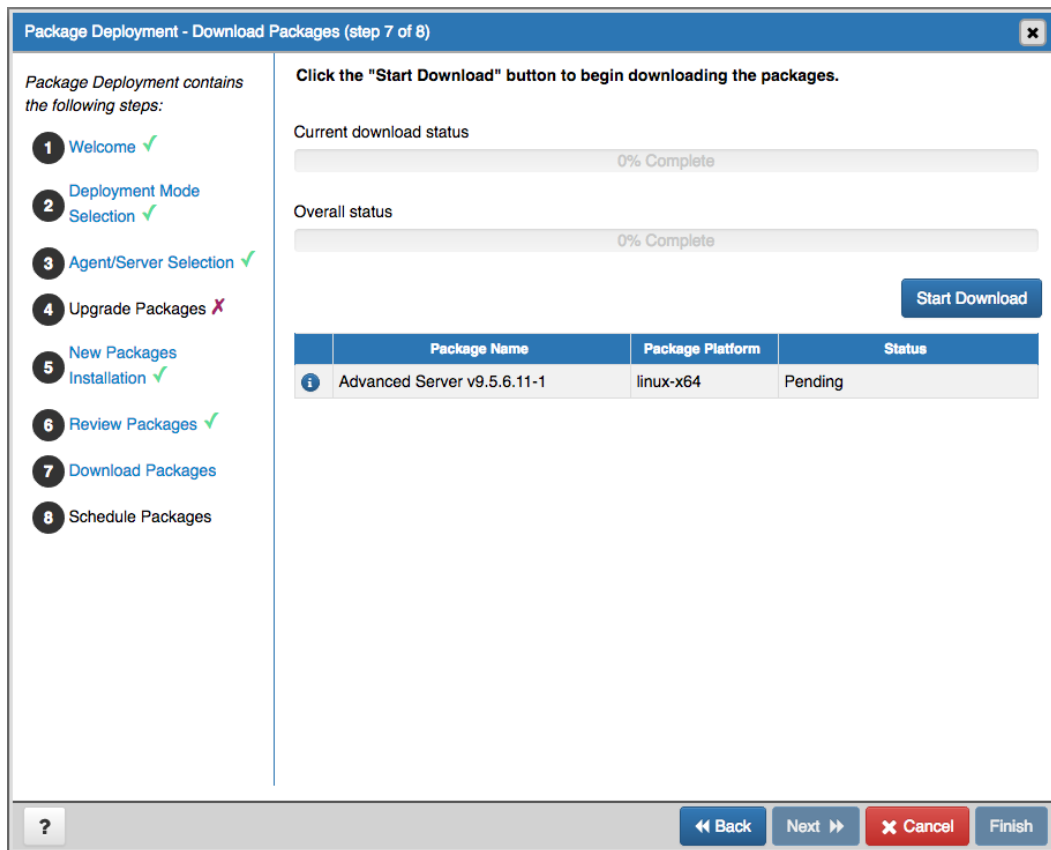


Figure 3.6 - Starting the installer download.

Click the `Start Download` button on the `Download Packages` dialog to instruct the `Package Deployment` wizard to download application installers (see Figure 3.6). During the download, you can click the `Cancel Download` button to abort the batch download. When the download completes, click `Next`.

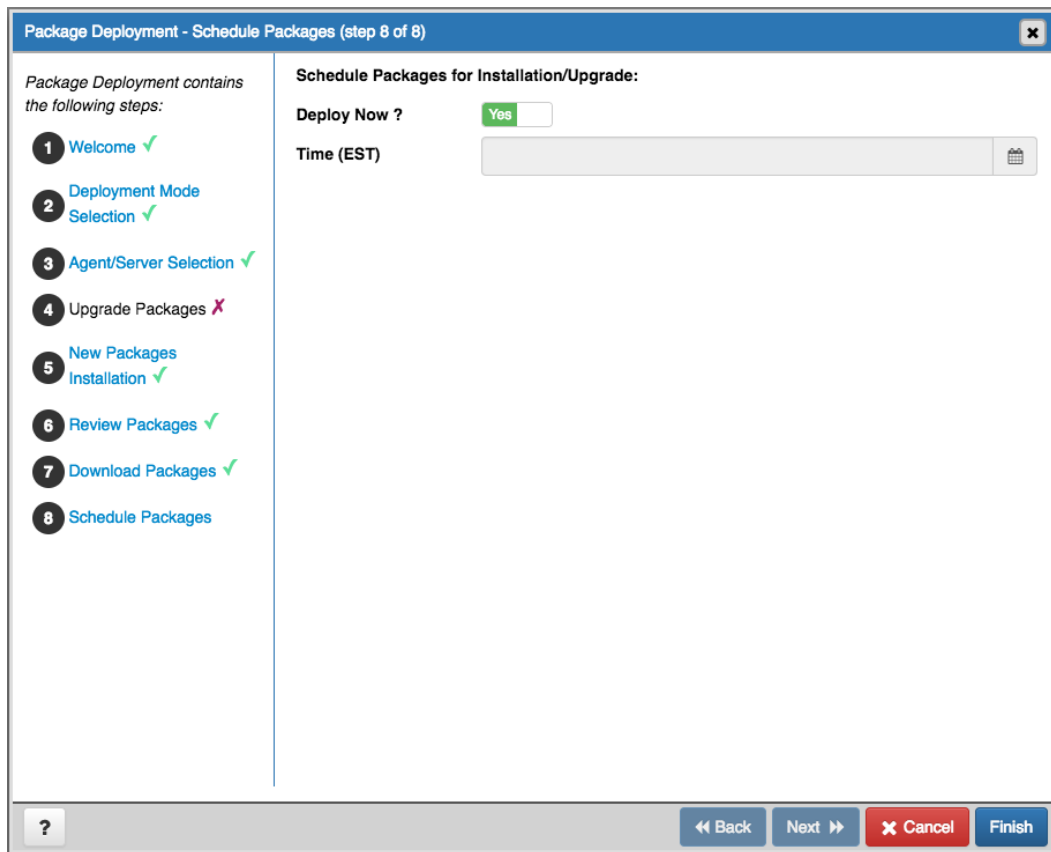


Figure 3.7 - The Package Deployment scheduling dialog.

Use the options on the scheduling dialog (see Figure 3.7) to schedule an installation time for the new packages:

- Set the slider next to `Deploy Now` to `Yes` to instruct the respective PEM agents to install the downloaded packages immediately. Please note that if a package requires a server restart, current user sessions may be interrupted.
- Set the slider next to `Deploy Now` to `No` and use the `Time` selector to specify a later date and time that you would like the package installation to begin.

Click `Finish` to install the downloaded packages or schedule the installation and exit the package deployment wizard.

If you have scheduled an installation for a later date/time, the scheduled task will be displayed on the `Scheduled Tasks` tab. To open the `Scheduled Tasks` tab, select `Scheduled Tasks...` from the `Management` menu.

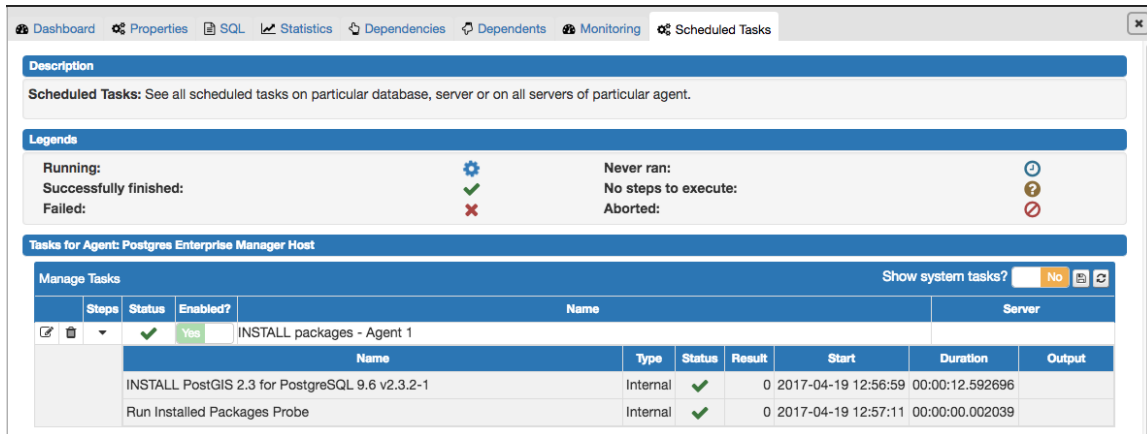


Figure 3.8 – The Scheduled Tasks tab.

3.1.1 Reviewing Scheduled Tasks

The `ScheduledTasks` tab (see Figure 3.8) features a legend, displaying the icons that identify the status of each task. The `ManageTasks` table displays a list of tasks that are pending execution or recently completed.

Set `Show system tasks?` to `Yes` to display system tasks; if it is set to `No`, only user-defined tasks are displayed. System tasks are displayed with a grey background, and may not be modified.

Highlight the name of a user-defined task and select the `Edit` icon to access detailed information about the selected task:

- Use the `Steps` drop-down to view a list of the steps performed during the selected task.
- The `Status` field lists the status of the current task.
- The `Enabled?` switch displays `Yes` if the task is enabled; `No` if the task is disabled.
- The `Name` field displays the name of the task.
- The `Agent` field displays the name of the agent responsible for executing the task.
- The `Last run` field displays the date and time of the last execution of the task.
- The `Next run` field displays the date and time of the next scheduled execution of the task.

- The `Created` field displays the date and time that the task was defined.

To remove a task, click the `Delete` icon located to the left of a task's name. The task will be marked for deletion, and removed when the tab refreshes.

3.2 Upgrading an Installed Package

If you select `Upgrade Packages` on the `Deployment Mode Selection` dialog, the `Package Deployment` wizard opens a dialog that allows you to specify which agents and packages will be updated.

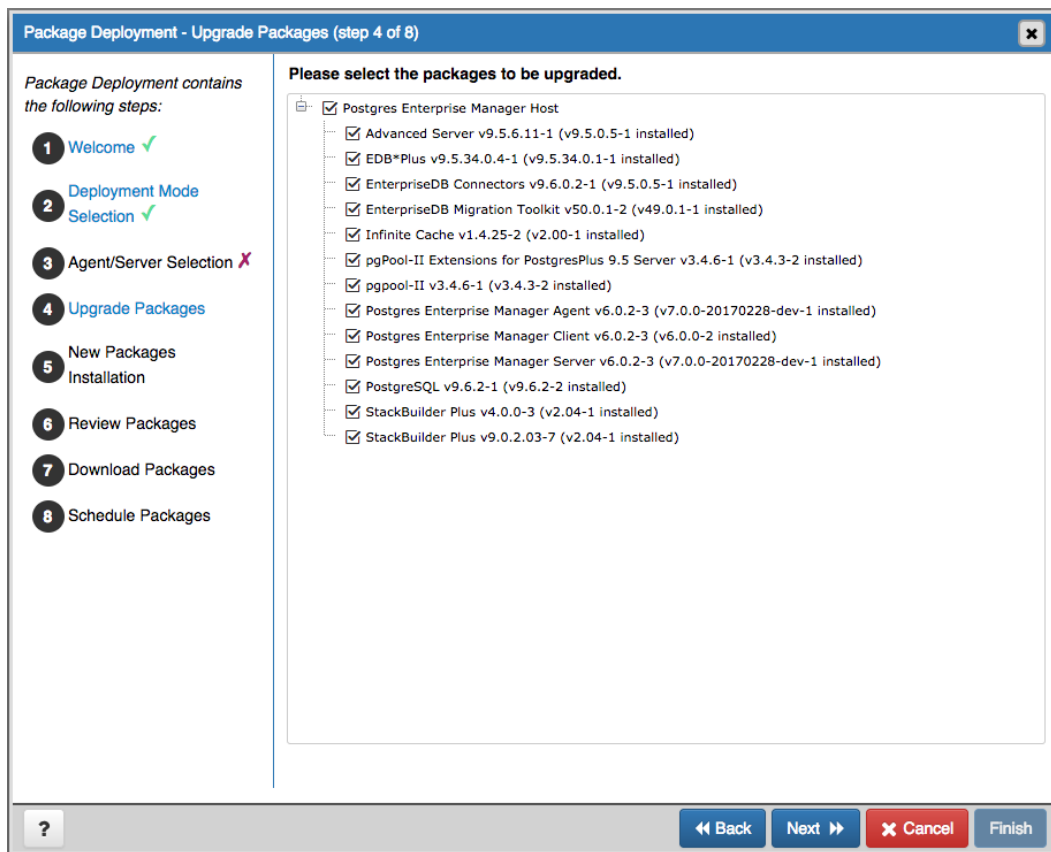


Figure 3.9 - Select the packages that will be upgraded.

Expand the tree control, and check the box next to package you wish to upgrade (see Figure 3.9). Click `Next` to continue.

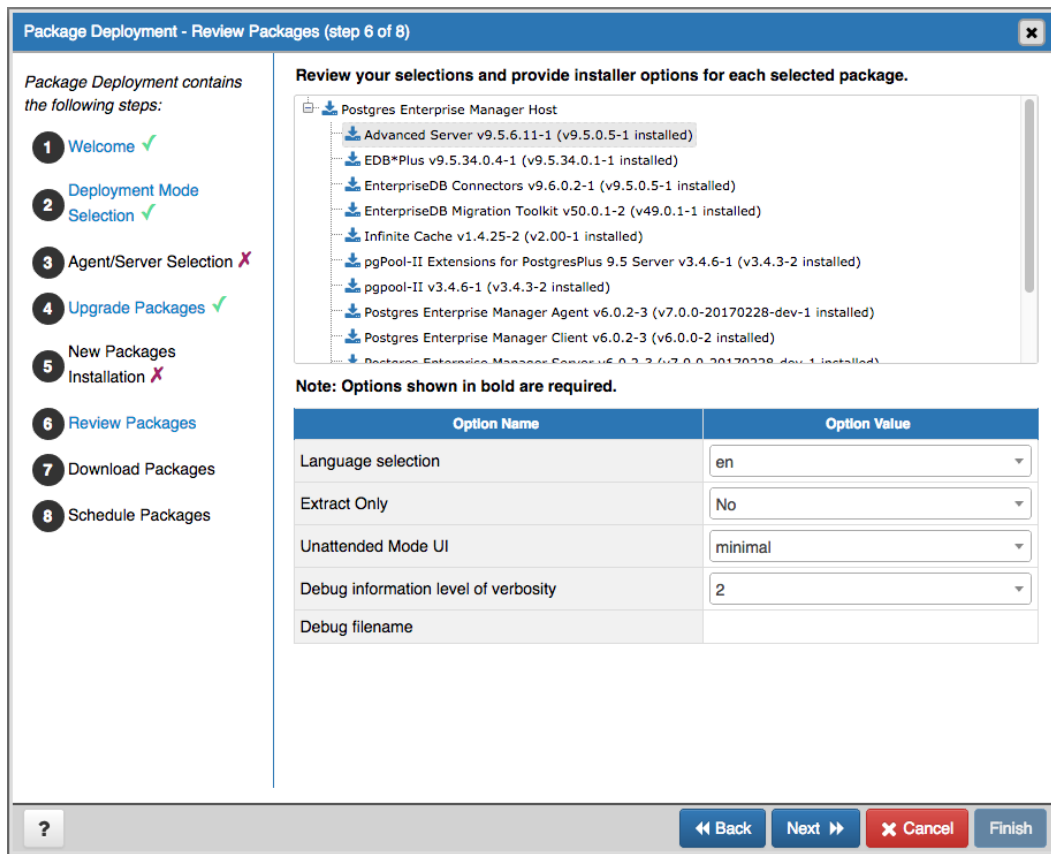


Figure 3.10 - Provide any requested installation options.

Review the list packages, and provide any installation options requested in the `Option Name/Option Value` fields (see Figure 3.10). When you've reviewed the list, click `Next` to continue.

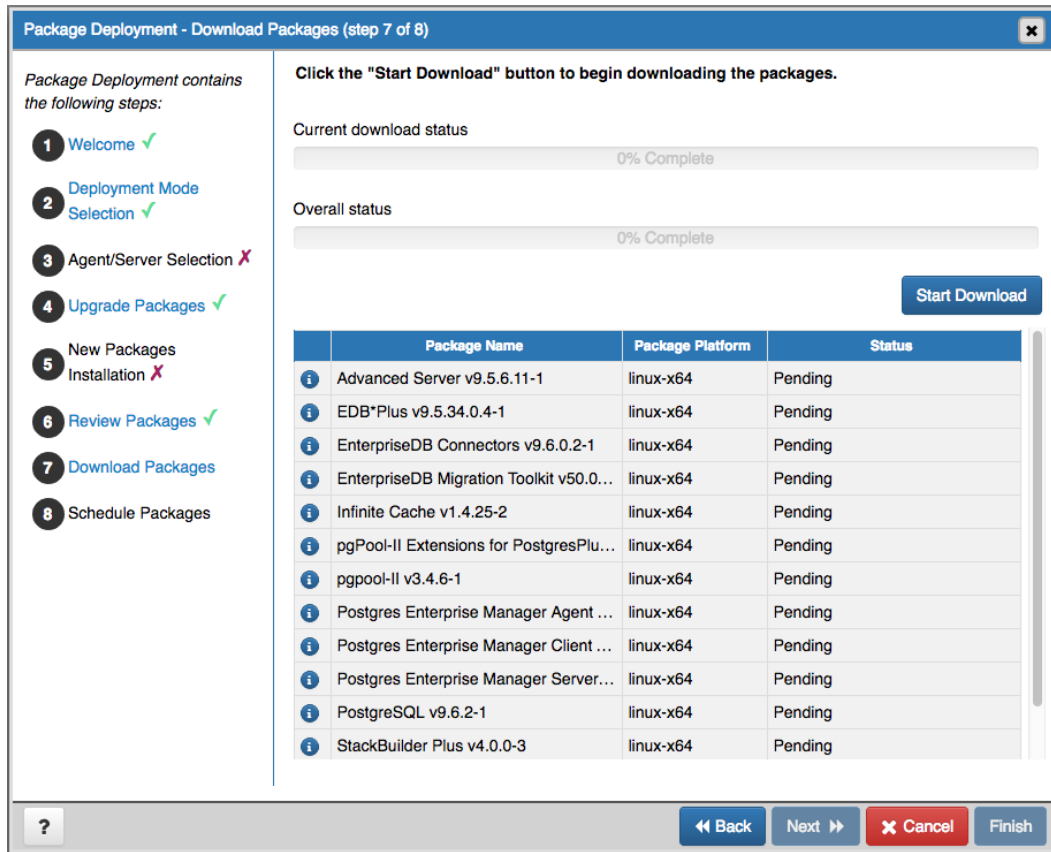


Figure 3.11 - Downloading the application installers.

Click the **Start Download** button to instruct the Package Deployment wizard to download application installers (see Figure 3.11). During the download, you can click the **Cancel Download** button to abort the batch download. When the download completes, click **Next**.

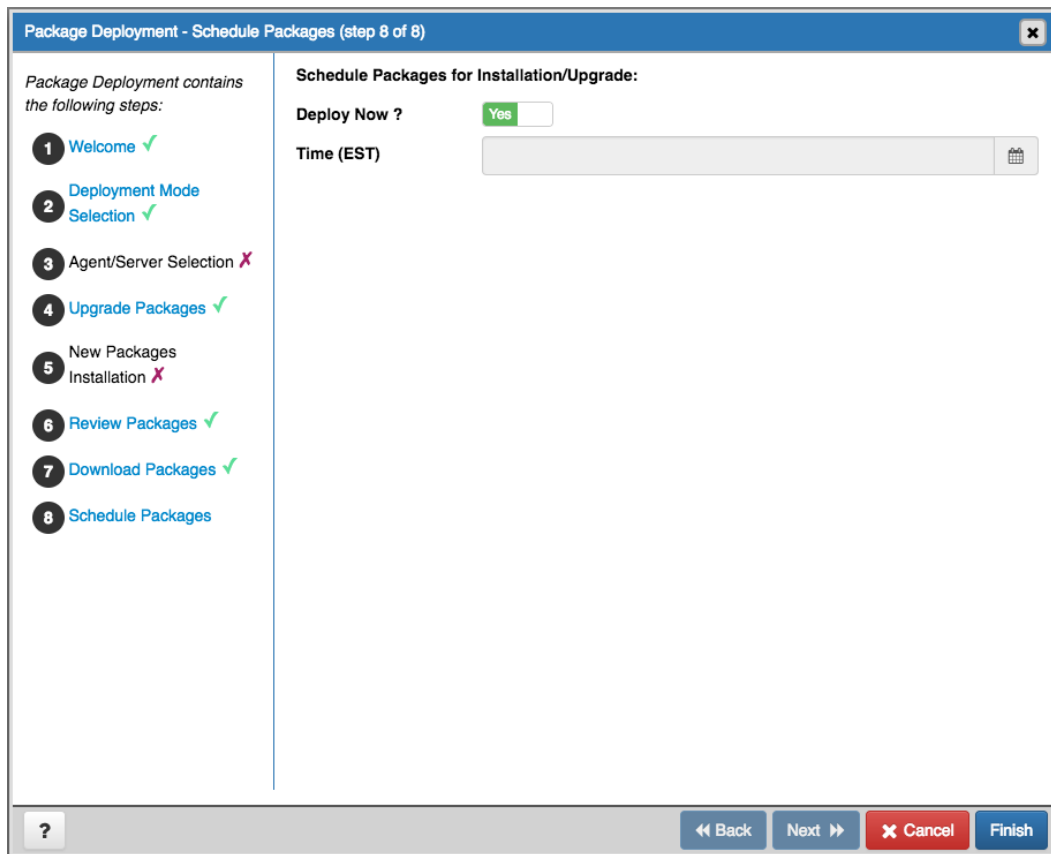


Figure 3.12 - Schedule a time for installation.

Use the options on the scheduling dialog (see Figure 3.12) to schedule an installation time for the new packages:

- Set the slider next to `Deploy Now` to `Yes` to instruct the respective PEM agents to install the downloaded packages immediately. Please note that if a package requires a server restart, current user sessions may be interrupted.
- Set the slider next to `Deploy Now` to `No` and use the `Time` selector to specify a later date and time that you would like the package installation to begin.

Click `Finish` to install the downloaded packages or schedule the installation and exit the package deployment wizard. If you have scheduled the update for a later date/time, the scheduled task will be included on the `Scheduled Tasks` dialog (accessed through the `Scheduled Tasks...` menu selection on the agent's context menu).

4 Performance Monitoring and Management

PEM contains built-in functionality that implements enterprise-wide performance monitoring of all managed servers. While you can customize many aspects of the various performance monitoring aspects of PEM, you can also elect to accept the recommended defaults that come out-of-the-box with the product.

4.1 Using Dashboards to View Performance Information

PEM displays performance statistics through a number of *dashboards*; each dashboard contains a series of summary views that contain charts, graphs and tables that display the statistics related to the selected object.

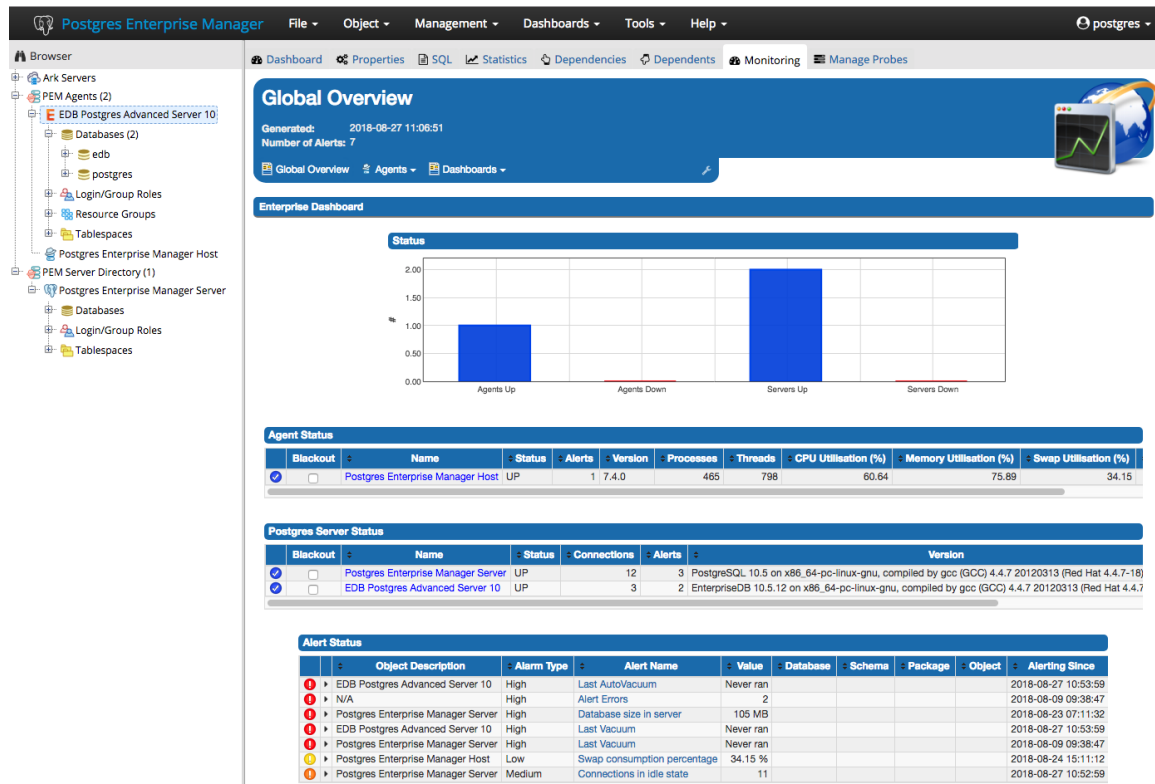


Figure 4.1 - The Global Overview dashboard.

The top-level dashboard is the `Global Overview` (shown in Figure 4.1). The `Global Overview` presents a status summary of all the servers and agents that are being monitored by the PEM server, a list of the monitored servers, and the state of any currently triggered alerts. The PEM client displays the `Global Overview` when it connects to the PEM server. Other dashboards provide statistical information about monitored objects. These include the:

Alerts Dashboard

The `Alerts` dashboard displays the currently triggered alerts. If opened from the `Global Overview`, the dashboard displays the current alerts for all monitored nodes on the system; if opened from a node within a server, the report will reflect alerts related to that node, and all monitored objects that reside below that object in the tree control.

Audit Log Analysis dashboard

For Advanced Server users, the `Audit Log Analysis` dashboard allows you to browse the audit logs that have been collected from instances that have audit logging and collection enabled.

Database Analysis dashboard

The `Database Analysis` dashboard displays performance statistics for the selected database.

I/O Analysis dashboard

The `I/O Analysis` dashboard displays I/O activity across various areas such as object DML activity, log operations and more.

Memory Analysis dashboard

The `Memory Analysis` dashboard supplies statistics concerning various memory-related metrics for the Postgres server.

Object Activity Analysis dashboard

The `Object Activity Analysis` dashboard provides performance details on tables/indexes of a selected database.

Operating System Analysis dashboard

The `Operating System Analysis` dashboard supplies information regarding the performance of the underlying machine's operating system.

Probe Log Analysis Dashboard

The `Probe Log Analysis` dashboard displays any error messages returned by a PEM agent.

Server Analysis dashboard

The `Server Analysis` dashboard provides general performance information about the overall operations of a selected Postgres server.

Server Log Analysis dashboard

The `Server Log Analysis` dashboard allows you to filter and review the contents of server logs that are stored on the PEM server.

Session Activity Analysis dashboard

The `Session Activity Analysis` dashboard provides information about the session workload and lock activity for the selected server

Session Waits Analysis dashboard

The `Session Waits Analysis` dashboard provides an overview of the current DRITA wait events for an Advanced Server session.

Storage Analysis dashboard

The `Storage Analysis` dashboard displays space-related metrics for tablespaces and objects.

System Waits Analysis dashboard

The `System Waits Analysis` dashboard displays a graphical analysis of system wait information for an Advanced Server session.

Streaming Replication Analysis dashboard

The `Streaming Replication Analysis` dashboard displays statistical information about WAL activity for a monitored server and allows you to monitor the status of Failover Manager clusters.

There are two ways to open a dashboard; you can:

- Select an active dashboard name from the `Dashboards` menu (accessed via the `Management` menu).
- Right click on the name of a monitored object in the tree control and select the name of the dashboard you would like to review from the `Dashboards` menu.

Each dashboard is displayed on the `Monitoring` tab in the main panel of the client window. After opening a dashboard, you can navigate to other dashboards within the same tab.

Each dashboard header includes navigation menus that allow you to navigate to other dashboards; use your browser's forward and back icons to scroll through previously-viewed dashboards. Use the `Refresh` icon to update the current dashboard.

Options on the `Dashboard Configuration` dialog allow you to link the time lines of all of the line graphs on the dashboard. To open the `Dashboard Configuration` dialog (see Figure 4.2), click the wrench icon displayed in the dashboard header.

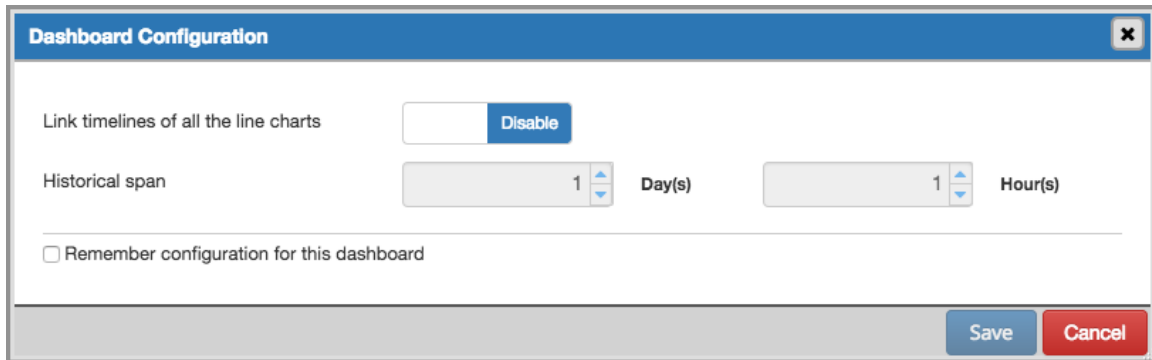


Figure 4.2 - The Dashboard Configuration dialog.

Use fields on the `Dashboard Configuration` dialog to control attributes of the charts displayed on the dashboard:

- Set the `Link timelines of all the line charts` slider to `Enable` to indicate that the specified timeline should be applied to line graphs displayed on the dashboard; if set to `Disable`, your preferences will be preserved for later use, but will not modify the amount of data displayed.
- Use the `Days` selector to specify the number of days of gathered data that should be displayed on line graphs.
- Use the `Hour(s)` selector to specify the number of hours of gathered data that should be displayed on line graphs.
- Check the box next to `Remember configuration for this dashboard` to indicate that the customized time span should be applied to the current dashboard only; if left unchecked, the time span will be applied globally to line graphs on all dashboards.

Please note that settings specified on the `Dashboard Configuration` dialog are applied only to the current user's session.

4.2 Managing Custom Dashboards

PEM displays performance statistics through a number of system-defined dashboards; each dashboard contains a series of summary views that contain charts, graphs and tables that display statistics related to the selected object. You can use the Manage Dashboards tab (see Figure 4.3) to create and manage custom dashboards that display the information that is most relevant to your system.

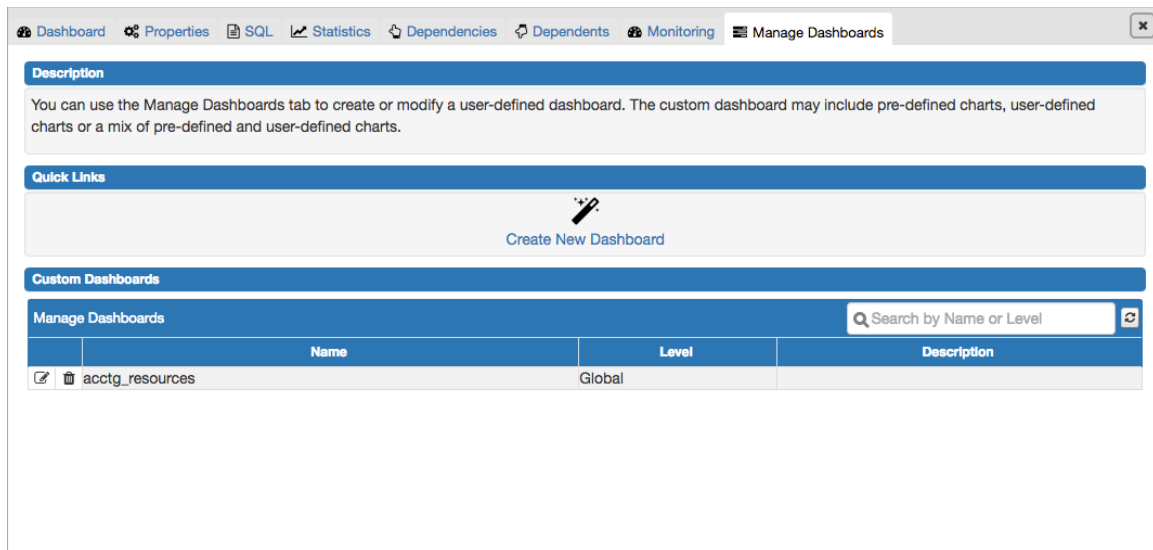


Figure 4.3 – The Manage Dashboards tab.

To create a custom dashboard, click the `Create New Dashboard` link (located in the `Quick Links` section of the `Manage Dashboards` tab).

To modify an existing dashboard, click the edit icon to the left of a dashboard name. The dashboard editor will open, displaying the definition of the dashboard. When you've finished modifying the dashboard's definition, click the `Save` button to preserve your changes; click `Cancel` to exit without saving your changes.

To delete a dashboard, click the delete icon to the left of a dashboard name. A popup will ask you to confirm that you wish to delete the dashboard; click `OK` to delete the selected dashboard.

4.2.1 Creating a Custom Dashboard

You can use the PEM dashboard editor to create or modify a user-defined dashboard. The custom dashboard may include pre-defined charts, user-defined charts or a mix of pre-defined and user-defined charts.

Figure 4.4 – The Create Dashboard editor.

Use the fields in the `Configure` section (see Figure 4.4) to specify general information about the dashboard:

- Specify a name for the dashboard in the `Name` field. The name specified will also be the title of the dashboard if the title is displayed.
- Use the `Level` drop-down listbox to specify the level of the PEM hierarchy within the PEM client on which the dashboard will be displayed. A dashboard

may be accessed via the `Dashboards` menu on a `Global` level, an `Agent` level, the `Server` level or the `Database` level. Each selected level within the list will expose a different set of metrics on which the custom dashboard's charts may be based.

- Provide a description of the dashboard in the `Description` field.

Provide information in the fields in the `Ops dashboard options` box if the dashboard will be used as an `Ops` dashboard:

- Set the `Ops Dashboard?` field to `Yes` to instruct the server to create a dashboard that is formatted for display on an `Ops` monitor.
- Set the `Show Title?` field to `Yes` to display the dashboard name at the top of the `Ops` dashboard.
- Use the `Font` drop-down list box to select a custom font style for the title. The selected font style will be displayed in the `Preview` box.
- Use the `Font size` drop-down list box to select a custom font size for the title. The selected font style will be displayed in the `Preview` box.

Use the `Permissions` box to specify the users that will be able to view the new dashboard:

- Set the `Share with all` slider to `Yes` to instruct the server to allow all `Teams` to access the dashboard, or set `Share with all` to `No` to enable the `Access permissions` field.
- Use the `Access permissions` field to specify which roles can view the new dashboard. Click in the field, and select from the list of users to add a role to the list of users with dashboard access.

When you've completed the `Configure Dashboard` section, click the arrow in the upper-right corner to close the section, and access the `Dashboard Layout Design` section.

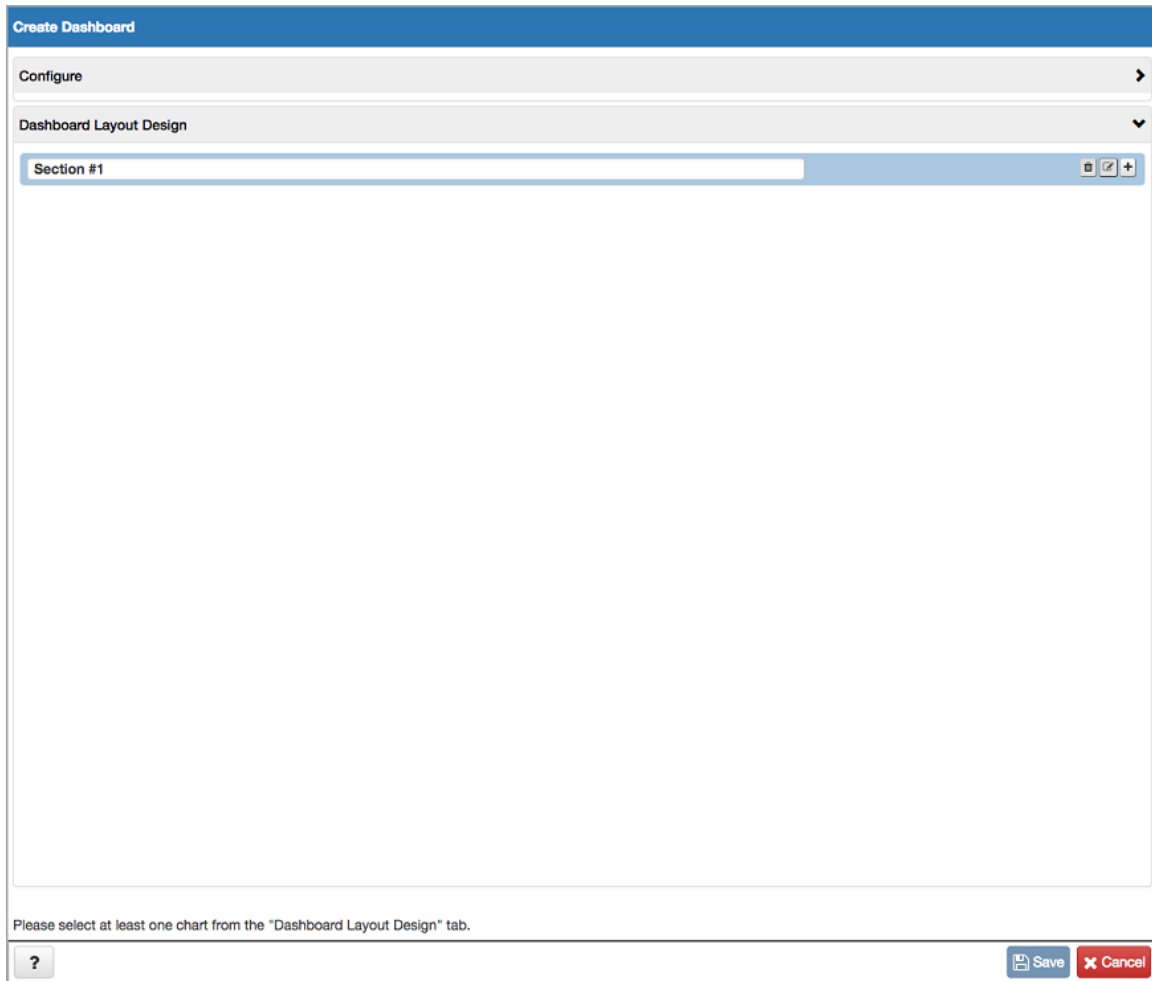


Figure 4.5 – Modifying a Section Header.

Click the edit icon in a section header to specify a section name (see Figure 4.5); then, click the add icon (+) to add a chart to the section (see Figure 4.6).

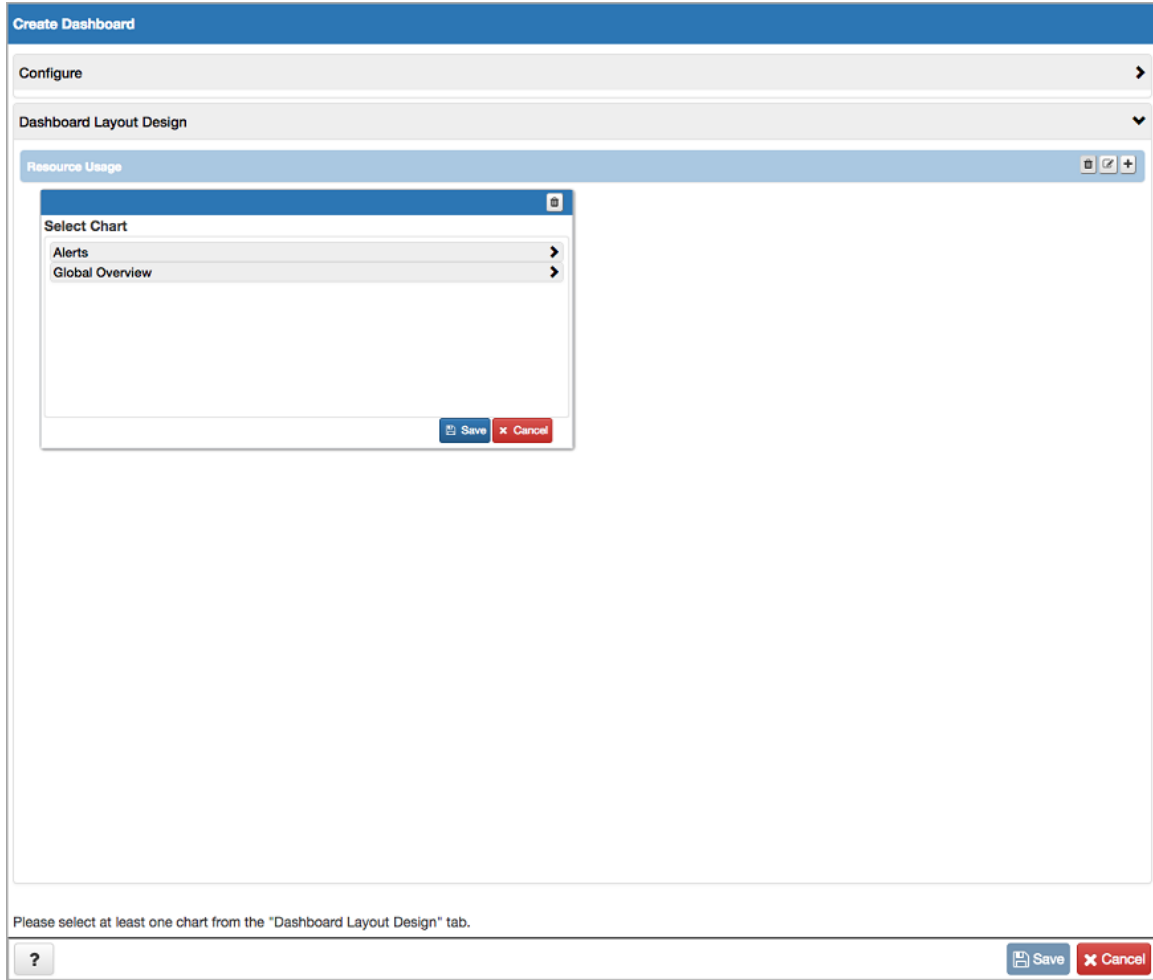


Figure 4.6 – Adding a Chart.

Use the arrows to the right of each chart category to display the charts available and select a chart.

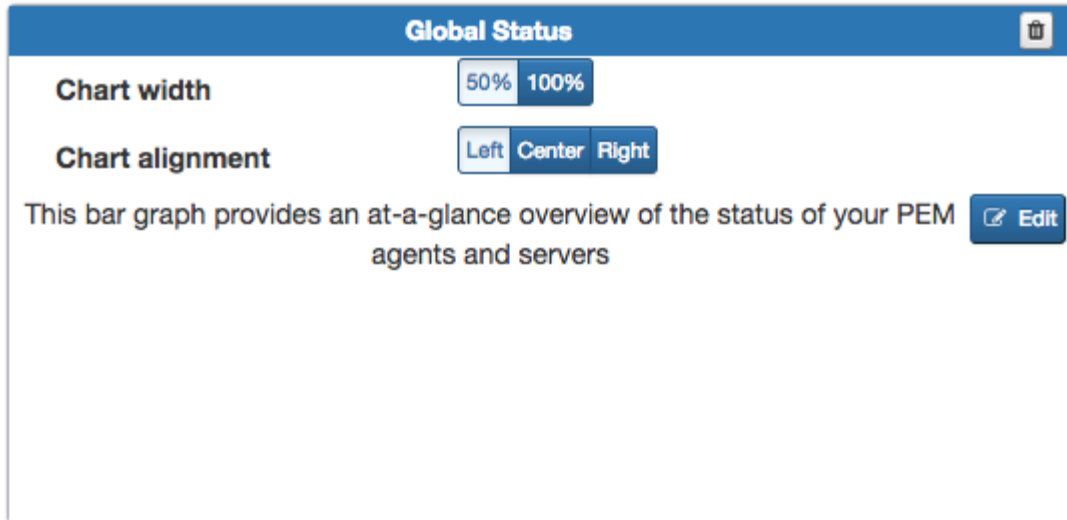


Figure 4.7 – Specifying placement details for a chart.

Use the chart detail selectors to specify placement details for the chart:

- Use the `Chart width` selector to indicate the width of the chart; select `50%` to display the chart in half of the dashboard, or `100%` to use the whole dashboard width.
- Use the `Chart alignment` selector to indicate the position of the chart within the section:

Select `Left` to indicate that the chart should be left-justified.

Select `Center` to indicate that the chart should be centered.

Select `Right` to indicate that the chart should be right-justified.

Please note that tables are always displayed centered.

When creating or editing a custom dashboard, you can use drag and drop to re-arrange the charts within a section or to move a chart to a different section.

To add another chart to your dashboard, click the add icon (+) in the section header.

When you've finished editing the dashboard, click the `Save` button to save your edits and exit.

To exit without saving your changes, click the `Cancel` button.

4.2.2 Creating an Ops Dashboard

You can use the PEM dashboard editor to create a custom dashboard formatted for display on an Ops monitor. An Ops dashboard displays the specified charts and graphs, while omitting header information and minimizing extra banners, titles, and borders.

The screenshot shows a dialog box titled "Ops dashboard options". It contains the following elements:

- Ops dashboard?**: A radio button followed by a blue button labeled "No".
- Show title?**: A radio button followed by a blue button labeled "No".
- Font**: A text input field with a dropdown arrow on the right.
- Font size**: A text input field with a dropdown arrow on the right.
- A large, empty rectangular area at the bottom, likely a preview box, with a small edit icon (pencil) in the bottom right corner.

Figure 4.8 –Ops dashboard options.

To create an Ops dashboard, provide detailed information about the Ops display in the `Ops dashboard options` section of the `Create Dashboard` dialog (see Figure 4.8).

- Set the `Ops Dashboard?` field to `Yes` to instruct the server to create a dashboard that is formatted for display on an Ops monitor.
- Set the `Show Title?` field to `Yes` to display the dashboard name at the top of the Ops dashboard.
- Use the `Font` drop-down list box to select a custom font style for the title. The selected font style will be displayed in the `Preview` box.
- Use the `Font size` drop-down list box to select a custom font size for the title. The selected font style will be displayed in the `Preview` box.

After adding charts and tables to the Ops dashboard, click the `Save` button to save your work. You can then access the dashboard by navigating through the `Dashboards` menu of the hierarchy level specified in the `Level` field on the `New Dashboard` dialog.

4.3 Using the Manage Charts tab

You can use the `Manage Charts` tab to access dialogs that allow you to create or modify a custom line chart or table, or import a Capacity Manager template for use in a custom chart. After defining a chart, you can display the chart on a custom dashboard. To open the `Manage Charts` tab, select `Manage Charts...` from the PEM client `Management` menu.

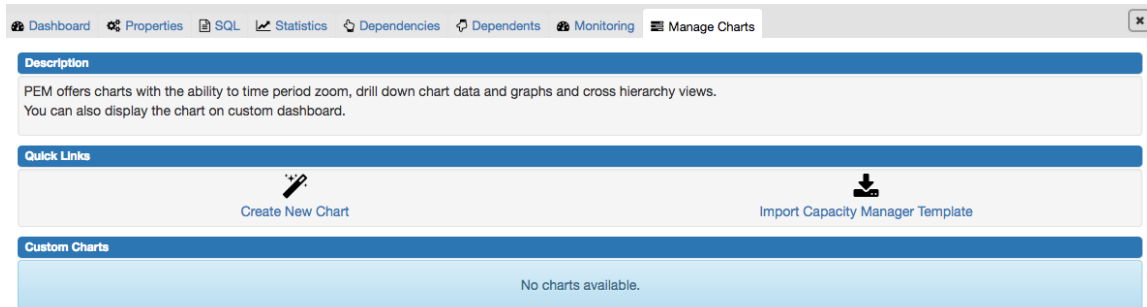


Figure 4.9 – The Manage Charts tab.

The `Manage Charts` tab (see Figure 4.9) provides a `Quick Links` menu that allows you to access dialogs to:

- Create a New Chart for use on a custom dashboard.
- Import a Capacity Manager template to use as a template for creating a custom chart.

The `Custom Charts` table (see Figure 4.10) displays a list of user-defined charts; when a chart is newly added, the font displays in green. When you add an additional chart or refresh the screen, the name of the chart is displayed in black.

Custom Charts				
<input type="text" value="Search chart by name, type, level or category"/>				
	Name	Type	Level	Metrics Category
	acctg_resource_usage	Line Chart	Agent	Database Server
	hr_resource_usage	Table Chart	Agent	Audit logs
	sales_resource_usage	Line Chart	Agent	Database Server

Figure 4.10 – The Custom Charts table.

Use the search box in the upper-right hand corner of the `Custom Charts` table to search through your custom charts. Specify a:

- Chart name
- Type
- Level
- Metrics Category

Use icons to the left of a chart's name in the `CustomCharts` table to manage a chart:

- Click the edit icon to open the Chart Configuration wizard and modify aspects of the chart or table.
- Click the delete icon to delete the selected chart.

4.3.1 Creating a Custom Chart

Click the `Create New Chart` icon in the `Quick Links` section of the `Manage Charts` tab to open the `Create Chart` wizard. The wizard will walk you through the steps required to define a new chart.

Figure 4.11 – Specifying general information about the chart.

Use the fields on the `Configure Chart` dialog (shown in Figure 4.11) to specify general information about the chart:

- Specify the name of the chart in the `Name` field.
- Use the drop-down listbox in the `Category` field to specify the category in which this chart will be displayed; when adding a custom chart to a custom dashboard, the chart will be displayed for selection in the category specified.
- Use the radio buttons in the `Type` field to specify if the chart will be a `Line chart` or a `Table`.

- Provide a description of the chart in the `Description` field. The description will be displayed to the user viewing the chart (on a custom dashboard) when they click the information icon.

When you've completed the fields on the `Configure Chart` dialog, click `Next` to continue.

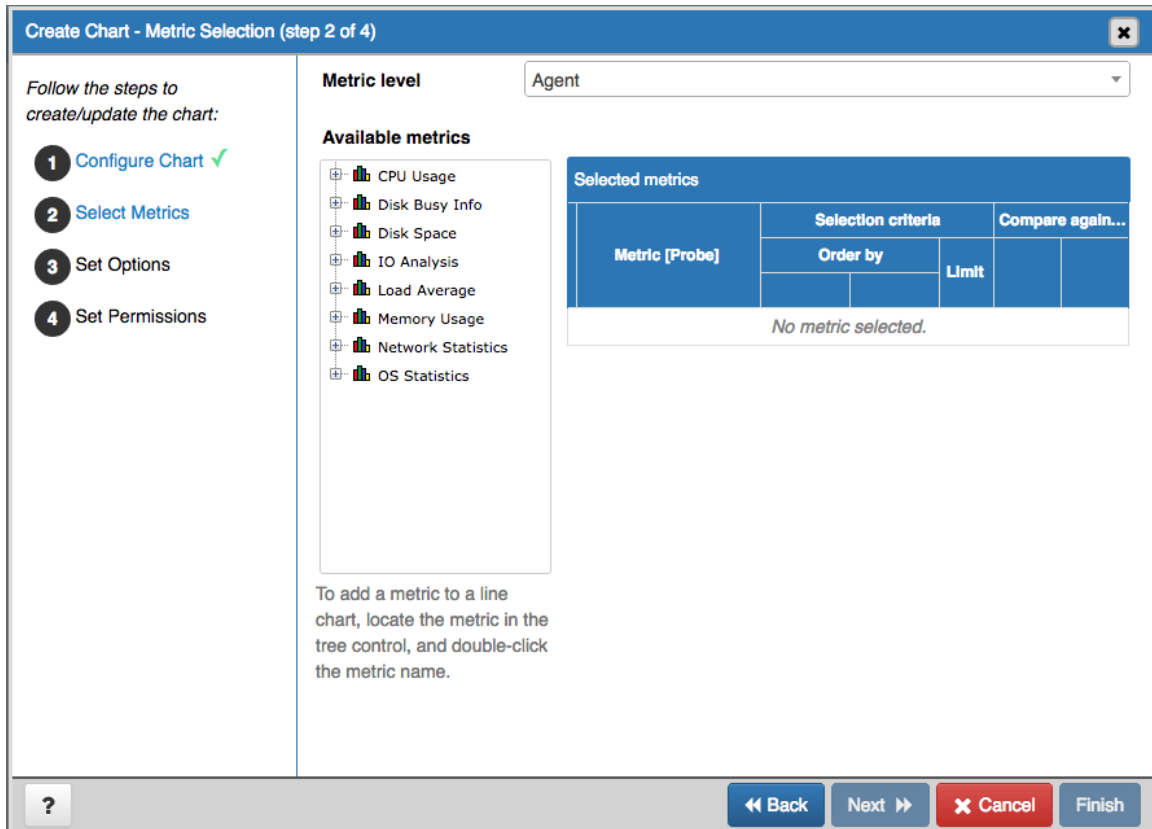


Figure 4.12 – Specifying the metrics that will be displayed.

Use the fields on the `Select Metrics` dialog (shown in Figure 4.12) to select the metrics that will be displayed on the chart.

- Use the `Metric level` drop-down listbox to specify the level of the PEM hierarchy from which you wish to select metrics. You can specify `Agent`, `Database`, or `Server`. Each level offers access to a unique set of probes and metrics.
- Use the tree control in the `Available metrics` box to select the metrics that will be displayed on the chart.

If you are creating a table, you may only select metrics from one probe; each node of the tree control lists the metrics returned by a single probe. Expand a node of

the tree control, and check the boxes to the left of a metric name to include that metric data in the table.

If you are creating a line chart, expand the nodes of the tree control and double-click each metric that you would like to include in the chart.

- Use the fields in the `SelectedMetrics` panel to specify how the metric data will be displayed in your chart. The selection panel displays the name of the metric in the (non-modifiable) `Metric [Probe]` column. You can:
 - Click the garbage can icon to delete a metric from the list of selected metrics.
 - Use the drop-down listboxes in the `SelectionCriteria` column to specify the order of the data displayed.
 - Use the `Limit` field to specify the number of rows in a table or lines in a chart:

The maximum number of lines allowed in a chart is 32.

The maximum number of rows allowed in a table is 100.

- If you are creating a line chart, PEM supports comparisons of cross-hierarchy metrics.
 - Click the compare icon to open a selection box that allows you to select one or more probe-specific attributes (i.e. CPUs, interfaces, databases, etc.) to compare in the chart.
 - Click the copy icon to apply your selections to all of the metrics for the same probe. When the popup opens, click `Yes` to confirm that other selections for the same probe will be overwritten, or `No` to exit the popup without copying the attributes.

When you've completed the fields on the `SelectMetrics` dialog, click `Next` to continue.

Create Chart - Chart Options (step 3 of 4)

Follow the steps to create/update the chart:

- 1 Configure Chart ✓
- 2 Select Metrics ✓
- 3 Set Options
- 4 Set Permissions

Auto refresh **Minute(s)**
Please specify the number of minutes between chart updates.

Line chart options

Points to plot
Please specify the maximum number of points to be plotted on the chart.

Historical span **Day(s)** **Hour(s)** **Minute(s)**
Please specify the number of minutes, hours or days of historical data that should be included on the chart.

Data extrapolation

☐ No Extrapolation ☒ Span ☐ Threshold

Data will not be extrapolated into the future.

Note: Extrapolated data is calculated using a linear trend analysis algorithm on the historical data. The accuracy of the extrapolated data depends on the available data in the historical span.

Navigation: ? Back Next Cancel Finish

Figure 4.13 – Specifying chart options.

Use the fields on the `Set Options` dialog (shown in Figure 4.13) to specify display options for your chart:

- Use the `AutoRefresh` field to specify the number of minutes between chart updates - choose a value from 1 to 120. The default auto refresh rate is 2 minutes.

Use fields under the `Line chart options` heading to specify display preferences for a line chart:

- Use the `Points to plot` field to specify the maximum number of points that will be plotted on the chart.
- Use the fields to the right of the `Historical span` label to specify how much historical data should be displayed on the chart:

Use the `Day(s)` field to specify the number of days of historical data that should be included on the chart.

Use the `Hour(s)` field to specify the number of hours of historical data that should be included on the chart.

Use the `Minute(s)` field to specify the number of minutes of historical data that should be included on the chart.

Use the fields in the `Data extrapolation` box to specify if PEM should generate extrapolated data based on historical data:

- Click the `No Extrapolation` label to omit extrapolated data from the chart.
- Click the `Span` label to use the `Days` and `Hours` selectors to specify the period of time spanned by the metrics on the chart.
- Click the `Threshold` label to use threshold selectors to specify a maximum or minimum value for the chart.

When you've completed the fields on the `Set Options` dialog, click `Next` to continue.

Create Chart - Security (step 4 of 4)

Follow the steps to create/update the chart:

- 1 Configure Chart ✓
- 2 Select Metrics ✓
- 3 Set Options ✓
- 4 Set Permissions

Share with all

Access permissions

Specify the user groups that will have access to the new chart.

◀ Back Next ▶

Figure 4.14 – Specifying access permissions.

Use the fields on the `Set Permissions` dialog (shown in Figure 4.14) to specify display options for your chart.

- Set the `Share with all` slider to `Yes` to indicate that the chart will be available to all authorized users, or `No` to restrict access to the users or groups specified in the `Access permissions` field.
- Use the `Access permissions` field to select the group or groups that will have access to the chart.

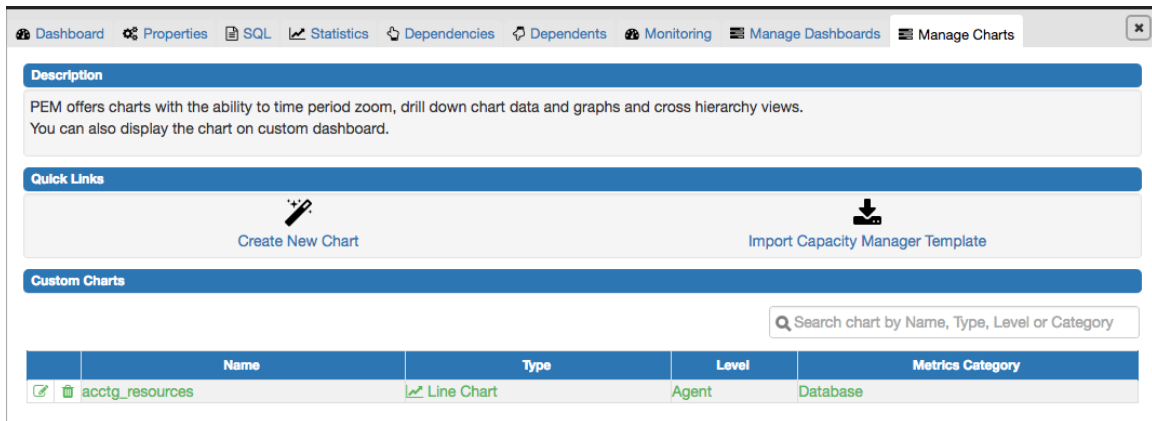


Figure 4.15 – The chart definition is displayed on the *Manage Charts* tab.

When you've finished defining the chart, click `Finish` to save your edits and add your chart to the list on the `Manage Charts` tab (see Figure 4.15).

4.3.2 Importing a Capacity Manager Template

Click the `Import Capacity Manager Template` icon in the `Quick Links` section of the `Manage Charts` tab to open the `Create Chart` dialog, and use a Capacity Manager template as a starting point for a chart or table.

Figure 4.16 – Importing a Capacity Manager template.

When the `Create Chart` dialog opens (see Figure 4.16), provide information about the custom chart:

- Use the drop-down listbox in the `Import capacity template` field to select the name of the template on which the chart will be based.
- Specify the name of the chart in the `Name` field.
- Use the drop-down listbox in the `Category` field to specify the category in which this chart will be displayed. When adding a custom chart to a custom dashboard, the chart will be displayed for selection in the `Category` specified.
- Use the radio buttons in the `Type` field to specify if the chart will be a `Line` chart or a `Table`.

- Provide a description of the chart in the `Description` field. The description will be displayed to the user viewing the chart (on a custom dashboard) when they click the information icon.

Click `Next` to continue to the `Select Metrics` dialog.

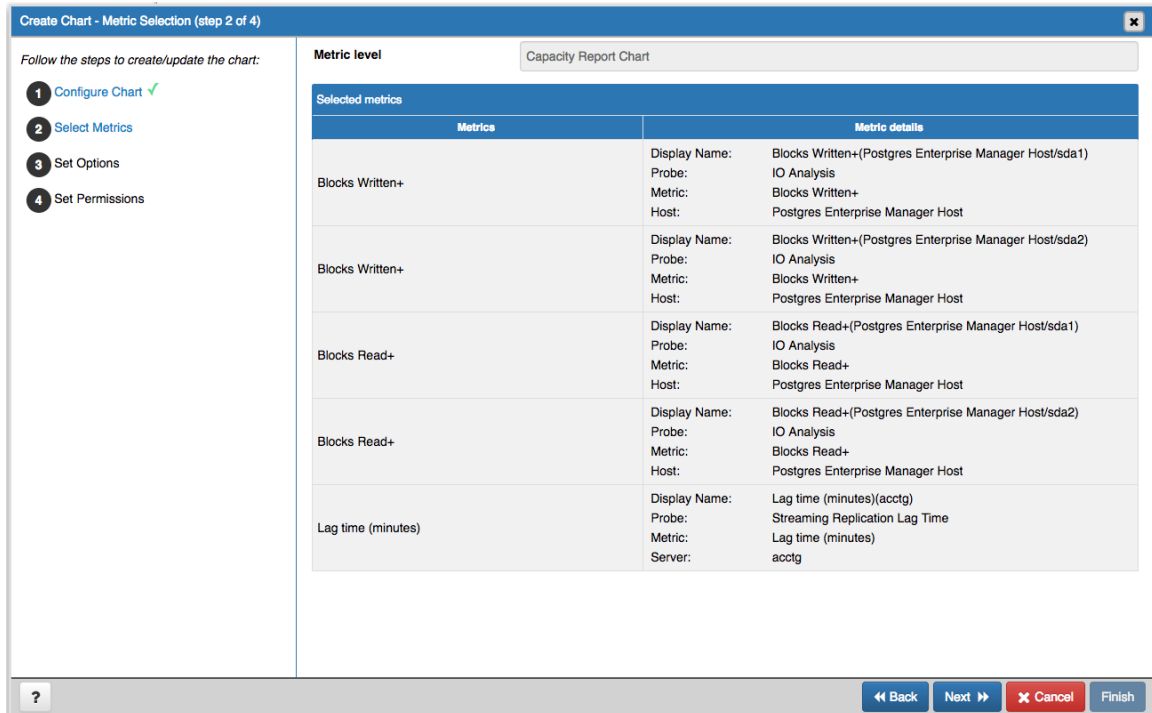


Figure 4.17 – The template metrics.

The `Select Metrics` window (see Figure 4.17) allows you to review the metrics specified by the selected template. The bottom panel of the chart editor displays the metrics that will be included in the chart. The metrics included in the chart are not modifiable via the chart editor; to modify the metrics, you must use the Capacity Manager utility to update the template.

When you've reviewed the metrics, click `Next` to continue to the `Set Options` dialog.

Figure 4.18 – Selecting chart options.

Use the fields on the `Set Options` window (see Figure 4.18) to specify display options for your chart:

- Use the `Auto Refresh` field to specify the number of minutes between chart updates - choose a value from 1 to 120. The default auto refresh rate is 2 minutes.

Use the fields in the `Data extrapolation` box to specify the time period covered by the chart. You can either:

- click the `Historical days and extrapolated days` label and provide:
 - the number of days of historical data that should be charted in the `Historical` field.
 - the number of projected days that should be charted in the `Extrapolated` field.
- or, click the `Historical days and threshold` label and provide:
 - the number of days of historical data that should be charted in the `Historical` field
 - the threshold value at which the chart will end.

When you've completed the `Set Options` window, click `Next` to continue.

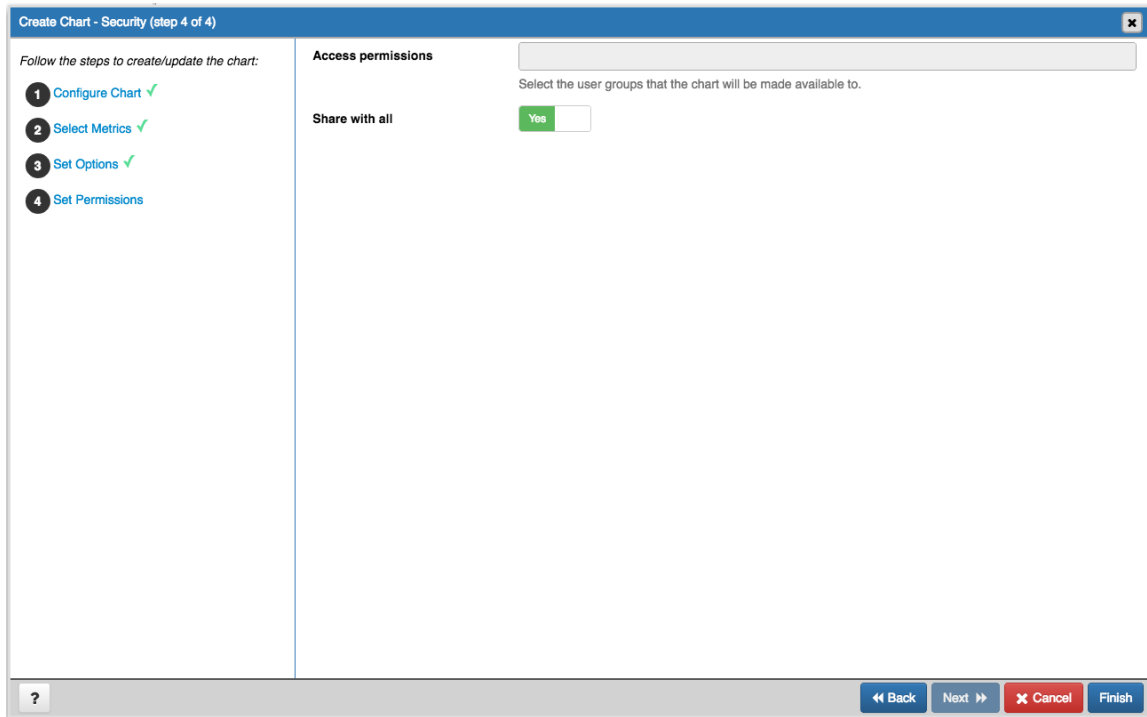


Figure 4.19 – Selecting permissions for the chart.

Use the fields on the `Set Permissions` window (see Figure 4.19) to specify display options for your chart:

- Set the `Share with all` slider to `Yes` to indicate that the chart will be available to all authorized users, or `No` to restrict access to the users or groups specified in the `Access permissions` field.
- Use the `Access permissions` field to select the group or groups that will have access to the chart.

When you've finished defining the chart, click `Finish` to save your edits and add your chart to the list on the `Manage Charts` tab.

4.4 Customizing Probes

A *probe* is a scheduled task that returns a set of performance metrics about a specific monitored object. A probe retrieves statistics from a monitored server, database, operating system or agent. You can use the `Manage Probes` tab (shown in Figure 6.13) to override the default configuration and customize the behavior of each probe.

To open the `Manage Probes` tab, select `Manage Probes...` from the `Management` menu. The `Manage Probes` tab opens in the PEM client (see Figure 4.20).

Probe name	Execution Frequency			Enabled?		Data Retention	
	Default	Minutes	Seconds	Default	Probe Enable?	Default	Days
Database Frozen XID	Yes	720	0	Yes	Yes	Yes	180
Database Size	Yes	30	0	Yes	Yes	Yes	180
Number of Prepared Transactions	Yes	5	0	Yes	Yes	Yes	180
Number of WAL Files	Yes	5	0	Yes	Yes	Yes	180
Object Catalog: Database	Yes	5	0	Yes	Yes	Yes	180
Object Catalog: Tablespace	Yes	5	0	Yes	Yes	Yes	180
Server Information	Yes	5	0	Yes	Yes	Yes	180
Settings	Yes	5	0	Yes	Yes	Yes	180
Tablespace Size	Yes	30	0	Yes	Yes	Yes	180
User Information	Yes	30	0	Yes	Yes	Yes	180

Figure 4.20 - The `Manage Probes` tab.

The `Manage Probes` tab provides a set of `Quick Links` that you can use to create and manage probes:

- Click the `Manage Custom Probes` icon to open the `Custom Probes` tab and create or modify a custom probe.
- Click the `Copy Probes` icon to open the `Copy Probe` dialog, and copy the probe configurations from the currently selected object to one or more monitored objects.

A probe monitors a unique set of metrics for each specific object type (server, database, database object, or agent); select the name of an object in the tree control to review the probes for that object.

To modify the properties associated with a probe, highlight the name of a probe, and customize the settings that are displayed in the `Probes` table:

- Move the `Default` switch in the `Execution Frequency` columns to `N` to enable the `Minutes` and `Seconds` selectors, and specify a non-default value for the length of time between executions of the probe.
- Move the `Default` switch in the `Enabled?` column to `No` to change the state of the probe, and indicate if the probe is active or not active.

Please note: If data from a `Disabled` probe is used in a chart, the chart will display an information icon in the upper-left corner that allows you to enable the probe by clicking the provided link.

- Move the `Default` switch in the `Data Retention` column to `No` to enable the `Day(s)` field and specify the number of days that information gathered by the probe is stored on the PEM server.

The `Manage Probes` tab may display information about probes that cannot be modified from the current node. If a probe cannot be modified from the current dialog, the switches are disabled. Generally, a disabled probe can be modified from a node that is higher in the hierarchy of the PEM client tree control; select another object in the tree control to modify which probes are displayed or enabled in the `Manage Probes` tab.

To review a list of pre-defined probes, see Section [15.3](#).

4.4.1 Creating a Custom Probe

You can use the PEM `Custom Probes` tab to create a new probe or modify an existing user-defined probe. To open the `Custom Probes` tab (see Figure 4.21), select the `Manage Custom Probes...` icon from the `Manage Probes` tab.

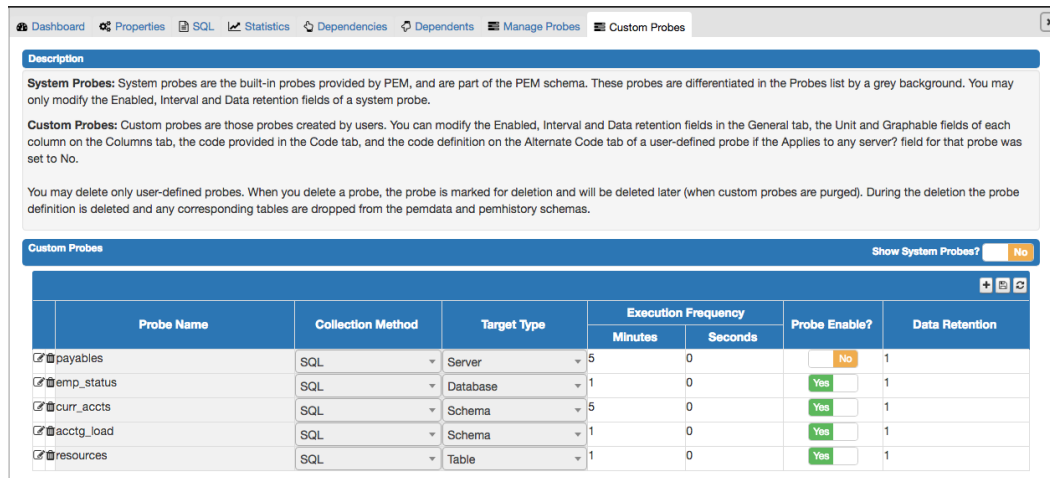


Figure 4.21 - The Custom Probes dialog.

Use the `Show System Probes?` switch to display or conceal the system probes on the `Custom Probes` tab.

You can use the `Custom Probes` tab to create a new probe or modify an existing probe. To create a new probe, click the `Add` icon in the upper-right corner of the tab; provide a name for the new probe in the `Probe Name` column. Then, select the `Edit` icon (located to the left of the probe name) to review or add the probe definition.

The screenshot shows the 'emp_status' window with the 'General' tab selected. The form contains the following fields and options:

- Probe Name:** A text input field.
- Collection Method:** A dropdown menu set to 'SQL'. Below it, a note states: 'Collection method field to specify the probe type. Use the drop-down to select:
 - SQL (the probe will gather information via a SQL statement)
 - WMI (the probe will gather information via a Windows Management Instrumentation extension)
 - Batch (the probe will use a command-script or shell-script to gather information). Please note that batch probes are platform-specific. If you specify a collection method of Batch, you must specify a platform type in the Platform field.
- Target Type:** A dropdown menu set to 'Server'. Below it, a note states: 'Target type drop-down to select the object type that the probe will monitor.'
- Minutes:** A numeric input field set to 5.
- Seconds:** A numeric input field set to 0.
- Probe Enable?:** A checkbox labeled 'Yes' which is checked. Below it, a note states: 'Enabled? to specify if the probe is enabled by default. Specify Yes to enable the probe by default, or No to specify that the probe is disabled by default.'
- Data Retention:** A numeric input field set to 1. Below it, a note states: 'Data retention field to specify the number of days that gathered information will be retained in the probe's history table.'
- Discard from history:** A checkbox labeled 'No' which is checked. Below it, a note states: 'Discard from history to specify if the server should create a history table for the probe. Select Yes to discard probe history, or No to retain the probe history in a table.'
- Platform:** A dropdown menu set to '*nix'. Below it, a note states: 'Platform drop-down to specify the type of platform that the probe will monitor. This field is enabled only when the Collection method is Batch'

Figure 4.22 – Defining a custom probe – the General tab.

Use the fields on the `General` tab (see Figure 4.22) to modify the definition of an existing probe or to specify the properties of a new probe:

- Use the `Probe Name` field to provide a name for a new probe.
- Use the `Collection method` field to specify the probe type. Use the drop-down listbox to select:
 - `SQL` - the probe will gather information via a SQL statement.
 - `WMI` - the probe will gather information via a Windows Management Instrumentation extension.
 - `Batch` - the probe will use a command-script or shell-script to gather information.

Before creating a batch probe on a Linux system, you must modify the `agent.cfg` file, setting the `allow_batch_probes` parameter equal to `true` and restart the PEM agent. The `agent.cfg` file is located in `/opt/PEM/agent/etc`.

On 64-bit Windows systems, agent settings are stored in the registry. Before creating a batch probe, modify the registry entry for the `AllowBatchProbes` registry entry and restart the PEM agent. PEM registry entries are located in

```
HKEY_LOCAL_MACHINE\\Software\\Wow6432Node\\EnterpriseDB\\
PEM\\agent.
```

Please note that batch probes are platform-specific. If you specify a collection method of `Batch`, you must specify a platform type in the `Platform` field.

- Use the `Target Type` drop-down list box to select the object type that the probe will monitor. `Target type` is disabled if `Collection method` is `WMI`.
- Use the `Minutes` and `Seconds` selectors to specify how often the probe will collect data.
- Use the `Probe Enable?` switch to specify if the probe is enabled by default. Specify `Yes` to enable the probe by default, or `No` to specify that the probe is disabled by default.

Please note: If data from a disabled probe is used in a chart, the chart will display an information icon in the upper-left corner that allows you to enable the probe by clicking the provided link.

- Use the `Data Retention` field to specify the number of days that gathered information will be retained in the probe's history table.
- Use the switch next to `Discard from history` to specify if the server should create a history table for the probe. Select `Yes` to discard probe history, or `No` to retain the probe history in a table.
- Use the `Platform` drop-down list box to specify the type of platform that the probe will monitor. This field is enabled only when the `Collection method` is `Batch`.

Name	Internal Name	Column Type	Data Type	Unit	Graphable	Is PIT	Calculate PIT
emp_status		Non Key	numeric		Yes	No	No

General

Name

Internal Name

Column Type Non Key
Column Type drop-down to specify if the column is a Key column (a primary key) or a Non key column. Non-key columns are generally metric items (values that can be graphed).

Data Type numeric
Data Type drop-down to specify the type of data that will be stored in the column.

Unit
Unit field to specify the unit of measure that applies to the metric stored in the column. This unit is displayed on the Y-Axis of a custom chart or a Capacity Manager chart. This is an optional field.

Graphable Yes
Graphable specify that the defined metric may be graphed, and that the probe should be accessible from the Capacity Manager or Manage Charts dialogs

Is PIT No
PIT specify that the metric is stored by point-in-time (by default). 'Point-in-time' metrics are those metrics that change (increase or decrease) at any given point of time. For example, database size is a point-in-time metric; at any given point-in-time, the size of the database is fluctuating. Metrics that are not point-in-time (also referred to as cumulative metrics) are metrics whose size always increases over time. For example, Blocks Read and Tuples Read are cumulative metrics; the value stays the same or increases.

Calculate PIT No
Calculate PIT specify that the server should calculate a point-in-time value for the metric data. Calculate PIT is disabled if Is PIT is Yes.

Figure 4.23 - The Columns tab of the Custom Probes dialog.

Use the **Columns** tab (see Figure 4.23) to define the columns in which the probe data will be stored. Navigate to the **Columns** tab, and click the **Add** button (in the upper-right corner) to define a new column. After providing a column name in the **Name** field, click the **Edit** button (to the left of the new column name) to provide information about the column:

- Provide a descriptive name for the column in the **Name** field.
- The **Internal Name** field is not enabled for user-defined probes.
- Use the **Column Type** drop-down listbox to specify if the column is a **Key** column (a primary key) or a **Non key** column. Non-key columns are generally metric items (values that can be graphed).
- Use the **Data Type** drop-down listbox to specify the type of data that will be stored in the column.
- Use the **Unit** field to specify the unit of measure that applies to the metric stored in the column. This unit is displayed on the Y-Axis of a custom chart or a Capacity Manager chart. This is an optional field.

- Use the `Graphable` switch to specify if the defined metric may be graphed, and that the probe should be accessible from the `Capacity Manager` or `Manage Charts` dialogs.
- Use the `Is PIT` switch to specify if the metric should be stored by point-in-time.

'Point-in-time' metrics are those metrics that change (increase or decrease) at any given point of time. For example, database size is a point-in-time metric; at any given point-in-time, the size of the database is fluctuating. Metrics that are not point-in-time (also referred to as cumulative metrics) are metrics whose size always increases over time. For example, `Blocks Read` and `Tuples Read` are cumulative metrics; the value stays the same or increases.

- Use the `Calculate PIT` switch to specify that the server should calculate a point-in-time value for the metric data. `Calculate PIT` is disabled if `Is PIT` is `Yes`.

PEM allows you to store point-in time-values of cumulative metrics as well. PEM subtracts the last collected value of a cumulative metric from the current value, and stores the difference as a point-in-time value.

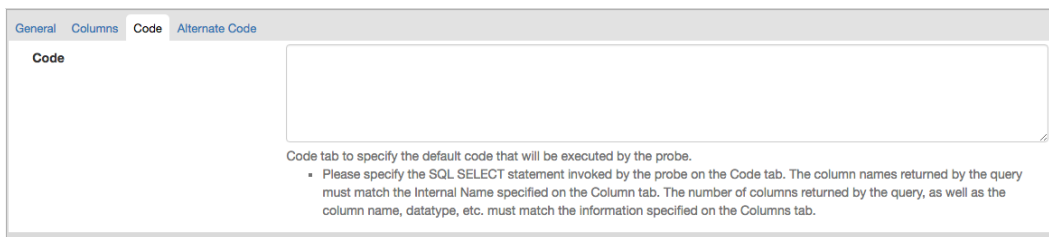


Figure 4.24 - The Code tab of the Custom Probes dialog.

Use the `Code` tab (see Figure 4.24) to specify the default code that will be executed by the probe:

- If the probe is a SQL probe, you must specify the SQL `SELECT` statement invoked by the probe on the `Code` tab. The column names returned by the query must match the `Internal Name` specified on the `Column` tab. The number of columns returned by the query, as well as the column name, data type, etc. must match the information specified on the `Columns` tab.
- If the probe is a batch probe, you must specify the shell or `.bat` script that will be invoked when the probe runs. The output of the script should be as follows:

The first line must contain the names of the columns provided on the `Columns` tab. Each column name should be separated by a tab (`\t`) character.

From the second line onwards, each line should contain the data for each column, separated by a tab character.

If a specified column is defined as key column, you should ensure that the script does not produce duplicated data for that column across lines of output. The number of columns specified in the `Columns` tab and their names, data type, etc. should match with the output of the script output.

- If the probe is a WMI probe, you must specify the WMI query as a `SELECT` WMI query. The column name referenced in the `SELECT` statement should be same as the name of the corresponding column specified on the `Column` tab. The column names returned by the query must match the `Internal Name` specified on the `Column` tab. The number of columns returned by the query, as well as the column name, data type, etc. must match the information specified on the `Columns` tab.

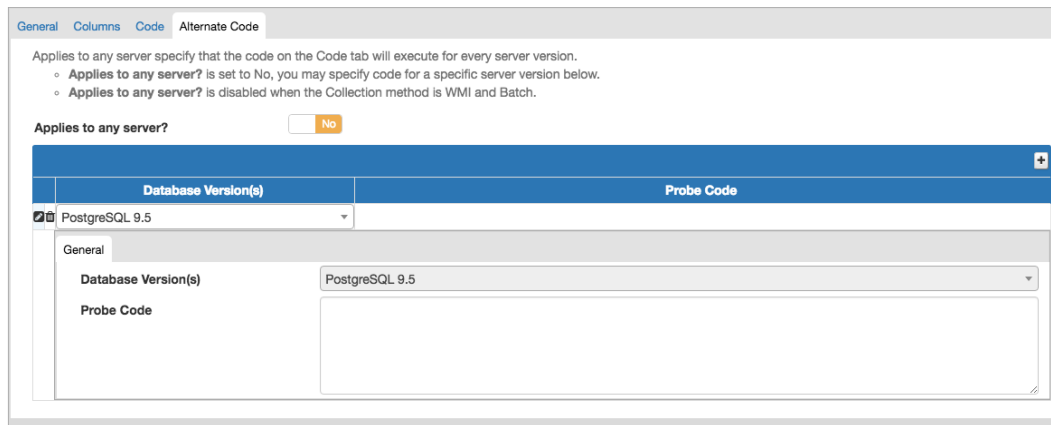


Figure 4.25 - The Alternate Code tab of the Custom Probes dialog.

Use the `Alternate Code` tab (see Figure 4.25) to provide code that will be invoked if the probe fires on a specific version of the server. To provide version-specific code, move the `Applies to any server version?` switch to `No`, and click the `Add` button. Then, select a version with the `Database Version(s)` drop-down listbox; click the `Edit` button (to the left of the version name) to provide the code that will execute when the probe fires.

If you select a database version, and leave the `Probe Code` column blank, PEM will invoke the code specified on the `Code` tab when the probe executes on a server that matches that version.

When you've finished defining the probe, click the `Save` icon (in the corner of the `Custom Probes` tab) to save the definition, and make the probe data available for use on custom charts and graphs.

4.4.2 Deleting a Probe

Use the `Delete` icon (located to the left of a `Probe Name`) to delete a user-defined probe. When you delete a probe, the probe is marked for deletion and will be deleted later (when custom probes are purged). During the deletion, the probe definition is deleted and any corresponding tables are dropped from the `pemdata` and `pemhistory` schemas.

Before deleting a probe, a PEM popup will ask you to confirm the deletion (see Figure 4.26).

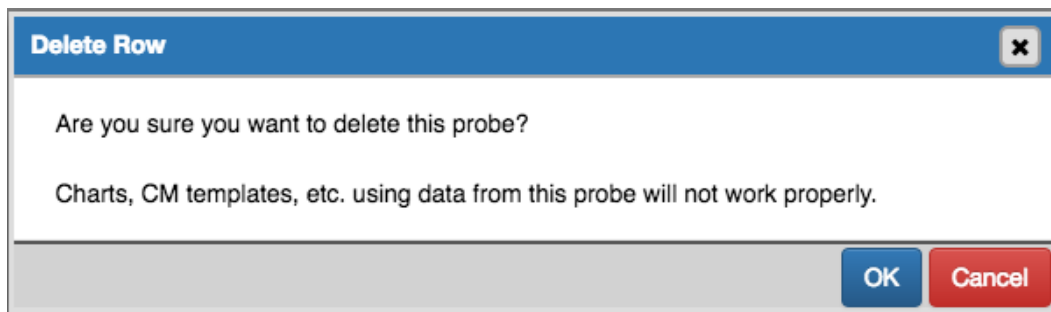


Figure 4.26 – Deleting a user-defined probe.

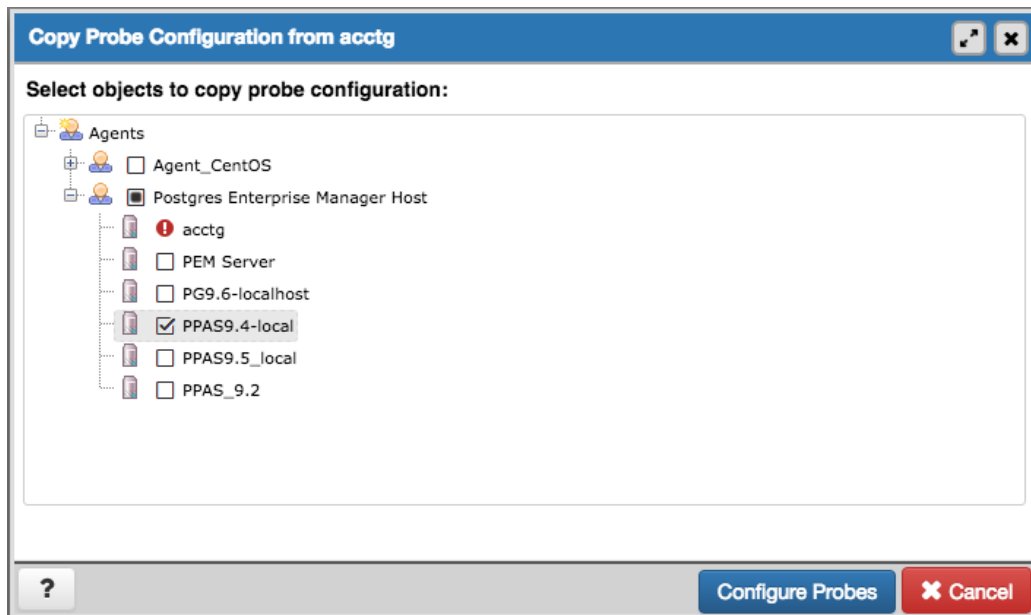
System probes are the built-in probes provided by PEM, and are part of the PEM schema. If you attempt to delete a system probe, the PEM client will display a notice, informing you that the probe cannot be deleted (see Figure 4.27).



4.27 – Attempting to delete a system probe.

4.4.3 Copying a Probe

You can use the `Copy Probe Configuration...` dialog to copy probe definitions from one monitored object to one or more monitored objects of the same type. To open the `Copy Probe Configuration...` dialog, highlight the object from which you are copying probes in the PEM client tree control, and select `Manage Probes` from the `Management` menu. When the `Manage Probes` tab opens, click on `Copy Probe` to open the `Copy Probe Configuration` dialog:



4.28 – The Copy Probe Configuration tree control.

The dialog will copy the probe definitions from the object through which the `Copy Probe Configuration` dialog was opened, to the location(s) selected on the tree control (see Figure 4.28).

Note that if you specify a parent node in the `Copy Probe Configuration` tree control, PEM will copy the probe configurations to each object (of the same type) that resides under that node in the tree control. For example, to copy the probe definitions from one schema to all schemas that reside within a database, select only the parent database of the target schemas. Please note that a red warning symbol is displayed to the left of the name of a listed target object if that object is the source of the probe that is being copied.

When you have selected the target object or objects, click the `Configure Probes` button to copy the probe definitions to the location selected on the dialog.

4.5 Alerting

PEM continually monitors registered servers and compares performance metrics against pre-defined and user-specified thresholds that constitute good or acceptable performance for each statistic. Any deviation from an acceptable threshold value triggers an *alert*. An *alert* is a system-defined or user-defined set of conditions that PEM compares to the system statistics. Alerts call your attention to conditions on registered servers that require your attention.

Reviewing alerts on the Global Overview

When your system statistics deviate from the boundaries specified for that statistic, the alert triggers, displaying a high (red), low (yellow), or medium (orange) severity warning in the left-most column of the `Alert Status` table on the `Global Overview` dashboard (see Figure 4.29).

Alert Status									
	Object Description	Alarm Type	Alert Name	Value	Database	Schema	Package	Object	Alerting Since
High	Postgres Plus Advanced Server 9.5	High	DB_bloat	821 MB					2018-01-30 13:54:58
High	EDB Postgres Advanced Server 10	High	Database size in server	489 MB					2018-01-29 11:22:12
High	EDB Postgres Advanced Server 9.6	High	Table size (sales)	180 MB					2018-01-30 13:31:48
Low	EDB Postgres Advanced Server 9.6	Low	Database size (resources)	215 MB					2018-01-30 14:26:08
Medium	Postgres Plus Advanced Server 9.5	Medium	DB_size_acctg	856 MB					2018-01-30 14:27:08
Medium	EDB Postgres Advanced Server 9.6	Medium	Last AutoVacuum	421475.14 hrs					2018-01-30 14:29:09

Figure 4.29 – The Alert Status table.

The PEM server includes a number of pre-defined alerts that are actively monitoring your servers. If the alert definition makes details available about the cause of the alert, you can click the down arrow to the right of the severity warning to access a dialog with detailed information about the condition that triggered the alert (see Figure 4.30).

EDB Postgres Advanced Server 10

High

Database size in server

489 MB

2018-01-29 11:22:12

Details

Parameters

Database name	Database size(MB)
edb	309
postgres	158
template1	11
template0	11

Figure 4.30 – Alert details.

PEM also provides an interface that allows you to create customized alerts. Each alert uses metrics defined on an *alert template*. An alert template defines how the server will evaluate the statistics for a resource or metric. The PEM server includes a number of pre-defined alert templates, or you can create custom alert templates. For more information about creating a custom alert template, see Section [4.5.2.3](#).

4.5.1 Using the Alerts Dashboard

Use the Dashboards menu (on the Monitoring tab) to access the Alerts Dashboard. The Alerts Dashboard (see Figure 4.31) displays a summary of the active alerts and the status of each alert:

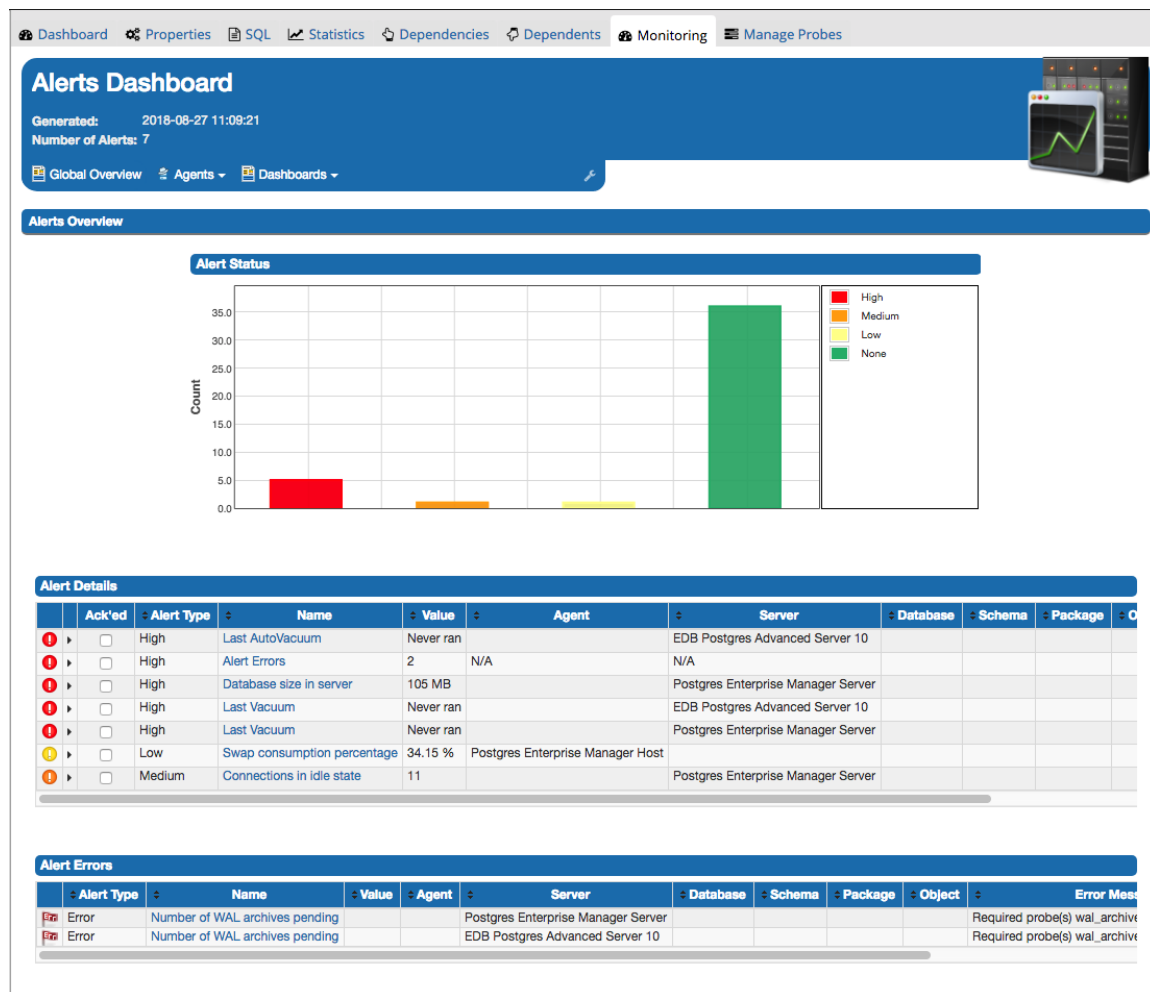


Figure 4.31 – The Alerts Dashboard.

The Alerts Dashboard header displays the date and time that the dashboard was last updated, and the number of current alerts.

The Alerts Overview section displays a graphic representation of the active alerts, as well as a count of the current high, low and medium alerts. The vertical bar on the left of the graph provides the count of the alerts displayed in each column. Hover over a bar to display the alert count for the selected alert severity in the upper-right hand corner of the graph.

The `Alert Details` table provides a list of the alerts that are currently triggered. The entries are prioritized from high-severity to lower-severity; each entry includes information that will allow you to identify the alert and recognize the condition that triggered the alert. Click the name of an alert to review detailed information about the alert definition.

The `Alert Errors` table displays configuration-related errors (eg. accidentally disabling a required probe, or improperly configuring an alert parameter). You can use the information provided in the `Error Message` column to identify and resolve the conflict that is causing the error.

Customizing the Alerts Dashboard

You can customize tables and charts that appear on the Alerts dashboard. To customize a table or chart, click the wrench icon located in the upper-right corner.

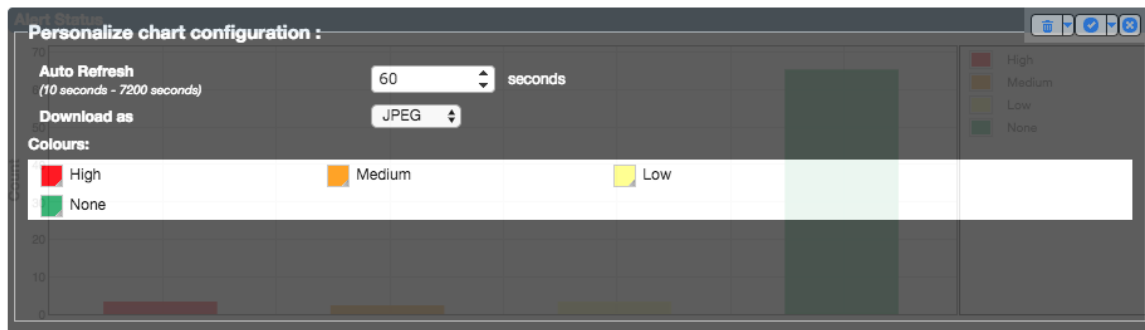


Figure 4.32 - Customizing a chart.

Use fields on the `Personalize chart configuration` dialog (Figure 4.32) to provide your display preferences:

- Use the `Auto Refresh` field to specify the number of seconds between updates of the data displayed in the table or chart.
- If applicable, use the `Download as` field to indicate if you would like a chart to be downloaded as a `JPEG` image or a `PNG` image.
- If applicable, use the `Colours` selectors to specify the display colors that will be used on a chart.
- If applicable, set the `Show Acknowledged Alerts` switch to `Yes` indicate that you would like the table to display alerts that you have acknowledged with a checkbox in the `Ack'ed` column. Set the field to `No` to indicate that the table should hide any acknowledged alerts. The switch acts as a toggle; acknowledged alerts are not purged from the table content until the time specified in the alert definition passes.

To save your customizations, click the `Save` icon (a check mark) in the upper-right corner; to delete any previous changes and revert to the default values, click the `Delete` icon. The `Save` and `Delete` drop-down menus allow you to specify if your preferences should be applied to `All Dashboards`, or to a selected server or database. Use the `Close` icon to close the `Personalize chart configuration` dialog without preserving your changes.

4.5.2 Using the Manage Alerts Tab

Use the PEM Client's `Manage Alerts` tab to define, copy, or manage alerts. To open the `Manage Alerts` tab (shown in Figure 4.33), select `Manage Alerts` from the `Management` menu.

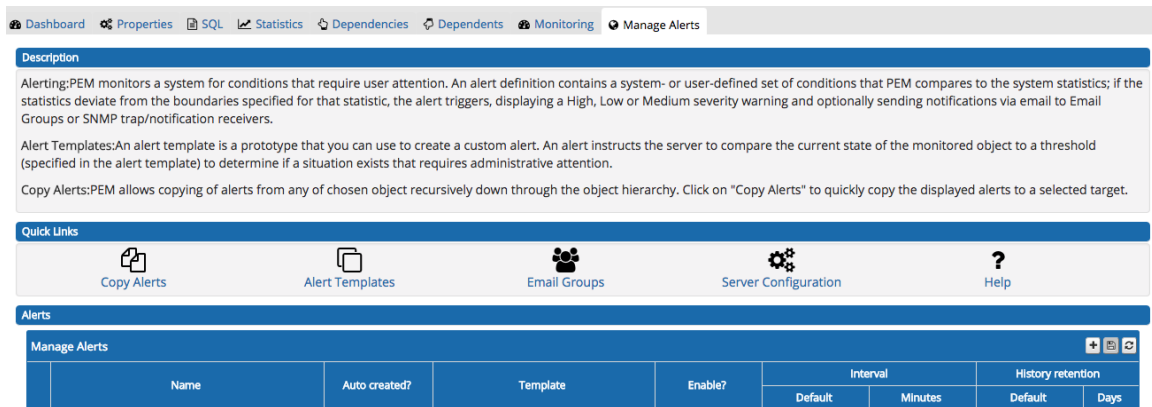


Figure 4.33 – The Manage Alerts tab.

Use the `Quick Links` toolbar to open dialogs and tabs that will assist you when managing alerts:

- Click `Copy Alerts` to open the `Copy Alert Configuration` dialog and copy an alert definition.
- Click `Alert Templates` to open the `Alert Template` tab, and modify or create an alert template.
- Click `Email Groups` to open the `Email Groups` tab, and modify or create an email group.
- Click `Server Configurations` to open the `Server configuration` dialog and review or modify server configuration settings.
- Click `Help` to open the PEM online help in a new tab of the PEM web interface.

Use the table in the `Alerts` section of the `Manage Alerts` tab to create new alerts or manage existing alerts.

4.5.2.1 Creating a Custom Alert Template

An alert template is a prototype that defines the properties of an alert. An alert instructs the server to compare the current state of the monitored object to a threshold (specified in the alert template) to determine if a situation exists that requires administrative attention.

You can use the **Alert Templates** tab (shown in Figure 4.34) to define a custom alert template or view the definitions of existing alert templates. To open the **Alert Templates** tab, select the **Manage Alerts...** menu option from the **Management** menu. When the **Manage Alerts** tab opens, select **Alert Templates** from the **Quick Links** toolbar.

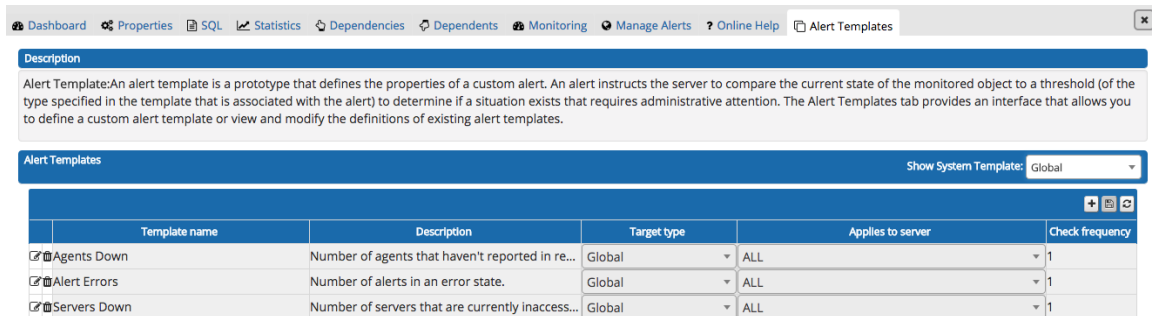


Figure 4.34 - The Alert Templates tab.

Use the **Show System Template** drop-down listbox to filter the alert templates that are displayed in the **Alert Templates** table. Use the listbox to select a level of the PEM hierarchy to view all of the templates for the selected level.

Defining a New Alert Template

To define a new alert template, use the **Show System Template** drop-down listbox to select **None**, and click the **Add** icon (+) located in the upper-right corner of the alert template table. The alert template editor opens as shown in Figure 4.35.

The screenshot shows the 'General' tab of an alert template configuration window. It includes fields for 'Template name' and 'Description' (both with red warning icons), a 'Target type' dropdown set to 'Server', an 'Applies to server' dropdown set to 'ALL', a 'History retention' spinner set to '30', and a 'Threshold unit' text field. Below these is an 'Auto create' section with an 'Auto create?' checkbox set to 'No' and a descriptive text block. At the bottom, there is an 'Operator' dropdown set to '>' and a 'Check frequency' spinner set to '1'. The 'Check frequency' field has a detailed instruction: 'Use the Check frequency field to specify the number of minutes between alert executions. This value specifies how often the server will invoke the SQL code specified on the SQL tab and compare the result to the threshold value specified in the alert definition.'

Figure 4.35 - The General tab.

Use fields on the `General` tab to specify general information about the template:

- Use the `Template name` field to specify a name for the new alert template.
- Use the `Description` field to provide a description of the alert template.
- Use the `Target type` drop-down listbox to select the type of object that will be the focus of the alert.
- Use the `Applies to server` drop-down listbox to specify the server type (EDB Postgres Advanced Server or PostgreSQL) to which the alert will be applied; you can specify a single server type, or `ALL`.
- Use the `History retention` field to specify the number of days that the result of the alert execution will be stored on the PEM server.
- Use the `Threshold unit` field to specify the unit type of the threshold value.
- Use fields in the `Auto create` box to indicate if PEM should use the template to generate an automatic alert. If enabled, PEM will automatically create an alert when a new server or agent (as specified by the `Target type` drop-down listbox) is added, and delete that alert when the target object is dropped.
 - Move the `Auto create?` slider to `Yes` to indicate that PEM should automatically create alerts based on the template. If you modify an

existing alert template, changing the `Auto create?` slider from `No` to `Yes`, PEM will create alerts on the existing agents and servers. Please note that if you change the slider from `Yes` to `No`, the default threshold values in existing alerts will be erased, and cannot be recovered.

- Use the `Operator` drop-down list box to select the operator that PEM will use when evaluating the current system values.

Select a greater-than sign (`>`) to indicate that the alert should be triggered when the system values are greater than the values entered in the `Threshold values` fields.

Select a less-than sign (`<`) to indicate that the alert should be triggered when the system values are less than the values entered in the `Threshold values` fields.

- Use the threshold fields to specify the values that PEM will compare to the system values to determine if an alert should be raised. Please note that you must specify values for all three thresholds (`Low`, `Medium`, and `High`):

Enter a value that will trigger a low-severity alert in the `Low` field.

Enter a value that will trigger a medium-severity alert in the `Medium` field.

Enter a value that will trigger a high-severity alert in the `High` field.

- Use the `Check frequency` field to specify the default number of minutes between alert executions. This value specifies how often the server will invoke the SQL code specified in the definition and compare the result to the threshold value specified in the template.

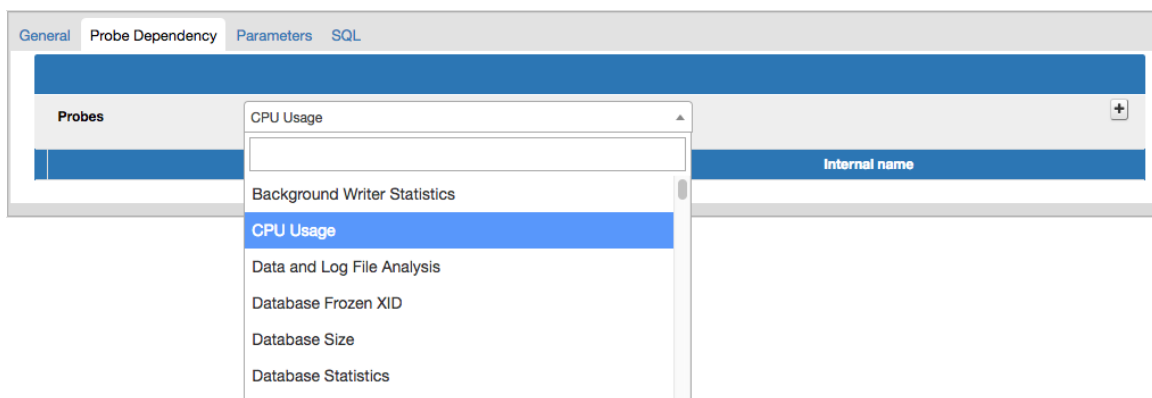


Figure 4.36 - The Probe Dependency tab of the Alert Templates dialog.

Use the fields on the `Probe Dependency` tab (see Figure 4.36) to specify the names of probes referred to in the SQL query specified on the `SQL` tab:

- Use the `Probes` drop-down listbox to select from a list of the available probes; highlight a probe name, and click the `Add` button to add the probe to the list of probes used by the alert template. To remove a probe from the selected probes list, highlight the probe name, and click the `Delete` icon.

Name	Data type	Unit
------	-----------	------

Figure 4.37 - The Parameters tab of the Alert Templates dialog.

- Use fields on the `Parameters` tab (see Figure 4.37) to define the parameters that will be used in the SQL code specified on the `SQL` tab. Click the `Add` icon (+) and:

Use the `Name` field to specify the parameter name.

Use the `Data type` drop-down listbox to specify the type of parameter.

Use the `Unit` field to specify the type of unit specified by the parameter.

- Use the `Code` field on the `SQL` tab (shown in Figure 4.38) to provide the text of the SQL query that the server will invoke when executing the alert. The SQL query will provide the result against which the threshold value is compared; if the alert result deviates from the specified threshold value, an alert will be raised.

Code

Use the `Code` field to provide the text of the SQL query that the server will invoke when executing the alert. The SQL query will provide the result against which the threshold value is compared; if the alert result deviates from the specified threshold value, an alert will be raised.

Detailed information SQL

Use the `Detailed information SQL` field to specify the SQL query that will provide the detailed information about that alert on Dashboard.

Figure 4.38 - The SQL tab of the Alert Templates dialog.

Within the query, parameters defined on the `Parameters` tab should be referenced sequentially by the variable `param_x`, where `x` indicates the position

of the parameter definition within the parameter list. For example, `param_1` refers to the first parameter in the parameter list, `param_2` refers to the second parameter in the parameter list, and so on.

The query can also include the following pre-defined variables:

Variable Description	Variable Name
agent identifier	'\${agent_id}'
server identifier	'\${server_id}'
database name	'\${database_name}'
schema name	'\${schema_name}'
Table	'\${object_name}'
index	'\${object_name}'
sequence	'\${object_name}'
function name	'\${object_name}'

- Use the `Detailed Information SQL` field to provide a SQL query that will be invoked if the alert is triggered. The result set of the query may be displayed as part of the detailed alert information on the `Alerts` dashboard or `Global Overview` dashboard.

Please Note: If the specified query is dependent on one or more probes from different levels within the PEM hierarchy (server, database, schema, etc.), and a probe becomes disabled, any resulting alerts will be displayed as follows:

- If the alert definition and the probe referenced by the query are from the same level within the PEM hierarchy, the server will display any alerts that reference the alert template on the `Alert Error` table of the `Global Alert Dashboard`.
- If the alert definition and the probe referenced by the query are from different levels of the PEM hierarchy, the server will display any triggered alerts that reference the alert template on the `Alert Details` table of the hierarchy on which the alert was defined.

Click the `Save` icon to save the alert template definition and add the template name to the `Alert Templates` list. After saving a custom alert template, you can use the `Alerting` dialog to define an alert based on the template.

Modifying or Deleting an Alert Template

To view the definition of an existing template (including PEM pre-defined alert templates), use the `Show System Template` drop-down list box to select the type of object monitored. When you select the object type, the `Alert Templates` table will display the currently defined alert templates that correspond with that object type.

Highlight a `TemplateName` in the list, and click the `Edit` icon (at the left end of the row) to review the template definition.

Use the tabs on the `Alert Templates` dialog to view detailed information about the alert template:

- General information is displayed on the `General` tab.
- The names of probes that provide data for the template are listed on the `Probe Dependency` tab.
- The names of any parameters referred to in the SQL code are listed on the `Parameters` tab.
- The SQL code that defines the behavior of the alert is displayed on the `SQL` tab.

To delete an alert template, highlight the template name in the alert templates table, and click the `Delete` icon. The alert history will persist for the length of time specified in the `History Retention` field in the template definition.

4.5.2.2 Creating a New Alert

The **Manage Alerts** tab displays a table of alerts that are defined on the object currently selected in the PEM client tree control. You can use the **Alerts** table to modify an existing alert, or to create a new alert (see Figure 4.39).

Description

Alerting: PEM monitors a system for conditions that require user attention. An alert definition contains a system- or user-defined set of conditions that PEM compares to the system statistics; if the statistics deviate from the boundaries specified for that statistic, the alert triggers, displaying a High, Low or Medium severity warning and optionally sending notifications via email to Email Groups or SNMP trap/notification receivers.

Alert Templates: An alert template is a prototype that you can use to create a custom alert. An alert instructs the server to compare the current state of the monitored object to a threshold (specified in the alert template) to determine if a situation exists that requires administrative attention.

Copy Alerts: PEM allows copying of alerts from any of chosen object recursively down through the object hierarchy. Click on "Copy Alerts" to quickly copy the displayed alerts to a selected target.

Quick Links

Copy Alerts Alert Templates Email Groups Server Configuration Help

Alerts

Manage Alerts								
	Name	Auto created?	Template	Enable?	Interval		History retention	
					Default	Minutes	Default	Days
<input checked="" type="checkbox"/>	Connections in idle-in-transaction sta...	No	Connections in idle-in-transactio...	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	DB usage	No	View Count	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	Last Autovacuum	No	Last AutoVacuum	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	Slow query	No	Long-running queries	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	Total Connections	No	Total connections	Yes	Yes	1	Yes	30

Figure 4.39 – The Manage Alerts tab.

To open the alert editor and create a new alert, click the **Add** icon (+) in the upper-right corner of the table. The editor opens as shown in Figure 4.40.

General Notification

Name

Description

Template

A template uses metrics to generate a value to which PEM compares user-specified alert boundaries. If the value returned by the template function evaluates to a value that is within the boundary of a user-defined alert, PEM raises an alert.

Enable? ☒ Yes ☐ No

Select Yes to enable the alert, and No to disable the alert.

Interval

Default ☒ Yes ☐ No Minutes

Use fields in the Interval box to specify how often the alert should confirm that alert conditions are satisfied.

History retention

Default ☒ Yes ☐ No Days

Use fields in the History retention box to specify the number of days that PEM will store data collected by the alert.

Threshold values

Operator Low Med High

Unit

The fields in the Threshold values box work together to define the triggering criteria for the alert.

Auto created? ☐ No ☒ Yes

Parameter Options	
Name	Value
server	EDB Postgres Advanced Server 9.6
database_name	edb

Figure 4.40 – The General tab of the alert editor.

Use the fields on the `General` tab to provide information about the alert:

- Enter the name of the alert in the `Name` field.
- Use the drop-down listbox in the `Template` field to select a template for the alert. An alert template is a function that uses one (or more) metrics or parameters to generate a value to which PEM compares user-specified alert boundaries. If the value returned by the template function evaluates to a value that is within the boundary of a user-defined alert (as specified by the `Operator` and `Threshold values` fields), PEM raises an alert, adds a notice to the `Alerts` overview display, and performs any actions specified on the template.
- Use the `Enable?` switch to specify if the alert is enabled (`Yes`) or disabled (`No`).

- Use the controls in the `Interval` box to specify how often the alert should confirm if the alert conditions are satisfied. Use the `Minutes` selector to specify an interval value. Use the `Default` switch to set or reset the `Minutes` value to the default (recommended) value for the selected template.
- Use controls in the `History retention` box to specify the number of days that PEM will store data collected by the alert. Use the `Days` selector to specify the number of days that the data will be stored. Use the `Default` switch to set or reset the `Days` value to the default value (30 days).
- Use controls in the `Threshold values` box to define the triggering criteria for the alert. When the value specified in the `Threshold Values` fields evaluates to greater-than or less-than the system value (as specified with the `Operator`), PEM will raise a `Low`, `Medium` or `High` level alert:
- Use the `Operator` drop-down list box to select the operator that PEM will use when evaluating the current system values:
 - Select a greater-than sign (`>`) to indicate that the alert should be triggered when the system values are greater than the values entered in the `Threshold values` fields.
 - Select a less-than sign (`<`) to indicate that the alert should be triggered when the system values are less than the values entered in the `Threshold values` fields.
- Use the threshold fields to specify the values that PEM will compare to the system values to determine if an alert should be raised. Please note that you must specify values for all three thresholds (`Low`, `Medium`, and `High`):
 - Enter a value that will trigger a low-severity alert in the `Low` field.
 - Enter a value that will trigger a medium-severity alert in the `Medium` field.
 - Enter a value that will trigger a high-severity alert in the `High` field.

The `Parameter Options` table contains a list of parameters that are required by the selected template; the table displays both pre-defined parameters, and parameters for which you must specify a value. Please note that you must specify a value for any parameter that displays a prompt in the `Value` column.

PEM can send a notification or execute a script if an alert is triggered, or if an alert is cleared. Use the **Notification** tab (see Figure 4.41) to specify how PEM will behave if an alert is raised.

General **Notification**

Email notification

All alerts ☐ No

Low alerts ☐ No

Medium alerts ☐ No

High alerts ☐ No

To configure notifications for an alert, use the fields in the Email notification box to specify the user or user group that will receive an email notification if the alert is triggered at the specified level. Use the drop-down listbox to select a pre-defined group that will be sent a notification if an alert of the selected level is triggered. Please note that you must configure the PEM Server to use an SMTP server to deliver email before PEM can send email notifications.

Trap notification

Send trap ☐ No **SNMP ver** **Low alert** ☐ No

Med alert ☐ No **High alert** ☐ No

Use the Trap notification options to configure trap notifications for this alert. Note that you must configure the PEM Server to send notifications to an SNMP trap/notification receiver before notifications can be sent.

Nagios notification

Submit passive service check result to Nagios ☐ No

Set "Submit passive service check result to Nagios" to "Yes" to instruct the PEM server to notify Nagios when the alert is triggered or cleared.

Script execution

Execute script ☐ No **Execute on alert cleared** ☐ No

Execute script on ☐ PEM Server ☒ Monitored Server

Code

Use the fields in the Script execution box to (optionally) define a script that will be executed if an alert is triggered, and to specify details about the script execution.

- Set the Execute script slider to Yes to instruct PEM to execute the provided script if an alert is triggered.
- Set the Execute on alert cleared slider to Yes to instruct PEM to execute the provided script when the situation that triggered the alert has been resolved.
- Use the selector to indicate if the script should execute on the PEM Server or the Monitored Server.
- Provide the script that PEM should execute in the Code field. You can provide a batch/shell script, or SQL code. Within the script, you can use placeholders for the following:
 - %AlertName% - this placeholder will be replaced with the name of the triggered alert.
 - %ObjectName% - this placeholder will be replaced with the name of the server or agent on which the alert was triggered.
 - %ThresholdValue% - this placeholder will be replaced with the threshold value reached by the metric when the alert triggered.
 - %CurrentValue% - this placeholder will be replaced with the current value of the metric that triggered the alert.
 - %CurrentState% - this placeholder will be replaced with the current state of the alert.
 - %OldState% - this placeholder will be replaced with the previous state of the alert.
 - %AlertRaisedTime% - this placeholder will be replaced with the time that the alert was raised, or the most recent time that the alert state was changed.

Figure 4.41 - The alert editor Notification tab.

Use the fields in the `Email notification` box to specify the email group that will receive an email notification if the alert is triggered at the specified level. Use the `Email Groups` tab to create an email group that contains the address of the user or users that will be notified when an alert is triggered. To access the `Email Groups` tab, click the `Email Groups` icon located in the `Quick Links` menu of the `Manage Alerts` tab.

- To instruct PEM to send an email when a specific alert level is reached, set the slider next to an alert level to `Yes`, and use the drop-down list box to select the pre-defined user or group that will be notified.

Please note that you must configure the PEM Server to use an SMTP server to deliver email before PEM can send email notifications.

Use the `Trap notification` options to configure trap notifications for this alert:

- Set the `Send trap` slider to `Yes` to send SNMP trap notifications when the state of this alert changes.
- Set the `SNMP Ver` slider to `v1` or `v2` to identify the SNMP version.
- Use the `Low alert`, `Med alert` and `High alert` sliders to select the level(s) of alert that will trigger the trap. For example, if you set the slider next to `High alert` to `Yes`, PEM will send a notification when an alert with a high severity level is triggered.

Please note that you must configure the PEM Server to send notifications to an SNMP trap/notification receiver before notifications can be sent.

Use the field in the `Nagios notification*` box to instruct the PEM server to notify Nagios network-alerting software when the alert is triggered or cleared. For detailed information about configuring and using Nagios with PEM, please see [Section 9](#).

- Set the `Submit passive service check result to Nagios` switch to `Yes` to instruct the PEM server to notify Nagios when the alert is triggered or cleared.

Use the fields in the `Script execution` box to (optionally) define a script that will be executed if an alert is triggered, and to specify details about the script execution.

- Set the `Execute script` slider to `Yes` to instruct PEM to execute the provided script if an alert is triggered.
- Set the `Execute on alert cleared` slider to `Yes` to instruct PEM to execute the provided script when the situation that triggered the alert has been resolved.

- Use the radio buttons next to `Execute script on` to indicate that the script should execute on the `PEM Server` or the `Monitored Server`.
- Provide the script that PEM should execute in the `Code` field. You can provide a batch/shell script, or SQL code. Within the script, you can use placeholders for the following:

`%AlertName%` - this placeholder will be replaced with the name of the triggered alert.

`%ObjectName%` - this placeholder will be replaced with the name of the server or agent on which the alert was triggered.

`%ThresholdValue%` - this placeholder will be replaced with the threshold value reached by the metric when the alert triggered.

`%CurrentValue%` - this placeholder will be replaced with the current value of the metric that triggered the alert.

`%CurrentState%` - this placeholder will be replaced with the current state of the alert.

`%OldState%` - this placeholder will be replaced with the previous state of the alert.

`%AlertRaisedTime%` - this placeholder will be replaced with the time that the alert was raised, or the most recent time that the alert state was changed.

When you have defined the alert attributes, click the edit icon to close the alert definition editor, and then the save icon (in the upper-right corner of the `Alerts` table). To discard your changes, click the refresh icon; a popup will ask you to confirm that you wish to discard the changes.

4.5.2.3 Modifying or Deleting an Alert

Use the `Alerts` table to manage an existing alert or create a new alert. Highlight an object in the PEM client tree control to view the alerts that monitor that object.

Alerts

Manage Alerts

	Name	Auto created?	Template	Enable?	Interval		History retention	
					Default	Minutes	Default	Days
<div><div><div></div><div></div><div></div></div><div><div></div><div></div><div></div></div></div>	Connections in idle-in-transaction sta...	<div>No</div>	Connections in idle-in-transactio... ▾	<div>Yes</div>	<div>Yes</div>	1	<div>Yes</div>	30
<div><div><div></div><div></div><div></div></div><div><div></div><div></div><div></div></div></div>	DB usage	<div>No</div>	View Count ▾	<div>Yes</div>	<div>Yes</div>	1	<div>Yes</div>	30
<div><div><div></div><div></div><div></div></div><div><div></div><div></div><div></div></div></div>	Last Autovacuum	<div>No</div>	Last AutoVacuum ▾	<div>Yes</div>	<div>Yes</div>	1	<div>Yes</div>	30
<div><div><div></div><div></div><div></div></div><div><div></div><div></div><div></div></div></div>	Slow query	<div>No</div>	Long-running queries ▾	<div>Yes</div>	<div>Yes</div>	1	<div>Yes</div>	30
<div><div><div></div><div></div><div></div></div><div><div></div><div></div><div></div></div></div>	Total Connections	<div>No</div>	Total connections ▾	<div>Yes</div>	<div>Yes</div>	1	<div>Yes</div>	30

Figure 4.42 - The Alerts table.

You can modify some properties of an alert in the `Alerts` table (see Figure 4.42):

- The `Alert name` column displays the name of the alert; to change the alert name, simply replace the name in the table, and click the save icon.
- The `Alert template` column displays the name of the alert template that specifies properties used by the alert. You can use the drop-down listbox to change the alert template associated with an alert.
- Use the `Alert enable?` switch to specify if an alert is enabled (`Yes`) or disabled (`No`).
- Use the `Interval` column to specify how often PEM should check to see if the alert conditions are satisfied. Set the `Default` switch to `No` and specify an alternate value (in `Minutes`), or return the `Default` switch to `Yes` to reset the value to its default setting. By default, PEM will check the status of each alert once every minute.
- Use the `History retention` field to specify the number of days that PEM will store data collected by the alert. Set the `Default*` switch to `No` and specify an alternate value (in `Days`), or return the `Default` switch to `Yes` to reset the value to its default setting. By default, PEM will recommend storing historical data for 30 days.

After modifying an alert, click the save icon (located in the upper-right corner of the table) to make your changes persistent.

Click the edit icon to the left of an alert name to open an editor that provides access to the complete alert definition to modify other alert attributes.

General **Notification**

Name: Connections in idle-in-transaction state

Description: Number of connections in database that are in idle-in-transaction state.

Template: Connections in idle-in-transaction state
 A template uses metrics to generate a value to which PEM compares user-specified alert boundaries. If the value returned by the template function evaluates to a value that is within the boundary of a user-defined alert, PEM raises an alert.

Enable? ☒ Yes
 Select Yes to enable the alert, and No to disable the alert.

Interval
 Default ☒ Yes Minutes

Use fields in the Interval box to specify how often the alert should confirm that alert conditions are satisfied.

History retention
 Default ☒ Yes Days

Use fields in the History retention box to specify the number of days that PEM will store data collected by the alert.

Threshold values
 Operator: > Low: Med: High:
 Unit: #

The fields in the Threshold values box work together to define the triggering criteria for the alert.

Auto created? ☐ No

Parameter Options	
Name	Value
server	EDB Postgres Advanced Server 9.6
database_name	edb

Figure 4.43 – The Alert details dialog.

Use fields on the `Alert details` dialog (see Figure 4.43) to modify the definition of the selected alert. When you've finished modifying the alert definition, click `Save` to preserve your changes, or `Cancel` to exit the dialog without saving any changes.

Deleting an Alert

To mark an alert for deletion, highlight the alert name in the `Alerts` table and click the delete icon to the left of the name; the alert will remain in the list, but in red strike-through font (see Figure 4.44).

Alerts								
Manage Alerts								
	Name	Auto created?	Template	Enable?	Interval		History retention	
					Default	Minutes	Default	Days
<input checked="" type="checkbox"/>	Connections in idle-in-transaction sta...	No	Connections in idle-in-transactio...	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	DB usage	No	View Count	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	Last Autovacuum	No	Last AutoVacuum	Yes	Yes	1	Yes	30
<input checked="" type="checkbox"/>	Slow query	No	Long-running queries	Yes	Yes	+	Yes	30
<input checked="" type="checkbox"/>	Total Connections	No	Total connections	Yes	Yes	1	Yes	30

Figure 4.44 – Deleting an alert.

The delete icon acts as a toggle; you can undo the deletion by clicking the delete icon a second time; when you click the `save` icon, the alert definition will be permanently deleted.

4.5.2.4 Copying an Alert

To speed up the deployment of alerts in the PEM system, you can copy alert definitions from one object to one or more target objects.

To copy alerts from an object, highlight the object in the PEM client tree control on the main PEM window, and select the `Copy Alerts...` option from the Management menu. When the Manage Alerts tab opens, click the `Copy Alerts` icon (located in the Quick Links toolbar) to open the Copy Alert Configuration dialog (see Figure 4.45).

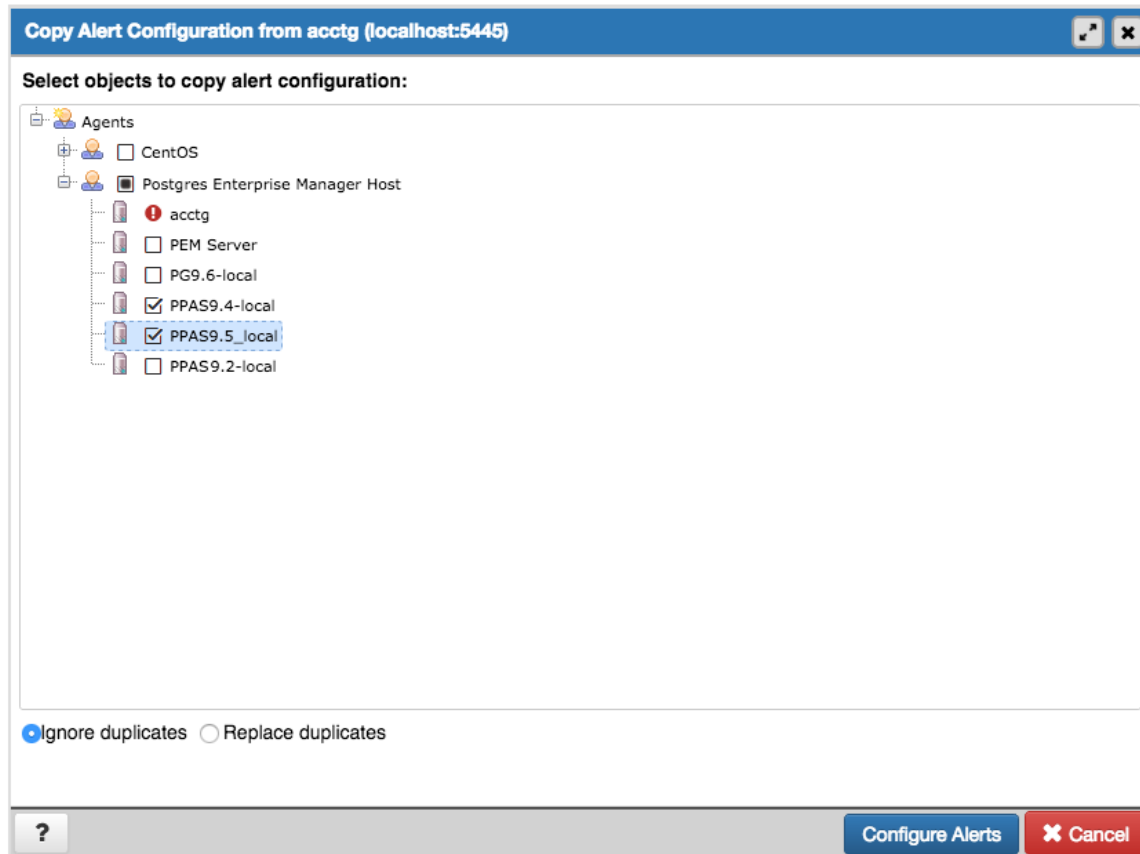


Figure 4.45 – The Copy Alert Configuration dialog.

The `Copy Alert Configuration` dialog copies all alerts from the object highlighted in the PEM client tree control to the object or objects selected on the dialog. Expand the tree control to select a node or nodes to specify the target object(s). The tree control displays a red warning indicator next to the source object.

To copy alerts to multiple objects at once, select a parent node of the target(s). For example, to copy the alerts from one table to all tables in a schema, you can simply select the checkbox next to the schema. PEM will only copy alerts to targets that are of the same type as the source object.

Check the `Ignore duplicates` radio button to prevent PEM from updating any existing alerts on the target objects with the same name as those being copied. Use the `Replace duplicates` option to replace existing alerts with alerts of the same name from the source object.

Click the `Configure Alerts` button to proceed to copy the alerts from the source object to all objects of the same type in, or under those objects selected on the `Copy Alert Configuration` dialog.

4.5.2.5 Audit Log Alerting

PEM provides alert templates that allow you to use the `Alerting` dialog to create an alert that will trigger when an `ERROR` or `WARNING` statement is written to a log file for a specific server or agent. To open the `Alerting` dialog, highlight the name of the server or agent in the PEM client `Object browser tree` control, and select `Alerting...` from the `Management` menu.

To create an alert that will notify you of `ERROR` or `WARNING` messages in the log file for a specific server, create an alert that uses one of the following alert templates:

Number of `ERRORS` in the logfile on server M in last X hours

Number of `WARNINGS` in the logfile on server M in last X hours

Number of `ERRORS` or `WARNINGS` in the logfile on server M in last X hours

To create an alert that will notify you of `ERROR` or `WARNING` messages for a specific agent, create an alert that uses one of the following alert templates:

Number of `ERRORS` in the logfile on agent M in last X hours

Number of `WARNINGS` in the logfile on agent M in last X hours

Number of `ERRORS` or `WARNINGS` in the logfile on agent M in last X hours

Please note that this functionality is supported only on Advanced Server.

4.5.2.6 Creating an Email Group

Postgres Enterprise Manager monitors your system for conditions that require user attention. You can use an email group to specify the email addresses of users that the server will notify if current values deviate from threshold values specified in an alert definition. An email group has the flexibility to notify multiple users, or target specific users during user-defined time periods.

Please note that you must configure the PEM Server to use an SMTP server to deliver email before PEM can send email notifications.

Use the **Email Groups** tab to configure groups of SMTP email recipients. To access the **Email Groups** tab, select **Manage Alerts...** from the PEM client's **Management** menu; when the **Manage Alerts** tab opens, select **Email Groups** from the **Quick Links** toolbar.

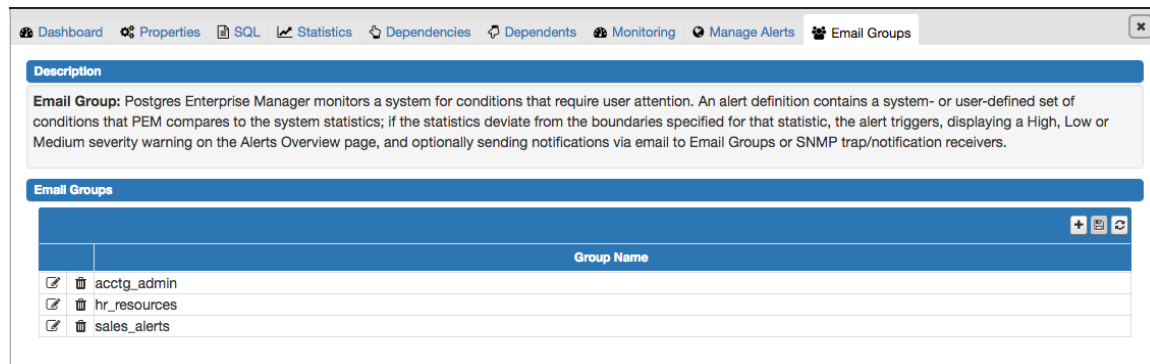


Figure 4.46 – The Email Groups tab.

The **Email Groups** tab displays a list of the currently defined email groups (see Figure 4.46). Highlight a group name and click the **Edit** icon (at the far left end of the row) to modify an existing group.

To define a new email group, click the Add icon (+) in the upper-right corner of the Email Groups table. The Email Group definition dialog opens, as shown in Figure 4.47.

Email Group

Group name

Email group options specify email notifications will be delivered to a specific group member (or members) during a selected time period.

- **To addresses:** Enter a comma-delimited list of recipient addresses in the To addresses field.
- **Reply to addresses:** Enter a comma-delimited list of recipient addresses in the Reply to addresses field.
- **Cc addresses:** Enter a comma-delimited list of addresses that will receive a copy of the email in the Cc addresses field.
- **Bcc addresses:** Enter a comma-delimited list of addresses that will receive a copy of the email (without the knowledge of other recipients) in the Bcc addresses field.
- **From address:** Enter the email address that messages to this group should be sent from in the From address field.
- **Subject prefix:** Enter the email subject prefix to this group in the Subject prefix field.
- **From time/To time(HH:MM:SS):** Use the From time and To time hour selectors to select a time range for a group member (or members). When a notification is sent, the server will evaluate the times specified within the group list and send the message to those members whose group entries include the current time. Provide the From time and To time values in the locale of the PEM client host, and the PEM server will translate the time into other time zones as required.

To addresses	From address	From time	To time
		00:00:00	23:59:59

Options

To addresses

Reply to addresses

Cc addresses

Bcc addresses

From address

Subject prefix

From time

To time

Please specify receiver email address

Figure 4.47 – Adding an email group.

Use the Email Group dialog to define an email group and its members:

- Provide a name for the email group in the Group Name field.

Each row within the email group definition will associate a unique set of email addresses with a specific time period. When an alert is triggered, the server will evaluate the times specified in each row and send the message to those group members whose definitions are associated with the time that the alert triggered.

Click the Add icon (+) in the group members table to open the Options tab, and add the member addresses that will receive notifications for the time period specified:

- Enter a comma-delimited list of recipient addresses in the `Reply to Addresses` field.
- Enter a comma-delimited list of addresses that will receive a copy of the email in the `CC Addresses` field.
- Enter a comma-delimited list of addresses that will receive a copy of the email (without the knowledge of other recipients) in the `Bcc Addresses` field.
- Enter the email address that messages to this group should be sent from in the `From Address` field.
- Use the `Subject prefix` field to provide a message that will be added to the start of each subject line when a notification is sent.
- Use the `From Time` and `To Time` time selectors to specify the time range for notifications to the group member(s) that are identified on this row. Provide the `From Time` and `To Time` values in the locale of the PEM client host, and the PEM server will translate the time into other time zones as required.

When you've identified the member or members that will receive an email during a specific time period, click the `Add` icon to add a row to the table, and specify another time period and the email addresses that will be notified during those hours. When you've finished defining the email group, click the `Save` icon.

To delete an email group, highlight the name of the group in the `Email Group` table and click the `Delete` icon (located to the left of the group name).




Figure 4.48 – Deleting an email group.

The group name will be displayed in the `Email Group` table in red; click the `Save` icon to make the change persistent and remove the group from the table (see Figure 4.48).

After creating the email group, you can use the `Manage Alerts` tab to set up the `Notification` details for an alert that will direct notifications to the group.

4.5.3 Using PEM with Nagios

The PEM server can send a passive alert result to Nagios network-alerting software when a user-defined alert is triggered. To instruct the PEM server to notify Nagios of a triggered alert, you must:

- Enable Nagios notification for each alert that will trigger a notification from the PEM server to Nagios. Please note that PEM alerting must be configured before you create the `host.cfg` file, the `services.cfg` file, or configure Nagios.
- Configure Nagios-related behaviors of the PEM server.
- Create the `host.cfg` and `services.cfg` configuration files.
- If necessary, modify the Nagios configuration file and restart the server.

After configuring the server to enable Nagios alerting, any triggered alerts will send a passive check result to the Nagios service. The syntax of a passive alert is:

```
[timestamp] PROCESS_SERVICE_CHECK_RESULT; host_name ;
service_name ; service_status ;
```

Where:

timestamp is the date and time that the alert was triggered.

host_name is the name of the server or agent.

service_name is the name of the alert.

service_status is the numeric service status value:

- 0 if the service status is OK
- 1 if the service status is WARNING
- 2 if the service status is CRITICAL
- 3 if the service status is UNKNOWN

The PEM server uses the following rules to evaluate the service status:

- If the PEM alert level is `CLEARED`, the warning message will read `OK`.
- If the PEM alert level is `LOW`, the warning message will read `WARNING`.

- If the `is_nagios_medium_alert_as_critical` flag (specified in the PEM server configuration dialog) is set to `FALSE` and the alert level `MEDIUM`, the warning message will read `WARNING`.
- If the `is_nagios_medium_alert_as_critical` flag (specified in the PEM server configuration dialog) is set to `TRUE` and the alert level is `MEDIUM`, the warning message will read `CRITICAL`.
- If the PEM alert level is `HIGH`, the warning message will read `CRITICAL`.

4.5.3.1 Enabling Nagios Notification for an Alert

The PEM server maintains a unique set of notification properties for each enabled alert. Use the `Notification` tab of the `Manage Alerts` tab to specify that (when triggered), a given alert will send an alert notice to Nagios.

To modify the notification properties of an alert, right-click on the name of the object monitored by the alert, and select `Manage Alerts...` from the `Management` menu. When the `Manage Alerts` tab opens, locate the alert, and then click the edit button to the left of the alert name in the `Alerts` list. When the edit pane opens, select the `Notification` tab (shown in Figure 4.49).

To configure notifications for an alert, use the fields in the Email Notification box to specify the user or user group that will receive an email notification if the alert is triggered at the specified level. Use the drop-down listbox to select a pre-defined group that will be sent a notification if an alert of the selected level is triggered. Please note that you must configure the PEM Server to use an SMTP server to deliver email before PEM can send email notifications.

Use the Trap notification options to configure trap notifications for this alert. Note that you must configure the PEM Server to send notifications to an SNMP trap/notification receiver before notifications can be sent.

Set "Submit passive service check result to Nagios" to "Yes" to instruct the PEM server to notify Nagios when the alert is triggered or cleared.

Use the fields in the Script execution to (optionally) define a script that will be executed if an alert is triggered, and to specify details about the script execution.

- Execute Script to instruct PEM to execute the provided script if an alert is triggered.
- Check the box next to Execute on alert cleared to instruct PEM to execute the provided script when the situation that triggered the alert has been resolved.
- Use the radio buttons to indicate that the script should execute on the PEM Server or the Monitored Server.
- Provide the script that PEM should execute in the Code field. You can provide a batch/shell script, or SQL code. Within the script, you can use placeholders for the following:
 - %AlertName% - this placeholder will be replaced with the name of the triggered alert.
 - %ObjectName% - this placeholder will be replaced with the name of the server or agent on which the alert was triggered.
 - %ThresholdValue% - this placeholder will be replaced with the threshold value reached by the metric when the alert triggered.
 - %CurrentValue% - this placeholder will be replaced with the current value of the metric that triggered the alert.
 - %CurrentState% - this placeholder will be replaced with the current state of the alert.
 - %OldState% - this placeholder will be replaced with the previous state of the alert.
 - %AlertRaisedTime% - this placeholder will be replaced with the time that the alert was raised, or the most recent time that the alert state was changed.

Figure 4.49 – The Notification tab.

To enable Nagios notification, move the slider next to `Submit passive service check result to Nagios` to `Yes`; before exiting the `Manage Alerts` tab, click the save icon to preserve your changes.

4.5.3.2 Configuring Nagios-related behavior of the PEM Server

You can use the `Server Configuration` dialog to provide information about your Nagios configuration to the PEM server. To open `Server Configuration` dialog (see Figure 4.50), select `Server Configuration...` from the PEM client's `Management` menu.

The screenshot shows the 'Server Configuration' dialog box. It has a search bar at the top right labeled 'Search by parameter name'. Below the search bar is a table of configuration parameters. The parameters are listed in two columns, with their values and units in the third column. Some parameters have toggle buttons for 'True' or 'False'.

Parameter Name	Value	Unit
nagios_cmd_file_name	/usr/local/nagios/var/rw/nagios.cmd	
nagios_enabled	<input checked="" type="checkbox"/> True	t/f
nagios_medium_alert_as_critical	<input type="checkbox"/> False	t/f
nagios_spool_retention_time	7	days
package_catalog_xml	http://sbp.enterprisedb.com/applications.xml	
package_download_chunk_size	1,048,576	Bytes
probe_log_retention_time	30	days
proxy_server	127.0.0.1	
proxy_server_authentication	<input type="checkbox"/> False	t/f
proxy_server_enabled	<input type="checkbox"/> False	t/f
proxy_server_password		
proxy_server_port	80	
proxy_server_username		
reminder_notification_interval	24	hours
server_log_retention_time	30	days
show_data_tab_on_graph	<input checked="" type="checkbox"/> True	
smtp_authentication	<input type="checkbox"/> False	t/f
smtp_enabled	<input checked="" type="checkbox"/> True	t/f
smtp_encryption	<input type="checkbox"/> False	t/f
smtp_password		
smtp_port	25	
smtp_server	127.0.0.1	
smtp_spool_retention_time	7	days
smtp_username		

At the bottom right of the dialog is a 'Close' button.

Figure 4.50 – Specify Nagios properties in the `Server Configuration` dialog.

Four server configuration parameters specify information about your Nagios installation and PEM server behavior related to Nagios:

- Use the `nagios_cmd_file_name` parameter to specify the location of the Nagios pipeline file that will receive passive check alerts from PEM. The default value of this parameter is `/usr/local/nagios/var/rw/nagios.cmd`. If

your `nagios.cmd` file resides in an alternate location, specify the file location in the `Value` field.

- Move the slider in the `nagios_enabled` parameter to `Yes` to instruct the PEM server to send passive check alerts to Nagios.
- Use the `nagios_medium_alert_as_critical` slider to specify the warning severity that the PEM server will pass to Nagios if a medium alert is triggered:

If the `is_nagios_medium_alert_as_critical` flag is set to `FALSE` and the alert level is `MEDIUM`, the warning message will read `WARNING`.

If the `is_nagios_medium_alert_as_critical` flag is set to `TRUE` and the alert level is `MEDIUM`, the warning message will read `CRITICAL`.

- Use the `nagios_spool_retention_time` parameter to specify the number of days of notification history that will be stored on the PEM server. The default value is 7 days.

After modifying parameter values, click the save icon (in the upper-right corner of the `Server Configuration` dialog) to preserve your changes.

4.5.3.3 Creating the `hosts.cfg` and `services.cfg` File

The `templates.cfg` file (by default, located in `/usr/local/nagios/etc/objects`) specifies the properties of a generic-host and generic-service. The properties specify the parameters used in the `hosts.cfg` and `services.cfg` files.

In most cases (when PEM is installed in a default configuration), you will not be required to modify the `templates.cfg` file before creating the `hosts.cfg` and `services.cfg` files. If necessary, you can modify the `templates.cfg` file to specify alternate values for parameters or to create new templates.

Before modifying the Nagios configuration file, use the following command to create a `hosts.cfg` file that contains information about the PEM hosts that reside on the local system:

```
psql -U postgres -p 5433 -d pem -A -t -c "select
pem.create_nagios_host_config('generic-host') " >
/usr/local/nagios/etc/objects/hosts.cfg
```

Then, use the following command to create a `services.cfg` file that contains information about the PEM services that reside on the local system:

```
psql -U postgres -p 5433 -d pem -A -t -c "select
pem.create_nagios_service_config('generic-service') " >
/usr/local/nagios/etc/objects/services.cfg
```

If you wish to use a custom `template.cfg` file entry, specify the entry name in place of `generic-host` or `generic-service` in the above commands.

4.5.3.4 Modifying the Nagios Configuration File

After creating the `host.cfg` and `services.cfg` files, you must specify their location in the Nagios configuration file (by default, `/usr/local/nagios/etc/nagios.cfg`). Modify the configuration file, adding entries that specify the location of the files:

```
cfg_file=/usr/local/etc/objects/hosts.cfg

cfg_file=/usr/local/etc/objects/services.cfg
```

You can use the following command to confirm that Nagios is properly configured:

```
/usr/local/nagios/bin/nagios -v
/usr/local/nagios/etc/nagios.cfg
```

After confirming that Nagios is configured correctly, restart the Nagios service:

```
/usr/local/nagios/bin/nagios -d
/usr/local/nagios/etc/nagios.cfg
```

5 Capacity Manager

PEM's Capacity Manager analyzes collected statistics (metrics) to generate a graph or table that displays the historical usage statistics of an object, and can project the anticipated usage statistics for an object. You can configure Capacity Manager to collect and analyze metrics for a specific host, server, database, or database object.

You can tailor the content of the Capacity Manager report by choosing a specific metric (or metrics) to include in the report, the time range over which the metrics were gathered, and a high or low threshold for the metrics analyzed. You can also specify a start and end date for the Capacity Manager report. If the end date of the report specifies a time in the future, Capacity Manager will analyze the historical usage of the selected object to extrapolate the projected object usage in the future.

To open Capacity Manager, select the `Capacity Manager . . .` option from the PEM client Management menu; the Capacity Manager wizard opens, displaying a tree control on the Metrics tab (see Figure 5.1).

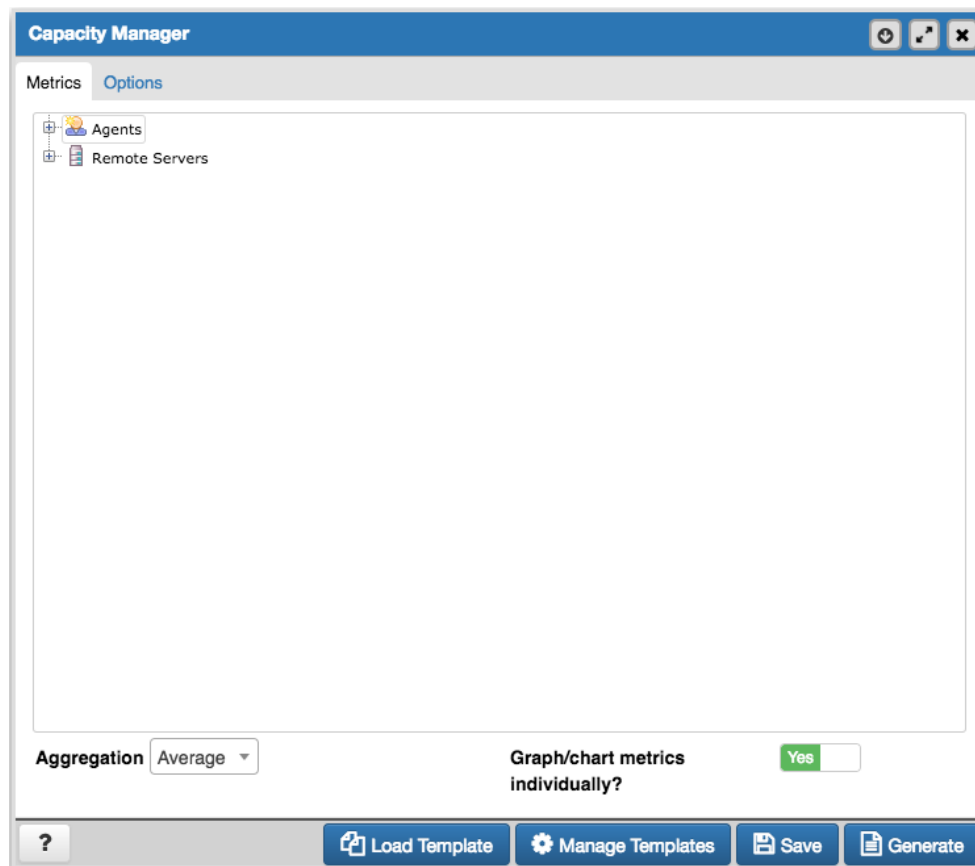


Figure 5.1 - The Capacity Manager dialog

Expand the tree control on the `Metrics` tab to review the metrics for the node that you wish to analyze (see Figure 5.2). Check the box to the left of the name of the metric to include the metric in your report.

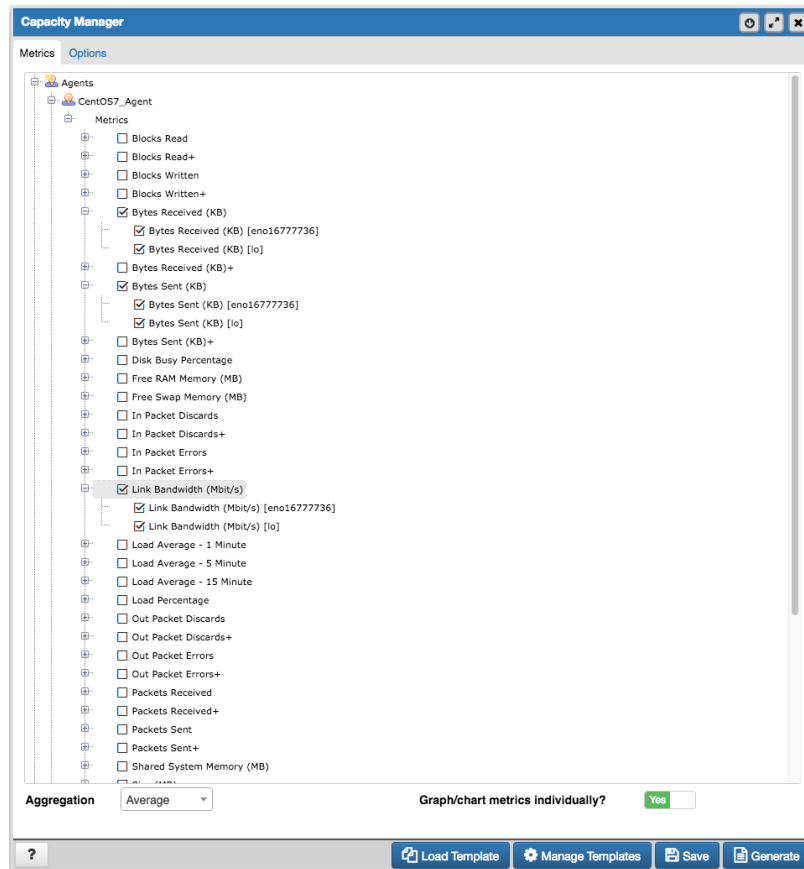


Figure 5.2 - The Metrics tree control.

Capacity Manager will use the aggregation method specified with the `Aggregation` drop-down listbox (located at the bottom of the `Metrics` tab). The aggregation method instructs Capacity Manager how to evaluate and plot the metric values. Select from:

- **Average:** Use the average of the values recorded during the time period.
- **Maximum:** Use the maximum value recorded during the time period.
- **Minimum:** Use the minimum value recorded during the time period.
- **First:** Use the first value recorded during the time period.

To remove a metric from the Capacity Manager report, uncheck the box to the left of the name of a metric.

Move the slider next to `Graph/chart metrics individually?` to `Yes` to instruct Capacity Manager to produce a separate report for each metric selected on the `Metrics` tab. If the option is set to `No`, all selected metrics will be merged into a single graph or table.

Click the `Generate` button to display the report onscreen (accepting the default configuration options), or use the `Options` tab (see Figure 5.3) to customize sampling boundaries, report type and report destination. Please note that the times displayed on the `Options` tab are the time zone in which the PEM client resides.

The screenshot shows the 'Capacity Manager' application window with the 'Options' tab selected. The 'Time Period' section on the left includes a 'Period' dropdown set to 'Start time and end time', 'Start time' and 'End time' date/time pickers, 'Historical day(s)' and 'Extrapolate day(s)' spinners both set to 7, a 'Threshold' dropdown set to 'Falls below', and a value field set to 0. The 'Report' section on the right includes 'Include on report' radio buttons for 'Graph' (selected), 'Table of data', and 'Graph and table of data'. The 'Report destination' section has radio buttons for 'New tab' (selected), 'Previous tab', and 'Download the report as a file:', with a text field below it containing 'cm_report.html'. At the bottom are buttons for '?', 'Load Template', 'Manage Templates', 'Save', and 'Generate'.

Figure 5.3 – Specify the time period, type, and destination of the report.

Use the fields within the `Time Period` box to define the boundaries of the Capacity Manager report:

- Use the `Period` drop-down list box to select the type of time period you wish to use for the report. You can select:

Start time and end time	Specify a start date and an end date/time for the report.
Start time and threshold	Specify a start date and time, and a threshold to determine the end time and date for the report.
Historical days and extrapolated days	Specify a start date for the report that is a number of days in the past, and an end date that is a number of days in the future. This option is useful for report templates that do not specify fixed dates.

Historical days and threshold	Specify a start date that is a number of days in the past, and end it when a threshold value is reached.
-------------------------------	--

After specifying the type of time period for the report, select from other options in the Time Period box to define the time period for the report:

- Use the date and time selectors next to the `Start time` field to specify the starting date and time of the sampling period, or select the number of `Historical day(s)` of data to include in the report. The date and time specified in the `Start time` field must not be later than the current date/time.

By default, Capacity Manager will select a start time that is one week prior to the current date and time.

- The end boundary for the report can be a time, a number of days in the future, or the point at which a selected metric reaches a user-specified threshold value. Use the date and time selectors next to the `End time` field to specify an end boundary for the report, or select the number of `Extrapolated day(s)` of data to include in the report. The time specified in the `End time` field must be later than the time specified in the `Start time` field.

Note that if you select an end date and time in the future, Capacity Manager will use historical usage information to extrapolate anticipated future usage. Since the projected usage is based on the sampling of historical data, the accuracy of the future usage trend will improve with a longer sampling period.

To specify a threshold value, use the drop-down listbox in the `Threshold` field to select a metric, an operator (`Exceeds*` or `Falls below`), and enter a target value for the metric. If you choose to define the end of the report using a threshold, the Capacity Manager report will terminate when the value for the selected metric exceeds or falls below the specified value.

The `cm_max_end_date_in_years` configuration parameter defines a default time value for the end boundary of a Capacity Manager report. If you specify a threshold value as the end boundary of a report, and the anticipated usage of the boundary is not met before the maximum time has passed, the report will terminate at the time specified by the `cm_max_date_in_years` parameter. By default, `cm_max_end_date_in_years` is 5; you can use the `Server Configuration` dialog to modify the value of `cm_max_end_date_in_years`. For more information about modifying the server configuration, see Section **Error! Reference source not found.**

The fields in the `Report` box specify the report type and destination. Use the `Include on report` radio buttons to specify the type of report produced by Capacity Manager:

- Select `Graph` to instruct Capacity Manager to display the report in the form of a line graph in the PEM client window.
- Select `Table of data` to instruct Capacity Manager to display a table containing the report data in the PEM client window.
- Select `Graph and table of data` to instruct Capacity Manager to display both a line graph and a data table in the PEM client window.

Use the `Report destination` radio buttons to instruct Capacity Manager where to display or save the report:

- Select `New tab` to instruct Capacity Manager to display the report on a new tab in the PEM client. You must select `New tab` to display the first generation of a Capacity Manager report; for subsequent reports, you may select `Previous tab`.
- Select `Previous tab` to instruct Capacity Manager to re-use a previously opened tab when displaying the report.
- Select `Download the report as a file` and specify a file name to instruct Capacity Manager to write the report to the specified file.

When you have specified the report boundaries and selected the type and destination of the Capacity Manager report, click the `Generate` button to create the report (see Figure 5.4).




Figure 5.4 - The Capacity Manager report

Reports saved to file are stored in HTML format. You can review a Capacity Manager report with any web browser that supports Scalable Vector Graphics (SVG). Browsers that do not support SVG will be unable to display a Capacity Manager graph and may include unwanted characters.

5.1 Capacity Manager Templates

After defining a report, you can save the definition as a template for future reports. Capacity Manager report templates may be accessed by all PEM users. To save a report definition as a template:

1. Use the `Metrics` and `Options` tabs to define your report.
2. Click the `Save` button to open the `Save Template` dialog (shown in Figure 5.5).

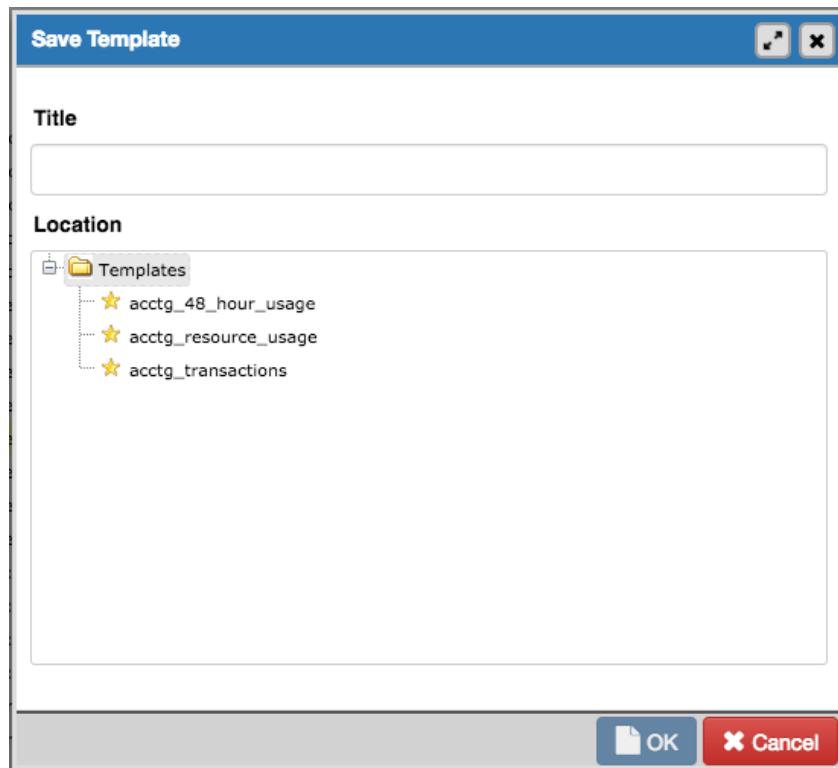


Figure 5.5 – Saving a Capacity Manager Template

3. Provide a report name in the `Title` field, select a location to store the template in the tree control.
4. Click `OK`.

When creating a report, you can use the `Load Template` button to browse and open an existing template. Once opened, the report definition may be modified if required, and optionally saved again, either as a new template, or overwriting the original template.

Use the `Manage Templates` button open a dialog that allows you to rename or remove unwanted templates.

6 Audit Manager

You can use the PEM Audit Manager to simplify audit log configuration for Advanced Server instances. With the Audit Manager, you can configure logging attributes such as:

- How often log files are to be collected by PEM
- The type of database activities that are included in the log files
- How often (and when) log files are to be rotated

Audit logs may include the following activities:

- All connections made to the database instance
- Failed connection attempts
- Disconnections from the database instance
- All queries (`SELECT` statements)
- All DML statements (`INSERT`, `UPDATE`, `DELETE`)
- All DDL statements (e.g., `CREATE`, `DROP`, `ALTER`)

Once the audit logs are stored on the PEM server, you can use the Audit Log dashboard to review the information in an easy-to-read form. The Audit Log dashboard allows you to filter the log file by timestamp range (when an activity occurred), the database on which the activity occurred, the user performing the activity, or the type of command being invoked.

6.1 Setting the Advanced Server Instance Service ID

To configure logging for an Advanced Server instance, the server must be a PEM-managed server with a bound agent, and the server registration must include the name of a service script. When registering a new server, include the service name in the `Service ID` field on the `Advanced` tab of the `New Server` dialog.

Before adding a service name to an existing (registered and connected) server, you must disconnect the server. Right click on the server name, and select `Disconnect server` from the context menu. Then, right click on the server name and select `Properties` from the context menu. Select the `Advanced` tab, and add a service name to the `Service ID` field (as shown in Figure 6.1).

The screenshot shows the 'Create - Server' dialog box with the 'Advanced' tab selected. The 'Service ID' field contains the text 'edb-as-9.6'. Other fields include 'Host Address', 'DB Restriction', 'EFM Cluster Name', and 'EFM Installation Path'. The bottom of the dialog has buttons for 'Save', 'Cancel', and 'Reset'.

Figure 6.1 - The Service ID of the Advanced Server instance.

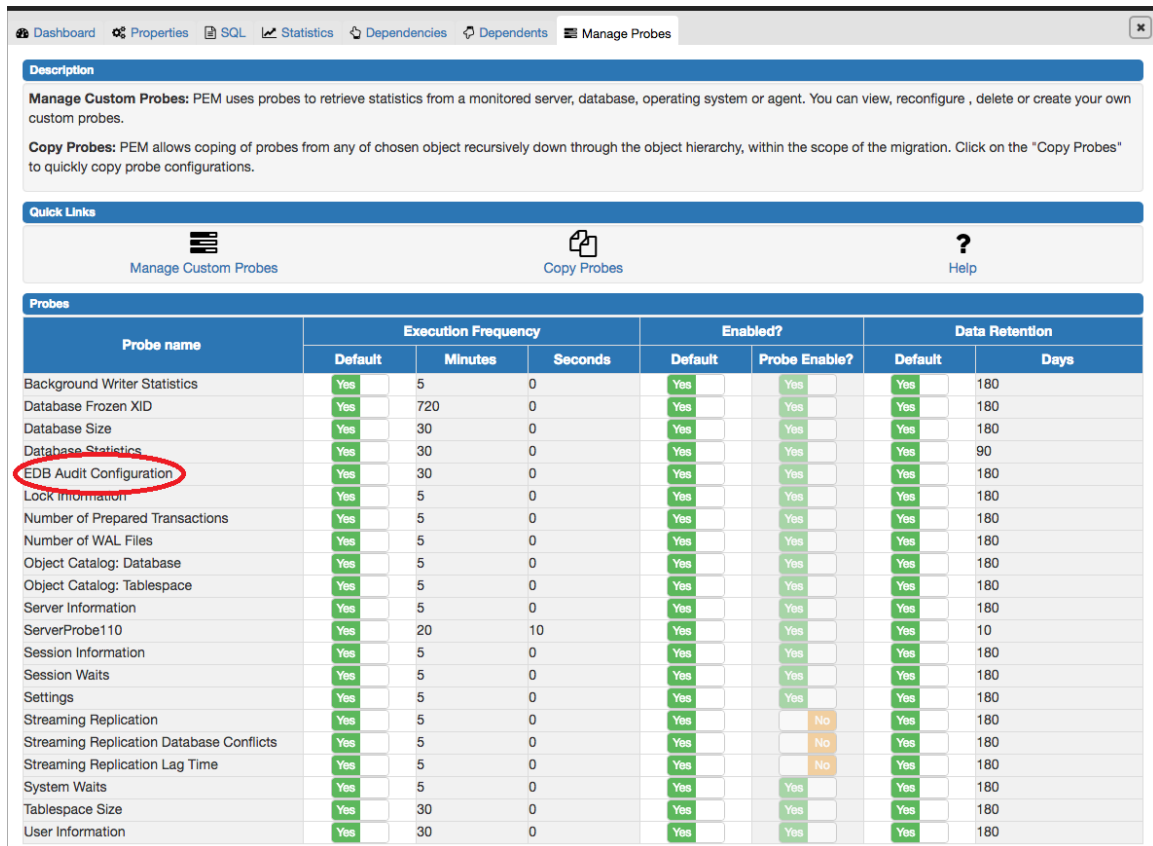
The `Service ID` field allows the PEM server to stop and start the service.

- The name of the Advanced Server 10 service script is `edb-as-10`.
- The name of the Advanced Server 9.6 service script is `edb-as-9.6`.
- The name of the Advanced Server 9.5 (or prior) service script is `ppas-9.x`, where *x* specifies the version.
- The name of the PostgreSQL 9.6 service script is `postgresql-9.6`.

6.2 Setting the EDB Audit Configuration Probe

Before configuring audit logging of Advanced Server servers, you must ensure that the EDB Audit Configuration probe is enabled. To open the Manage Probes tab and check the status of the probe, right click on the name of a registered Advanced Server server in the tree control, and select Manage Probes... from the Management menu.

Ensure that the Enabled column in the Probe Configuration dialog is set to Yes for the EDB Audit Configuration probe (see Figure 6.2).



Probe name	Execution Frequency			Enabled?		Data Retention	
	Default	Minutes	Seconds	Default	Probe Enable?	Default	Days
Background Writer Statistics	Yes	5	0	Yes	Yes	Yes	180
Database Frozen XID	Yes	720	0	Yes	Yes	Yes	180
Database Size	Yes	30	0	Yes	Yes	Yes	180
Database Statistics	Yes	30	0	Yes	Yes	Yes	90
EDB Audit Configuration	Yes	30	0	Yes	Yes	Yes	180
Lock Information	Yes	5	0	Yes	Yes	Yes	180
Number of Prepared Transactions	Yes	5	0	Yes	Yes	Yes	180
Number of WAL Files	Yes	5	0	Yes	Yes	Yes	180
Object Catalog: Database	Yes	5	0	Yes	Yes	Yes	180
Object Catalog: Tablespace	Yes	5	0	Yes	Yes	Yes	180
Server Information	Yes	5	0	Yes	Yes	Yes	180
ServerProbe110	Yes	20	10	Yes	Yes	Yes	10
Session Information	Yes	5	0	Yes	Yes	Yes	180
Session Waits	Yes	5	0	Yes	Yes	Yes	180
Settings	Yes	5	0	Yes	Yes	Yes	180
Streaming Replication	Yes	5	0	Yes	No	Yes	180
Streaming Replication Database Conflicts	Yes	5	0	Yes	No	Yes	180
Streaming Replication Lag Time	Yes	5	0	Yes	No	Yes	180
System Waits	Yes	5	0	Yes	Yes	Yes	180
Tablespace Size	Yes	30	0	Yes	Yes	Yes	180
User Information	Yes	30	0	Yes	Yes	Yes	180

Figure 6.2 - The EDB Audit Configuration probe.

If EDB Audit Configuration is not enabled, use the Enabled? switch on the Manage Probes tab to enable it.

6.3 Configuring Audit Logging with the Audit Manager

To open the Audit manager wizard, select Audit Manager... from the Management menu. The Audit manager - Welcome dialog opens as shown in Figure 6.3.

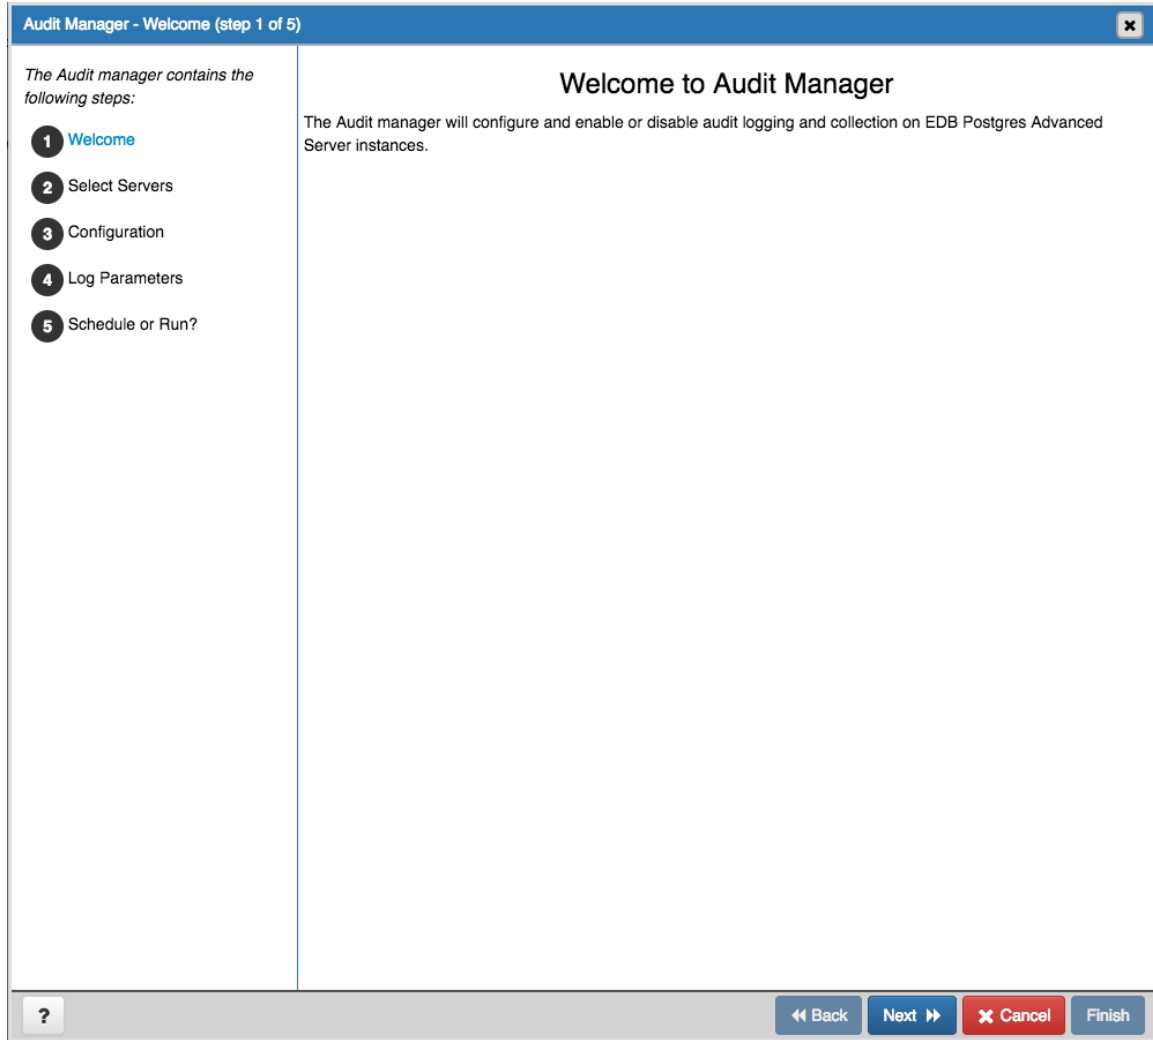


Figure 6.3 - The Audit Manager Welcome dialog.

Click **Next** to continue.

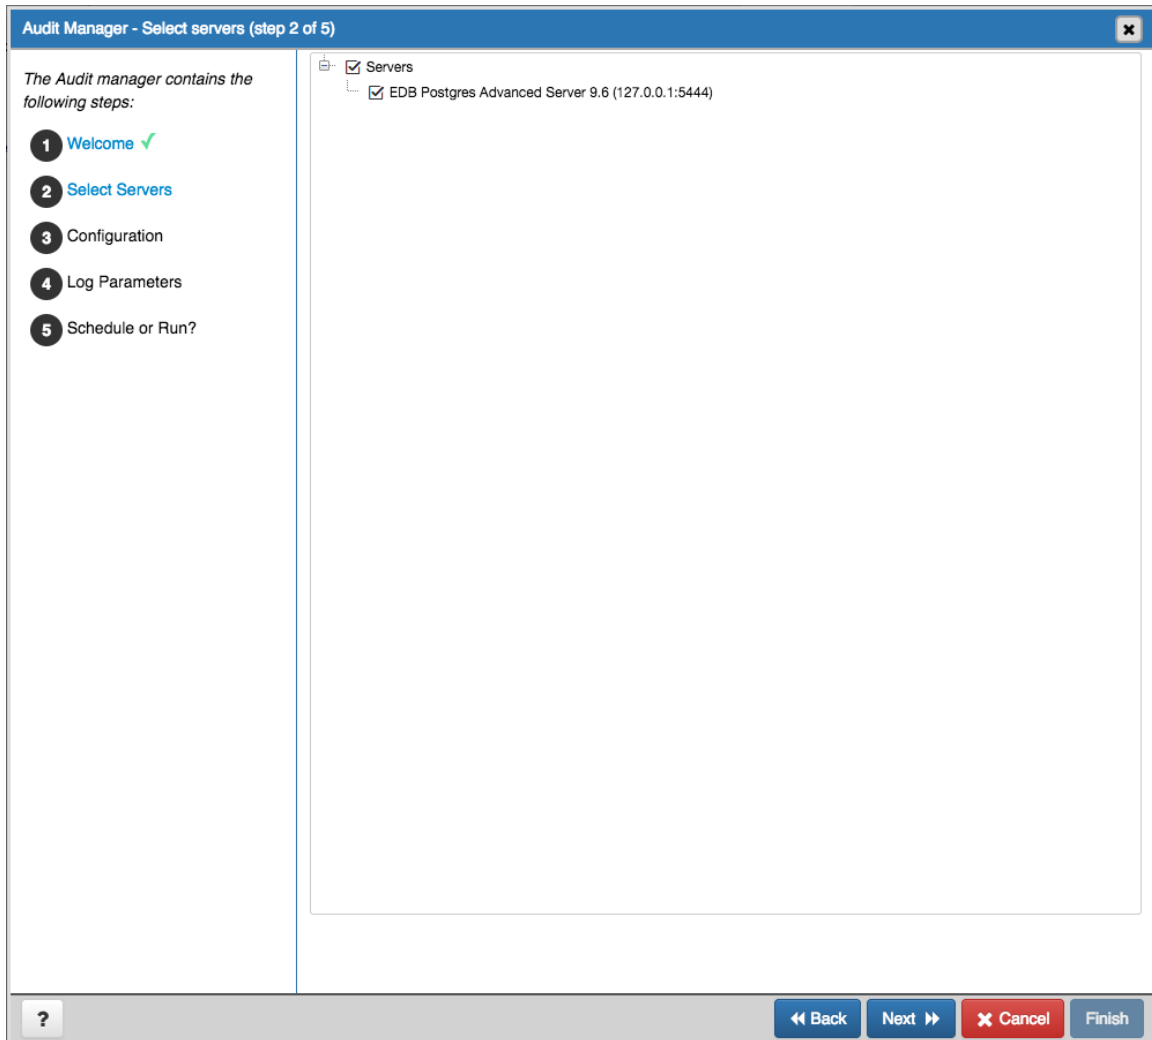


Figure 6.4 – Select the servers you wish to configure for auditing.

Use the `Select servers` tree control (shown in Figure 6.4) to specify the servers to which the auditing configuration will be applied. To make a server available in the tree control, you must provide the `Service ID` on the `Advanced` tab of the `Create – Server` dialog when registering a server for monitoring by PEM. Note that only EDB Postgres Advanced Server supports auditing; PostgreSQL servers will not be included in the tree control.

Click `Next` to continue.

The Auditing Parameters Configuration dialog lets you enable or disable auditing and choose how often log records are collected into PEM (see Figure 6.5).

Audit Manager - Audit parameters configuration (step 3 of 5)

The Audit manager contains the following steps:

- 1 Welcome ✓
- 2 Select Servers ✓
- 3 Configuration
- 4 Log Parameters
- 5 Schedule or Run?

Auditing ☒ Enable

Audit destination

Import logs to PEM ☐ No

Import frequency 1 Hour

Log format CSV

File name audit-%Y-%m-%d_%H%M%S
Specifies the file name of audit file where the auditing information will be stored.

Log directory

Change log directory for selected servers? ☐ No

Directory name edb_audit
Specifies the directory where the log files will be created. The path of the directory can be relative or absolute to the data folder.

Navigation: ? << Back Next >> Cancel Finish

Figure 6.5 – The Auditing Parameters Configuration dialog.

Use the fields on the Auditing parameters configuration dialog to specify auditing preferences:

- Use the Auditing switch to Enable or Disable auditing on the specified servers.
- Use the Audit destination drop-down to select a destination for the audit logs; select File or Syslog. Please note this feature is supported on Advanced Server 10 and newer releases only.

- Use the `Import logs to PEM` switch to instruct PEM to periodically import log records from each server to the PEM Server. Set the switch to `Yes` to import log files; the default is `No`.
- Use the `Collection frequency` drop-down list box to specify how often PEM will collect log records from monitored servers when log collection is enabled.
- Use the `Log format` drop-down list box to select the raw log format that will be written on each server. If log collection is enabled, the PEM server will use CSV format.
- Use the `File name` field to specify the format used when generating log file names. By default, the format is set to `audit-%Y-%m-%d_%H%M%S` where:

`audit` is the file name specified in the `Audit Directory Name` field
`Y` is the year that the log was stored
`m` is the month that the log was stored
`d` is the day that the log was stored
`H` is the hour that the log was stored
`M` is the minute that the log was stored
`S` is the second that the log was stored

- Check the box next to `Change Log Directory for selected servers?` and use the `Audit Directory Name` field to specify a directory name to which the audit logs will be written. The directory will reside beneath the `data` directory on the PEM server.

Use fields in the `Log directory` box to specify information about the directory in which the log files will be saved:

- Move the `Change log directory for selected servers?` switch to `Yes` to enable the `Directory name` field.
- Use the `Directory name` field to specify the name of the directory on each server into which audit logs will be written. The directory specified will be created as a sub-directory of the `data` directory on the server.

Click `Next` to continue.

The Audit log configuration dialog (see Figure 6.6) is only available if you have enabled auditing on the Auditing parameters configuration dialog.

Audit Manager - Audit log configuration (step 4 of 5)

The Audit manager contains the following steps:

- 1 Welcome ✓
- 2 Select Servers ✓
- 3 Configuration ✓
- 4 Log Parameters
- 5 Schedule or Run?

Connection attempts
 Enables auditing of database connection attempts by users. To disable auditing of all connection attempts, set edb_audit_connect to none. To audit all failed connection attempts, set the value to failed.

Disconnection attempts
 Enables auditing of database disconnections by connected users. To enable auditing of disconnections, set the value to all. To disable, set the value to none.

Log statements
 This configuration parameter is used to specify auditing of different categories of SQL statements.

Audit tag
 Use this configuration parameter to specify a string value. that will be included in audit log when the edb_audit parameter is set to csv or xml.

Log rotation

Enable ?

Day

Size(MB)

Time(seconds)

Navigation: ? << Back Next >> X Cancel Finish

Figure 6.6 – The Audit Log Configuration dialog.

Use the controls on the Audit log configuration dialog to specify log configuration details that will be applied to each server:

- Use the Connection attempts switch to specify if connection attempts should be logged. Specify:

None to disable connection logging.

All to indicate that all connection attempts will be logged.

Failed to log any connection attempts that fail.

- Use the `Disconnection attempts` switch to specify if disconnections should be logged. Specify:

`None` to specify that disconnections should not be logged.

`All` to enable disconnection logging.

- Use the `Log statements` field to specify the statement types that will be logged. Click within the field, and select from:

`Select` - All statements that include the `SELECT` keyword will be logged.

`Error` - All statements that result in an error will be logged.

`DML` - All DML (Data Modification Language) statements will be logged.

`DDL` - All DDL (Data Definition Language) statements (those that add, delete or alter data) will be logged.

Check the box next to `Select All` to select all statement types.

Check the box next to `Unselect All` to deselect all statement types.

- Use the `Audit tag` field to specify a tracking tag for the collected logs. Please note that audit tagging functionality is available only for Advanced Server versions 9.5 and later. If you are defining auditing functionality for multiple servers, and one or more of the servers are version 9.5 or later, this field will be enabled, but if selected, tagging functionality will only apply to those servers that are version 9.5 or later.

Use the fields in the `Log rotation` box to specify how the log files are managed on each server:

- Use the `Enable?` switch to specify that logfiles should be rotated. Please note that a new log file should be used periodically to prevent a single file becoming unmanageably large.
- Use the `Day` drop-down listbox to select a day or days on which the log file will be rotated.
- Use the `Size (MB)` field to specify a size in megabytes at which the log file will be rotated.
- Use the `Time (seconds)` field to specify the number of seconds between log file rotations.

Click `Next` to continue:

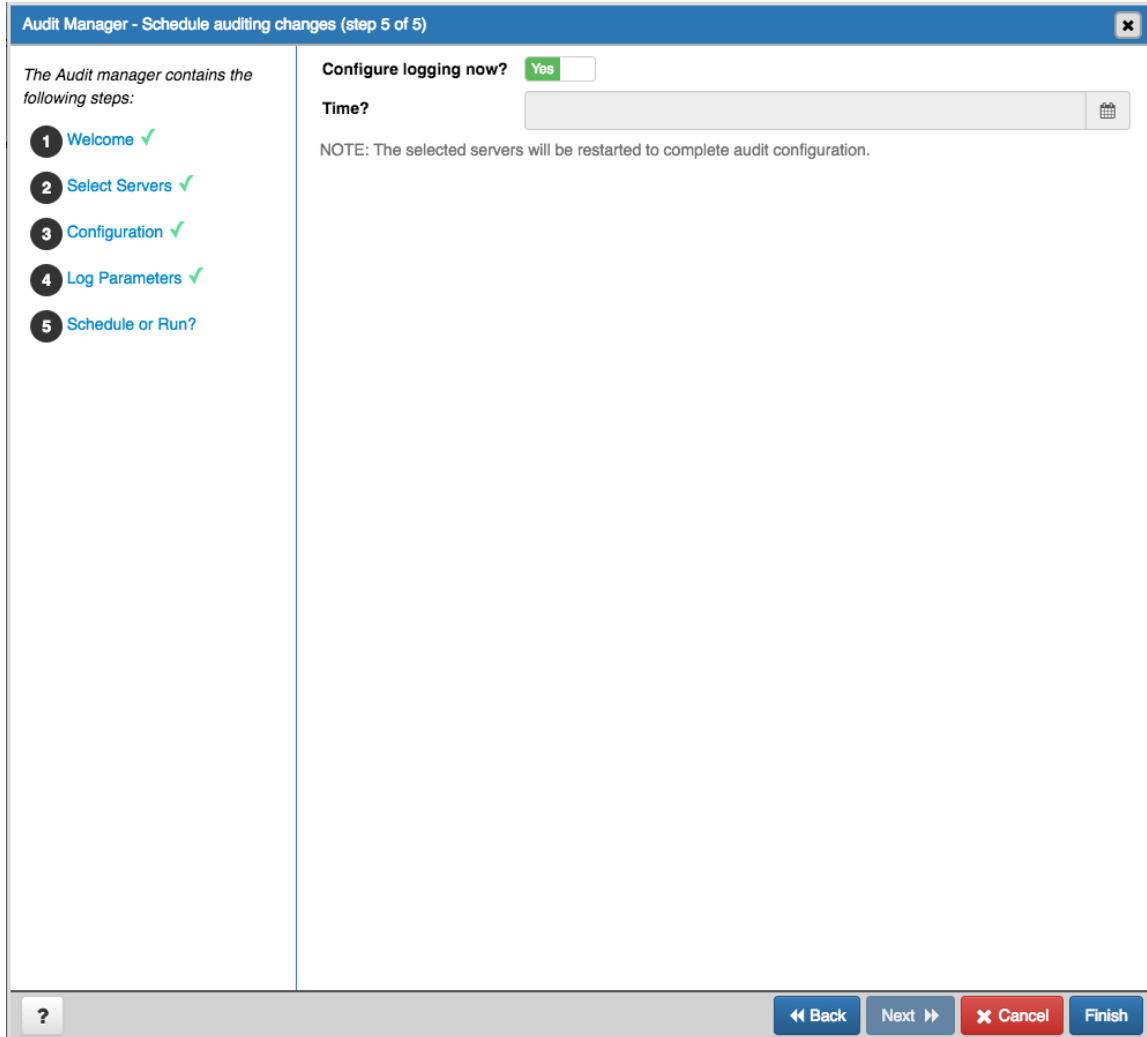


Figure 6.7 – The Schedule Auditing Changes dialog.

Use the `Schedule Auditing Changes` dialog (see Figure 6.7) to determine when auditing configuration changes are to take effect.

- Select `Configure logging now?` if you want the auditing configuration changes to take place immediately. The affected database servers will be restarted so the auditing changes can take effect.
- Use the `Time?` selector to schedule the auditing configuration changes to take place at some point in the future. Select the desired date and time from the drop-down lists. The affected database servers will be restarted at the specified date/time to put the auditing changes into effect.

Click `Finish` to complete the auditing configuration process.

The Audit Manager will schedule a job to apply the configuration to each server. The job will consist of two tasks: one to update the audit logging configuration on the server, and one to restart the server with the new configuration.

You can use the `ScheduledTasks` tab to review a list of Scheduled jobs. To open the `Scheduled Tasks` tab, highlight the name of a server or agent and select `Scheduled Tasks...` from the `Management` menu.

6.4 Viewing the Log with the Audit Log Dashboard

Use the `Audit Log` dashboard to view the audit log from Advanced Server database instances (see Figure 6.8).

To open the `Audit Log` dashboard, right click on a server or agent node, and select `Audit Log Analysis` from the `Dashboards` menu. You can also open the `Audit Log` dashboard by navigating through the `Dashboards` menu (located on the `Management` menu).

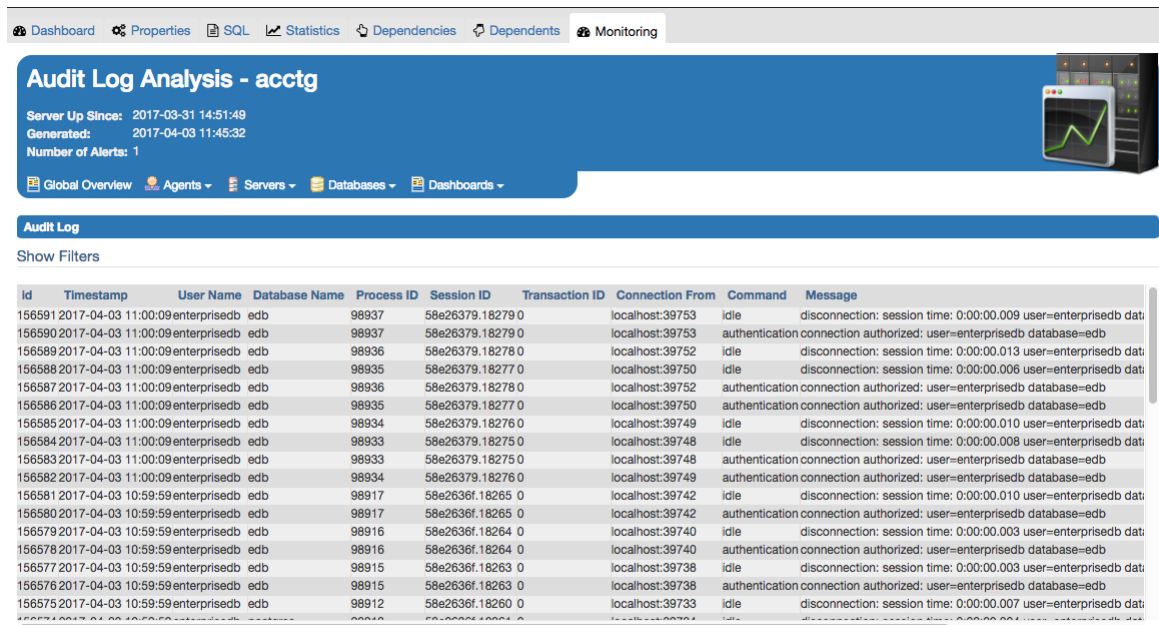


Figure 6.8 – The Audit Log dashboard.

The `Audit Log` dashboard displays the audit records in reverse chronological order (newest records at the top, oldest records towards the bottom).

To view older audit records that do not appear in the window, use the vertical scroll bar controlling the list of audit records (the innermost scroll bar of the two located on the right-hand side of the window). As you move the scroll bar towards the bottom of the window, older audit records are continuously loaded and displayed.

You can use filtering to limit the number of audit records that are displayed. Click `Show Filters` to expose the filters panel (see Figure 6.9).

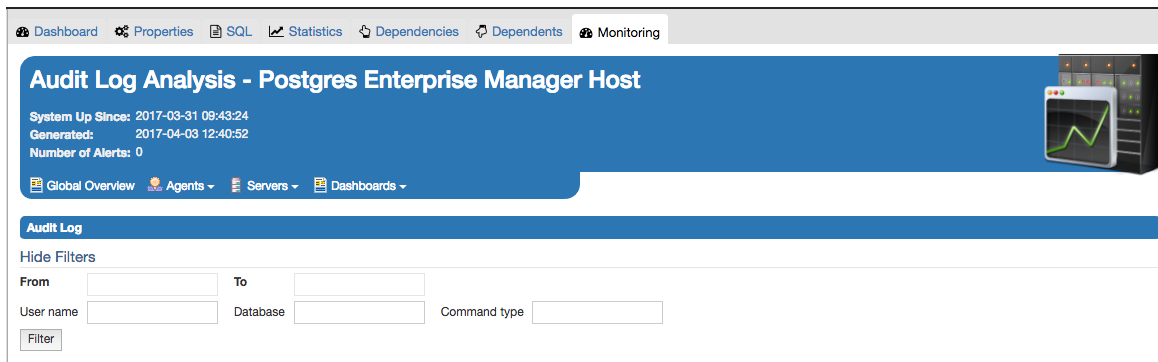


Figure 6.9 – The Audit Log dashboard filters panel.

Use the fields in the filters panel to provide certain selection criteria for the audit records you wish to display.

- Use the `From` field to specify a start date for the report. Click the mouse button in the field to open a calendar and select a start date.
- Use the `To` field to specify an end date for the report. Click the mouse button in the field to open a calendar and select an end date.
- Use the `Username` field to display only those entries where the activity was initiated by the given Postgres user.
- Use the `Database` field to display only those entries where the activity was issued on the given database.
- Use the `Command type` field to display only those entries where the activity was of the given type. Command types you can specify are `idle`, `authentication`, and `SELECT`. (For viewing SQL statements from user applications, specify the `idle` command type.)

Click `Filter` to apply the filtering criteria to the log entries.

7 Log Manager

You can use the PEM Log Manager to simplify server log configuration for Postgres instances. With the Log Manager, you can modify all of your server log parameters with a click:

- Where log files are written
- How often log files are written
- The type of information written to log files
- The format of log file entries
- Log rotation properties

To configure logging for a Postgres instance, the server must be registered as a PEM-managed server, and the registration information must include the name of a service script.

To open the Log Manager, select the `Log Manager...` option from the `Management` menu of the PEM client. The wizard opens, welcoming you to the Log Manager (as shown in Figure 7.1).

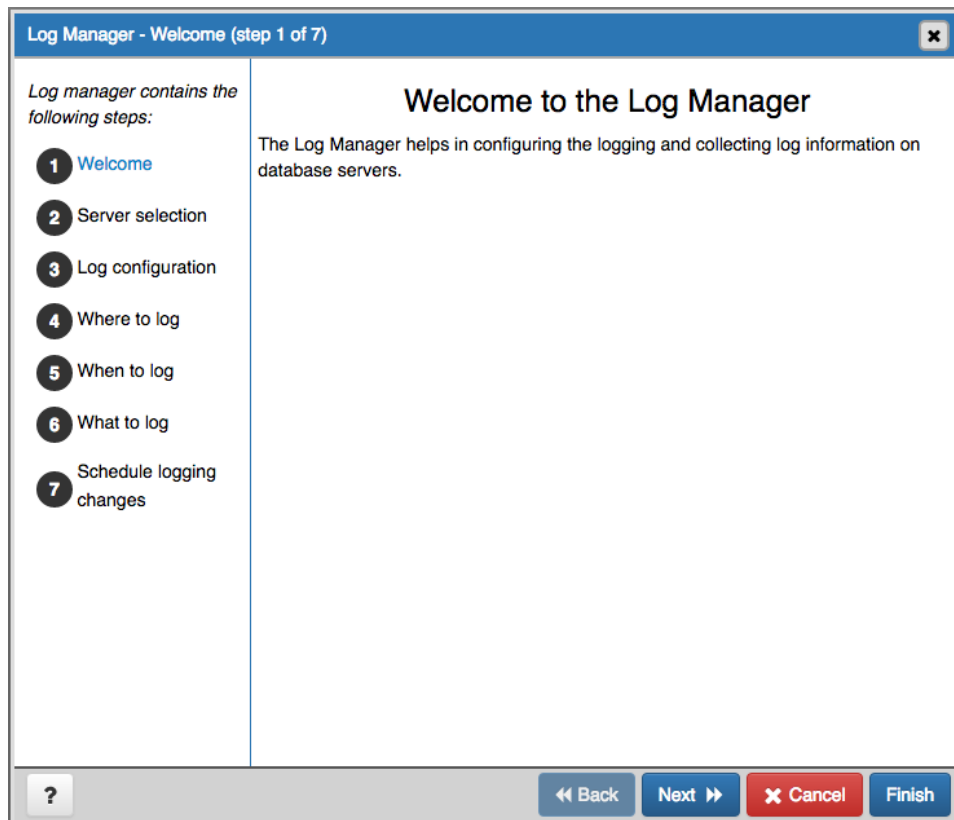


Figure 7.1 - The Log Manager welcome dialog.

Click `Next` to continue to the `Server selection` dialog (see Figure 7.2).

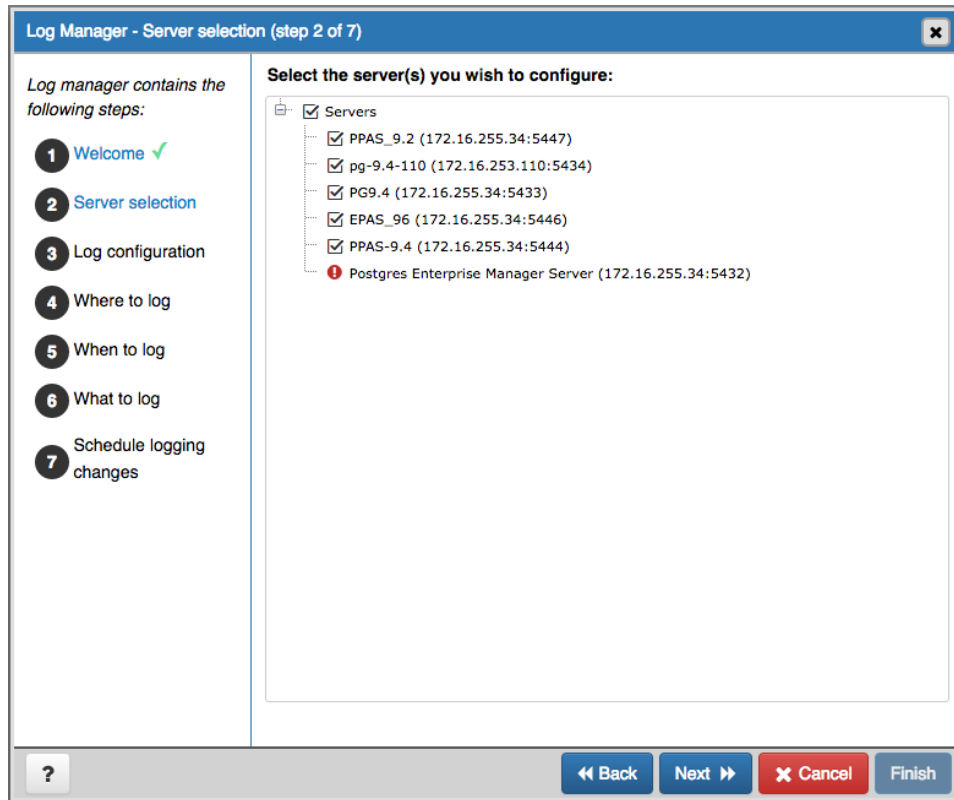


Figure 7.2 - The Log Manager Server selection dialog.

The `Server selection` dialog displays a list of the server connections monitored by PEM. Check the box next to the name of a server (or servers) to which the Log Manager wizard will apply the specified configuration. Log Manager is disabled for any server displaying a red exclamation mark to the left of its name in the `Server selection` tree control; there are several reasons that a server may not be enabled:

- Only a server that specifies a `Service ID` on the `Advanced` tab of the `Properties` dialog can be configured by Log Manager.

To provide a service ID, right click on the server name in the tree control, and select `Disconnect Server` from the context menu; if prompted, provide a password. Then, open the context menu for the server, and select `Properties`. Navigate to the `Advanced` tab, and provide the name of the service in the `Service ID` field; click `Save` to save your change and exit the dialog.

- If the PEM agent bound to the server does not have sufficient privileges to restart the server, the server will be disabled.

- If the PEM agent bound to the server is an older version than the associated PEM server, the server will be disabled.

Click `Next` to continue.

Log Manager - Log configuration (step 3 of 7)

Log manager contains the following steps:

- 1 Welcome ✓
- 2 Server selection ✓
- 3 Log configuration
- 4 Where to log
- 5 When to log
- 6 What to log
- 7 Schedule logging changes

Import logs

Import logs to PEM ☐ No

Import frequency 5 Minutes

Log rotation configuration

Rotation size (in MB) 10

Automatic rotation of log files will happen after that much log output. 0 disables.

Rotation time (in days) 1

Automatic rotation of log files will happen after that time. 0 disables.

Truncate on rotation ☐ Off

If on, an existing log file with the same name as the new log file will be truncated rather than appended to. But such truncation only occurs on time-driven rotation.

Back Next Cancel Finish

Figure 7.3 - The Log Manager Log configuration dialog.

Use the options on the `Log configuration` dialog (shown in Figure 7.3) to specify how often log files will be imported to PEM and to specify log rotation details:

Options within the `Import Logs` box specify how often log files will be imported to PEM:

- Use the switch next to the `Import logs to PEM` label to specify if log files will be imported to PEM and displayed on the Server Log Analysis dashboard.
- Use the `Import Frequency` drop-down list box to specify how often log files are imported to PEM.

Use the fields in the `Log rotation configuration` box to specify the maximum length (lifespan or size) of a log file:

- Use the `RotationSize` field to specify the maximum size in megabytes of an individual log file. The default value is 10 MB; when set to 0, no limit is placed on the maximum size of a log file.
- Use the `RotationTime` field to specify the number of whole days that should be stored in each log file. The default value is 1 day.

Use the `Truncation on Rotation` switch to specify server behavior for time-based log file rotation:

- Select `ON` to specify that the server should overwrite any existing log file that has the same name that a new file would take.
- Select `OFF` to specify that the server should append any new log file entries to an existing log file with the same name that a new log file would take. This is the default behavior.

Click `Next` to continue.

Log Manager - Where to log (step 4 of 7)

Log manager contains the following steps:

- Welcome ✓
- Server selection ✓
- Log configuration ✓
- Where to log
- When to log
- What to log
- Schedule logging changes

Log destination

stderr ☒ Yes ☐ No csvlog ☐ No ☒ No

syslog ☐ No ☒ Yes eventlog ☐ No ☒ Yes

Log collection

Log collector ☒ Enable ☐ Disable Log silent mode ☐ Disable ☒ Enable

Log directory

Change log directory for selected servers? ☐ No ☒ Yes

Directory name

Log file name

Syslog facility
 This parameter determines the syslog "facility" to be used.

Syslog ident
 This parameter determines the program name used to identify PostgreSQL messages in syslog logs.

◀ Back Next ▶ ✕ Cancel Finish

Figure 7.4 – The Where to Log dialog.

Use the fields on the `Where to log` dialog (shown in Figure 7.4) to specify where log files should be written.

- Select an option from the `Log Destination` box to specify a destination for the server log output:
 - Set the `stderr` switch to `Yes` to specify that log files should be written to `stderr`.
 - Set the `csvlog` switch to `Yes` to specify that log files should be written to file in a comma-separated value format. This option is automatically enabled (and no longer editable) if you have selected `Import logs to PEM` on the `Schedule` dialog; if you are not importing server log files to PEM, this option is editable.
 - Set the `syslog` switch to `Yes` to specify that log files should be written to the system log files.
 - On Windows, set the `eventlog` switch to `Yes` to specify that log files should be written to the event log.
- Use the options within the `Log collection` box to specify your collection preferences:
 - Set the `Log Collector` switch to `Enable` to instruct the server to re-direct captured log messages (directed to `STDERR`) into log files.
 - Set the `Log Silent Mode` switch to `Enable` to instruct the server to run silently in the background, disassociated from the controlling terminal.
- Use options in the `Log Directory` box to specify log file location preferences:
 - Set the `Change log directory for selected servers?` switch to `Yes` to specify that each set of log files should be maintained in a separate directory.
 - Use the `Directory name` field to specify the directory to which log files will be written. The directory will reside beneath the `pg_log` directory under the installation directory of the monitored server.
- Use the `Log File Name` field to specify a format for the log file name. If set to `DEFAULT`, the format is `enterprisedb-%Y-%m-%d_%H%M%S`, where:

`enterprisedb` is the file name prefix
`Y` is the year that the log was stored

m is the month that the log was stored
 d is the day that the log was stored
 H is the hour that the log was stored
 M is the minute that the log was stored
 S is the second that the log was stored

When logging to `syslog` is enabled:

- Use the `Syslog Facility` drop-down list box to specify which syslog facility should be used.
- Use the `Syslog Ident` field to specify the program name that will identify Advanced Server entries in system logs.

Click `Next` to continue.

Log Manager - When to log (step 5 of 7)

Log manager contains the following steps:

- Welcome ✓
- Server selection ✓
- Log configuration ✓
- Where to log ✓
- When to log**
- What to log
- Schedule logging changes

Client min messages	notice	Message levels sent to the client.
Log min messages	warning	Message levels written to the server log.
Log min error statement	error	SQL statements that cause an error condition are recorded in the server log.
Log min duration statement (in milliseconds)	-1	Logs only statements running at least this number of milliseconds. -1: disabled, 0 logs all with durations.
Log temp files (in kilobytes)	-1	Log temporary files equal or larger than the specified size; -1 disables, 0 logs all temp files.
Log autoVacuum min duration (in milliseconds)	-1	

? << Back Next >> X Cancel Finish

Figure 7.5 - The Log Manager When to Log dialog.

Use the fields on the `When to log` dialog (shown in Figure 7.5) to specify which events will initiate a log file entry. The severity levels (in order of severity, from most severe to least severe) are:

`panic` - Errors that cause all database sessions to abort.

`fatal` - Errors that cause a session to abort.

`log` - Information messages of interest to administrators.

`error` - Errors that cause a command to abort.

`warning` - Error conditions in which a command will complete but may not perform as expected.

`notice` - Items of interest to users. This is the default.

`info` - Information implicitly requested by the user.

`debug5` through `debug1` - Detailed debugging information useful to developers.

- Use the `Client min messages` drop-down list box to specify the lowest severity level of messages sent to the client application.
- Use the `Log min messages` drop-down list box to specify the lowest severity level that will be written to the server log.
- By default, when an error message is written to the server log, the text of the SQL statement that initiated the log entry is not included. Use the `Log min error statement` drop-down list box to specify a severity level that will trigger SQL statement logging. If a message is of the specified severity or higher, the SQL statement that produced the message will be written to the server log.
- Use the `Log min duration statement` drop-down list box to specify a statement duration (in milliseconds); any statements that exceed the specified number of milliseconds will be written to the server log. A value of `-1` disables all duration-based logging; a value of `0` logs all statements and their duration.
- Use the `Log temp files` field to specify a file size in kilobytes; when a temporary file reaches the specified size, it will be logged. A value of `-1` (the default) disables this functionality.
- Use the `Log autoVacuum min duration` field to specify a time length in milliseconds; if auto-vacuuming exceeds the length of time specified, the activity will be logged. A value of `-1` (the default) disables this functionality.

Click `Next` to continue.

Figure 7.6 - The Log Manager What to Log dialog.

Use the fields on the `What to log` dialog (shown in Figure 7.6) to specify log entry options that are useful for debugging and auditing.

The switches in the `Debug options` box instruct the server to include information in the log files related to query execution that may be of interest to a developer:

- Set the `Parse tree` switch to `Yes` to instruct the server to include the parse tree in the log file.
- Set the `Rewriter output` switch to `Yes` to instruct the server to include query rewriter output in the log file.
- Set the `Execution plan` switch to `Yes` to instruct the server to include the execution plan for each executed query in the log file.

When the `Indent Debug Options Output in Log` switch is set to `Yes`, the server indents each line that contains a parse tree entry, a query rewriter entry or query execution plan entry. While indentation makes the resulting log file more readable, it does result in a longer log file.

Use the switches in the `General Options` box to instruct the server to include auditing information in the log file:

- Set the `Checkpoints` switch to `Yes` to include checkpoints and restartpoints in the server log.
- Set the `Connections` switch to `Yes` to include each attempted connection to the server (as well as successfully authenticated connections) in the server log.
- Set the `Disconnections` switch to `Yes` to include a server log entry for each terminated session that provides the session information and session duration.
- Set the `Duration` switch to `Yes` to include the amount of time required to execute each logged statement in the server log.
- Set the `Hostname` switch to `Yes` to include both the IP address and host name in each server log entry (by default, only the IP address is logged). Please note that this may cause a performance penalty.
- Set the `Lock Waits` switch to `Yes` to instruct the server to write a log entry for any session that waits longer than the time specified in the `deadlock_timeout` parameter to acquire a lock. This is useful when trying to determine if lock waits are the cause of poor performance.

Use the `Error verbosity` drop-down list box to specify the detail written to each entry in the server log:

- Select `default` to include the error message, `DETAIL`, `HINT`, `QUERY` and `CONTEXT` in each server log entry.
- Select `terse` to log only the error message.
- Select `verbose` to include the error message, the `DETAIL`, `HINT`, `QUERY` and `CONTEXT` error information, `SQLSTATE` error code and source code file name, the function name, and the line number that generated the error.

Use the `Prefix string` field to specify a printf-style string that is written at the beginning of each log file entry. For information about the options supported, please see the `log_line_prefix` documentation (in the Postgres core documentation), available at:

<http://www.postgresql.org/docs/10/static/runtime-config-logging.html>

Use the `Statements` drop-down list box to specify which SQL statements will be included in the server log. The default is `none`; valid options are:

- Specify `none` to disable logging of SQL statements.
- Specify `ddl` to instruct the server to log ddl (data definition language) statements, such as `CREATE`, `ALTER`, and `DROP`.
- Specify `mod` to instruct the server to log all ddl statements, as well as all dml (data modification language) statements, such as `INSERT`, `UPDATE`, `DELETE`, `TRUNCATE` and `COPY FROM`.
- Specify `all` to instruct the server to log all SQL statements.

Click `Next` to continue.

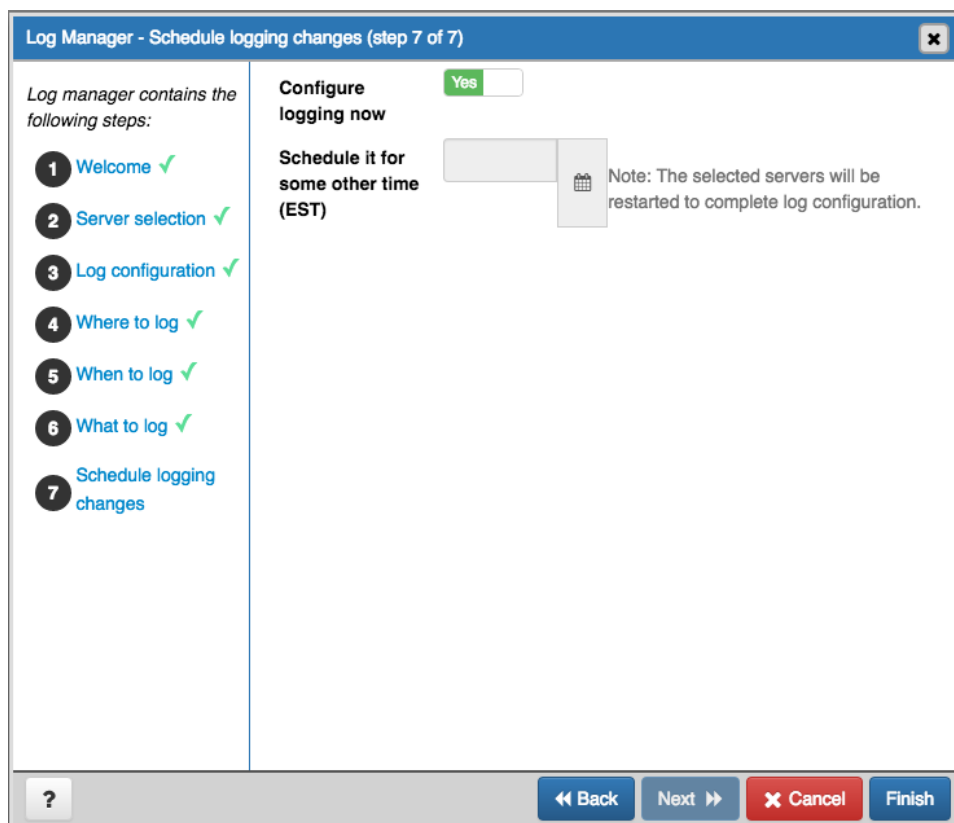


Figure 7.7 - The Schedule Logging Changes dialog.

Use options on the `Schedule logging changes` dialog (see Figure 7.7) to specify when logging configuration changes will be applied:

- Set the `Configure logging now` switch to `Yes` to specify that your configuration preferences will be enabled, and the server will restart when you have completed the Log Manager wizard.

- Set `Configure logging now` to `No` to use the `Schedule it for some other time` calendar selector to specify a convenient time for logging configuration preferences to be applied, and the server to restart.

Note that when you apply the configuration changes specified by the Log Manager wizard, the server restart will temporarily interrupting use of the database server for users.

Click `Finish` to exit the wizard, and either restart the server, or schedule the server restart for the time specified on the scheduling dialog.

7.1 Reviewing the Server Log Analysis Dashboard

After invoking the Log Manager wizard, and importing your log files to PEM, you can use the Server Log Analysis dashboard to review the log files for a selected server. To open the Server Log Analysis dashboard, right-click on the name of a monitored server in the PEM client tree control, and navigate through the Dashboards menu, selecting Server Log Analysis.

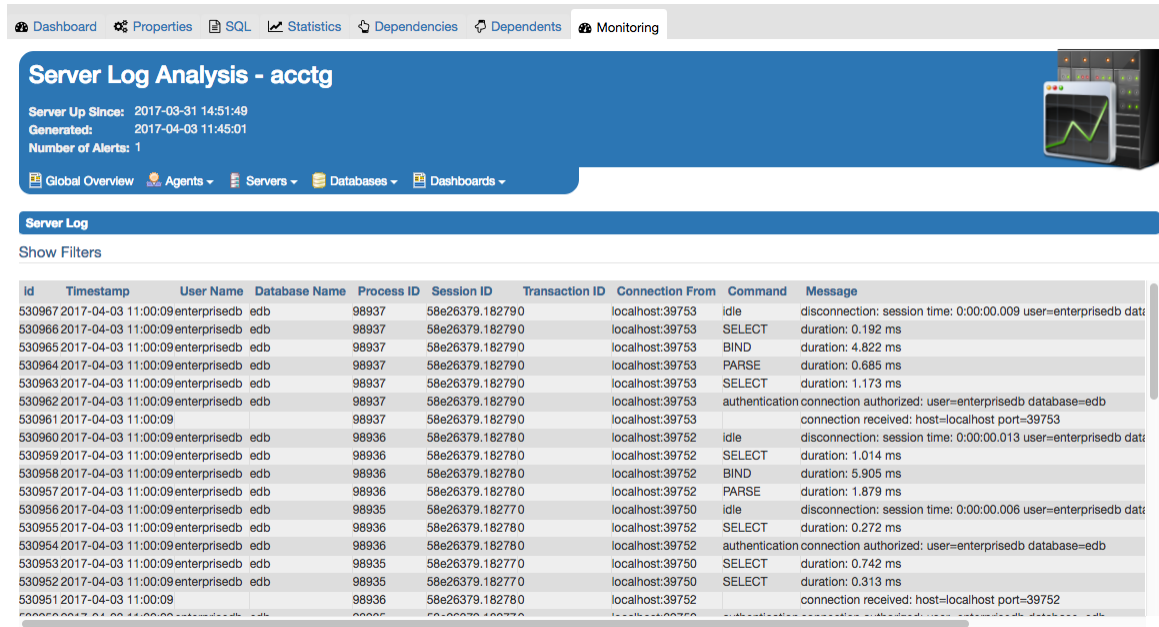


Figure 7.8 - The Server Log Analysis dashboard.

The header information on the Server Log Analysis dashboard (shown in Figure 7.8) displays the date and time that the server was started, the date and time that the page was last updated, and the current number of triggered alerts.

Entries in the Server Log are displayed in chronological order, with the most-recent log entries first. Use the scroll bars to navigate through the log entries, or to view columns that are off of the display.

Headings at the top of the server log table identify the information stored in each column; hover over a column heading to view a tooltip that contains a description of the content of each column.

You can use filtering to limit the number of server log records that are displayed. Click Show Filters to expose the filters panel and define a filter (see Figure 7.9).

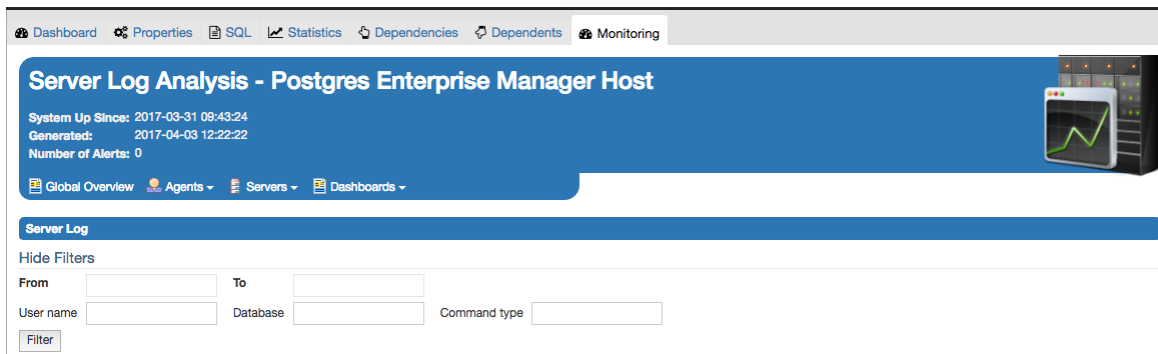


Figure 7.9 - Defining a Server Log filter.

Use the fields within the filter definition box to describe the selection criteria that PEM will use to select a subset of a report for display:

- Use the `From` field to specify a starting date for the displayed server log.
- Use the `To` field to specify an ending date for the displayed server log.
- Enter a role name in the `Username` field display only transactions performed by that user.
- Enter a database name in the `Database` field to specify that the server should limit the displayed records to only those transactions that were performed against the specified database.
- Use the `Command Type` field to specify a selection criteria for the commands that will be displayed in the filtered report.

When you've described the criteria by which you wish to filter the server logs, click `Filter` to display the filtered server log in the `Server Log` table.

8 Postgres Log Analysis Expert

The PEM Log Analysis Expert analyzes the log files of servers that are registered with Postgres Enterprise Manager, and produces a report that provides an analysis of your Postgres cluster's usage based on log file entries. You can use information on the Log Analysis Expert reports to make decisions about optimizing your cluster usage and configuration to improve performance.

Before using the PEM Log Analysis Expert, you must specify the `Service ID` on the `Advanced` tab of the `Server Properties` dialog, and use the Log Manager wizard to enable log collection by the PEM server.

To open the Postgres Log Analysis Expert wizard, select the `Postgres Log Analysis Expert...` option from the `Management` menu of the PEM client. The wizard's `Welcome` dialog opens as shown in Figure 8.1, click `Next` to continue:

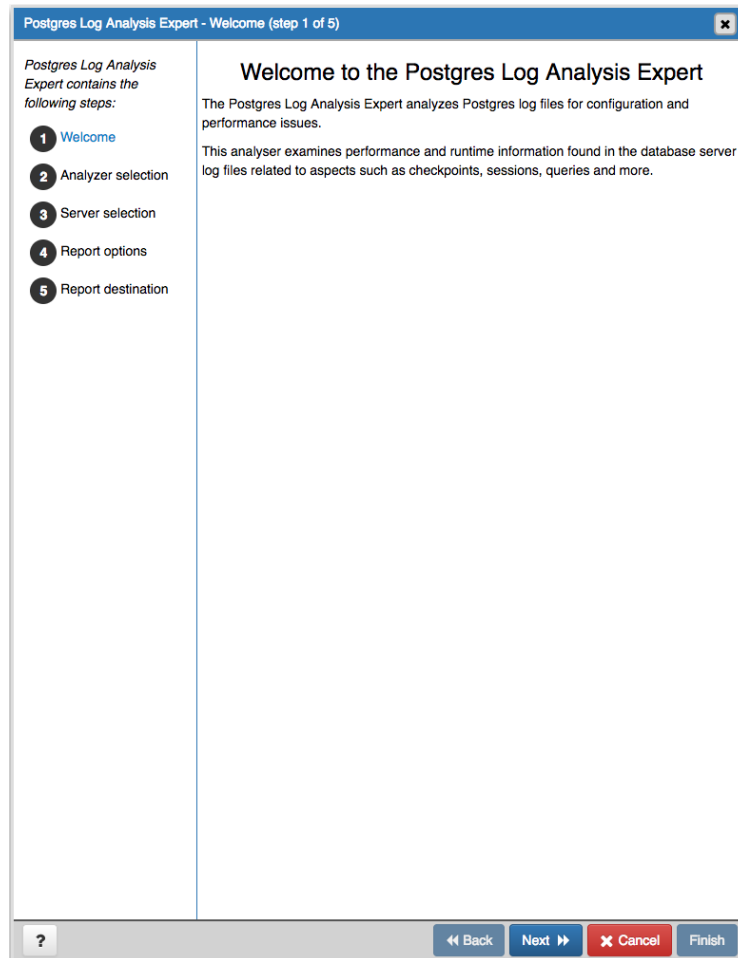


Figure 8.1 – The Log Analysis Expert Welcome dialog.

The wizard's `Analyzer selection` dialog (see Figure 8.2) displays a list of `Analyzers` from which you can select. Each `Analyzer` generates a corresponding table, chart, or graph that contains information gleaned from the log files.

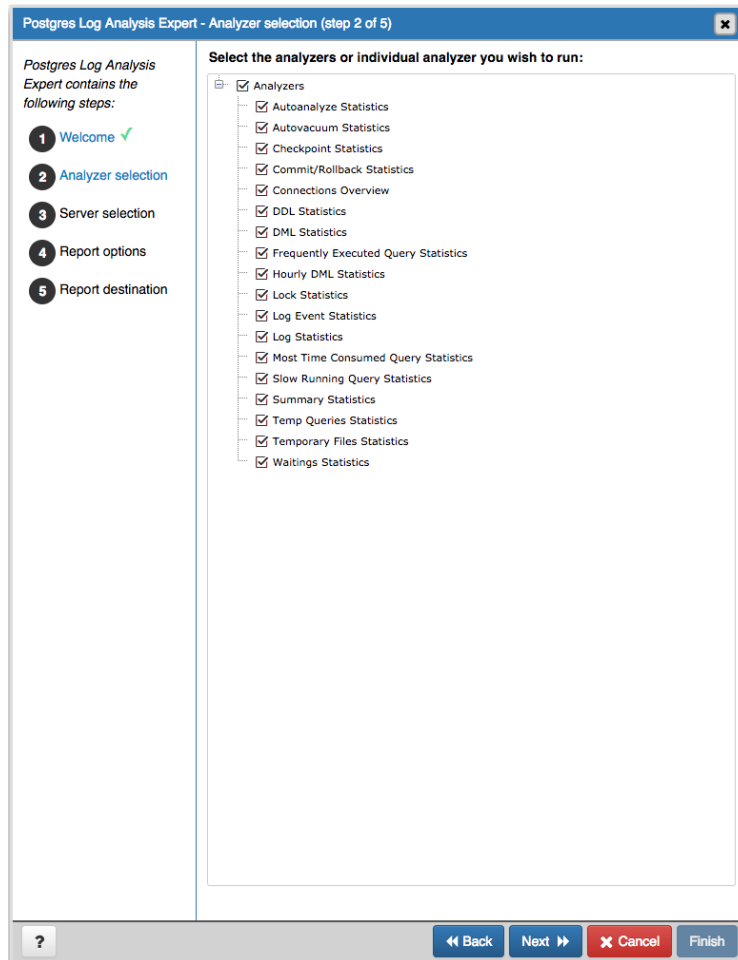


Figure 8.2 – The Analyzer selection dialog.

Check the box to the left of an `Analyzer` to indicate that the Log Analysis Expert should prepare the corresponding table, chart or graph. After making your selections, click `Next` to continue to the `Server selection` tree control (see Figure 8.3).

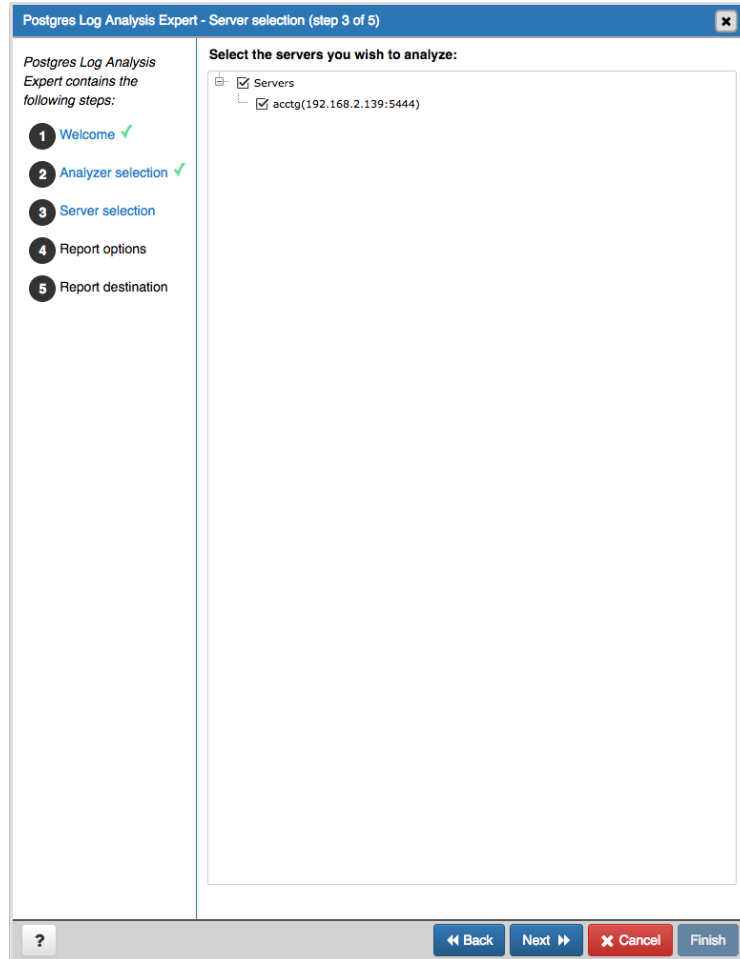


Figure 8.3 – The Server selection dialog.

Use the tree control to specify which servers you would like the Postgres Log Analysis Expert to analyze. If you select multiple servers, the resulting report will contain the corresponding result set for each server in a separate (but continuous) list. Click **Next** to continue to the **Report options** dialog (see Figure 8.4).

Figure 8.4 – The Report options dialog.

Use the fields in the `Options` section to specify the analysis method and the maximum length of any resulting tables:

- Use the `Aggregate method` drop-down to select the method used by the Log Analysis Expert to consolidate data for the selected time span. You can select from:
 - `SUM` instructs the analyzer to calculate a value that is the sum of the collected values for the specified time span.
 - `AVG` instructs the analyzer to calculate a value that is the average of the collected values for the specified time span.
 - `MAX` instructs the analyzer to use the maximum value that occurs within a specified time span.

- `MIN` instructs the analyzer to use the minimum value that occurs within a specified time span.
- Use the `Time span` field to specify the number of minutes that the analyzer will incorporate into each calculation for a point on a graph. For example, if the `Time span` is 5 minutes, and the `Aggregate method` is `AVG`, each point on the given graph will contain the average value of the activity that occurred within a five minute time span.
- Use the `Rows limit` field to specify the maximum number of rows to include in a table.

Use the fields in the `Time Intervals` section to specify the time range that the Log Analysis Expert will analyze:

- Set `Relative days` to `Yes` to enable the `(+/-) From date` field and specify the number of days before or after the date and time selected in the `From` field.
- Use the `From` field to specify the starting date and time for the analysis.
- Use the `To` field to specify the ending date and time for the analysis.
- Use the `(+/-) From date` selector to specify the number of days before or after the `From date` that should be included in the analysis.

When you've specified the report options, click `Next` to continue to the `Report destination` dialog (see Figure 8.5).

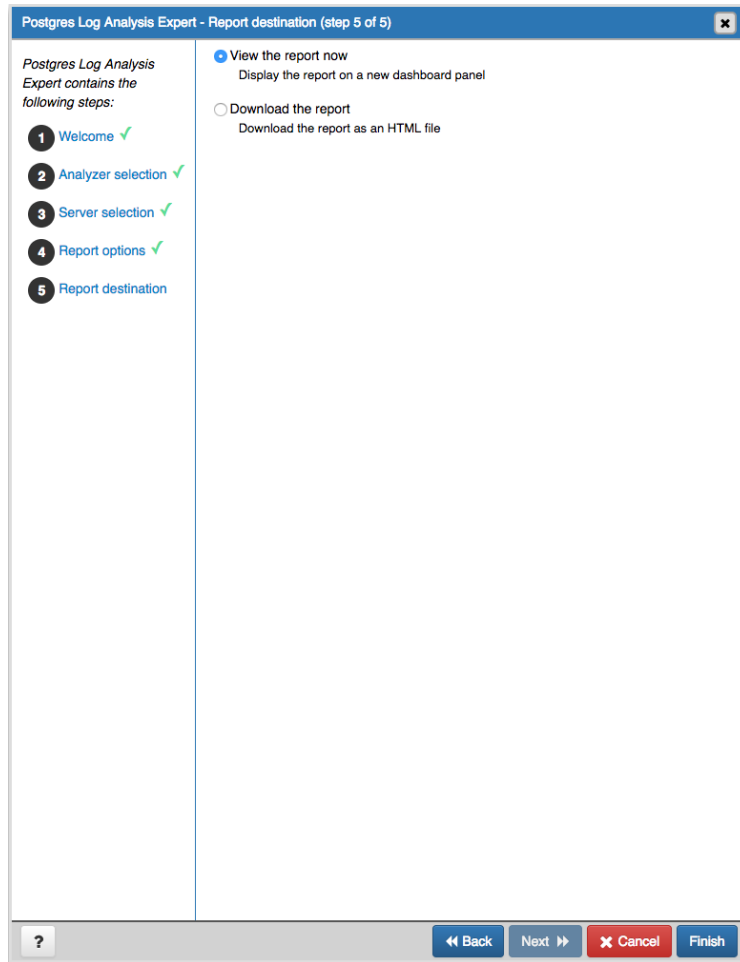


Figure 8.5 – The Report destination dialog.

You can choose the default option and select `Finish` to view the Log Analysis Expert report in the PEM client's tabbed browser, or click the radio button next to `Download the report` to save a copy of the report to an HTML file for later use.

If you have specified that the report should be saved to a file, the report will be downloaded.

8.1 Reviewing the Postgres Log Analysis Expert Report

If you've elected to review the report immediately, the Postgres Log Analysis Expert report will be displayed in the PEM Client window (see Figure 8.6). The report header displays the date and time that the report was generated, the time period that the report spans, and the aggregation method specified when defining the report. The name of the server for which information is displayed is noted at the start of each section of the report.

The report displays the tables, graphs and charts that were selected in the Log Analysis Expert wizard. Use the **Jump To** button (located in the lower-right hand corner of the screen) to navigate to a specific graphic.

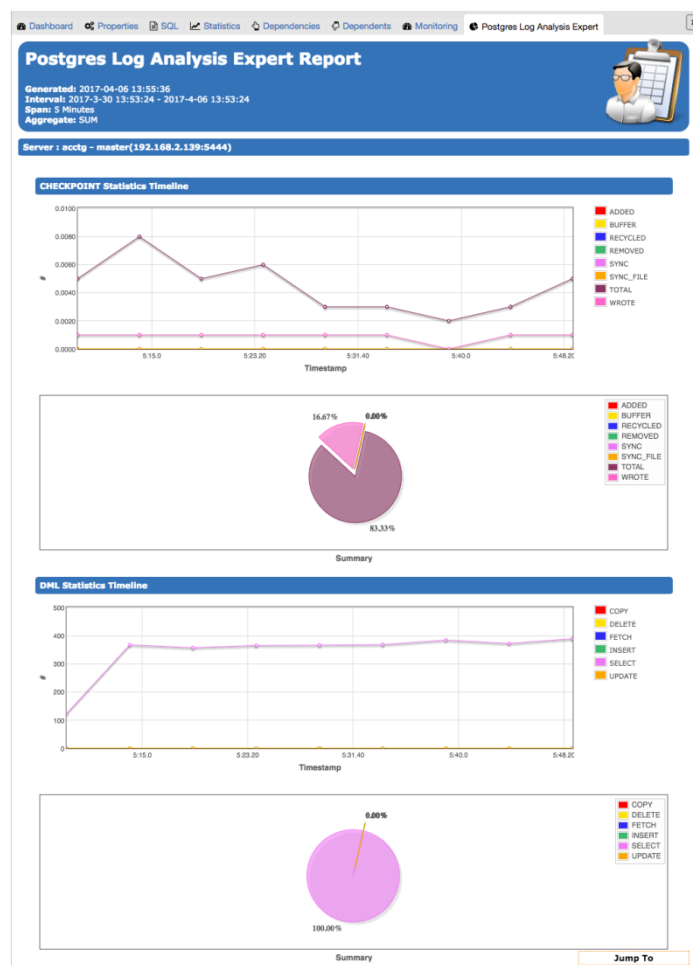


Figure 8.6 - The Postgres Log Analysis Expert Report.

If the report contains an analysis of more than one monitored server, charts and tables will be displayed in sets; first the graphs, tables and charts that display statistics for one server, then the graphics for the next server in the report. For detailed information about each graphic, please see the online documentation.

9 SQL Profiling and Analysis

Most RDBMS experts agree that inefficient SQL code is the leading cause of most database performance problems. The challenge for DBAs and developers is to locate the poorly-running SQL code in large and complex systems, and then optimize that code for better performance.

The SQL Profiler component allows a database superuser to locate and optimize poorly-running SQL code. Users of Microsoft SQL Server's Profiler will find PEM's SQL Profiler very similar in operation and capabilities. SQL Profiler is installed with each Advanced Server instance; if you are using PostgreSQL, you must download the SQL Profiler installer, and install the SQL Profiler product into each managed database instance you wish to profile.

For each database monitored by SQL Profiler, you must:

1. Edit the `postgresql.conf` file; you must include the SQL Profiler library in the `shared_preload_libraries` configuration parameter.

For Linux installations, the parameter value should include:

```
$libdir/sql-profiler
```

on Windows, the parameter value should include:

```
$libdir/sql-profiler.dll
```

2. Create the functions used by SQL Profiler in your database. The SQL Profiler installation program places a SQL script (named `sql-profiler.sql`) in the `share/postgresql/contrib` subdirectory of the main PostgreSQL installation directory on Linux systems. On Windows systems, this script is located in the `share` subdirectory. You must invoke this script on the maintenance database specified when registering the server with PEM.
3. Stop and re-start the server for the changes to take effect.

Please note: if you have connected to the PEM server with the PEM client before configuring SQL Profiler, you must disconnect and reconnect with the server to enable SQL Profiler functionality. For more detailed information about installing and configuring the SQL Profiler plugin, please refer to the PEM Installation Guide, available from the EnterpriseDB website at:

<http://enterprisedb.com/products-services-training/products/documentation>

9.1 Creating a New SQL Trace

SQL Profiler captures and displays a specific SQL workload for analysis in a *SQL trace*. You can start and review captured SQL traces immediately, or save captured traces for review at a later time. You can use SQL Profiler to create and store up to 15 named traces; use menu options to create and manage traces.

9.1.1 Creating a Trace

You can use the `Create trace...` dialog to define a SQL Trace for any database on which SQL Profiler has been installed and configured. To access the dialog, highlight the name of the database in the PEM client tree control; navigate through the `Management` menu to the `SQL Profiler` pull-aside menu, and select `Create trace...`

Figure 9.1 – The Trace options tab.

Use the fields on the `Trace options` tab (see Figure 9.1) to specify details about the new trace:

- Provide a name for the trace in the `Name` field.
- Click in the `User filter` field to specify the roles whose queries will be included the trace; optionally, check the box next to `*Select All*` to include queries from all roles.

- Click in the `Database filter` field to specify which databases to trace; optionally, check the box next to `Select All` to include queries against all databases.
- Specify a trace size in the `Maximum Trace File Size` field; SQL Profiler will terminate the trace when it reaches approximately the size specified.
- Specify `Yes` in the `Run Now` field to start the trace when you select the `Create` button; select `No` to enable fields on the `Schedule` tab.

Figure 9.2 – The Create trace Schedule tab.

Use the fields on the `Schedule` tab (see Figure 9.2) to specify scheduling details for the new trace:

- Use the `Start time` field to specify the starting time for the trace.
- Use the `End time` field to specify the ending time for the trace.
- Specify `Yes` in the `Repeat?` field to indicate that the traces should be repeated every day at the times specified; select `No` to enable fields on the `Periodic job options` tab.

Create trace...

Trace options | **Schedule** | Periodic job options

Days

Week Days

Month Days

Months

Times

Hours

Minutes

Schedules are specified using a **cron-style** format.

- For each selected time or date element, the schedule will execute.
e.g. To execute at 5 minutes past every hour, simply select '05' in the Minutes list box.
- Values from more than one field may be specified in order to further control the schedule.
e.g. To execute at 12:05 and 14:05 every Monday and Thursday, you would click minute 05, hours 12 and 14, and weekdays Monday and Thursday.
- For additional flexibility, the Month Days check list includes an extra Last Day option. This matches the last day of the month, whether it happens to be the 28th, 29th, 30th or 31st.

Please provide valid name for trace

?

Figure 9.3 – The Create trace Periodic job options tab.

Fields on the `Periodic job options` tab (see Figure 9.3) specify scheduling details about a recurring trace. Use fields in the `Days` section to specify the days on which the job will execute:

- Click in the `Week days` field to select the days of the week on which the trace will execute.
- Click in the `Month days` field to select the days of the month on which the trace will execute.
- Click in the `Months` field to select the months in which the trace will execute.

Use fields in the `Times` section to specify a time schedule for the trace execution:

- Click in the `Hours` field to select the hours at which the trace will execute.
- Click in the `Minutes` field to select the hours at which the trace will execute.

When you've completed the `Create trace...` dialog, click `Create` to start the newly defined trace or to schedule the trace for a later time.

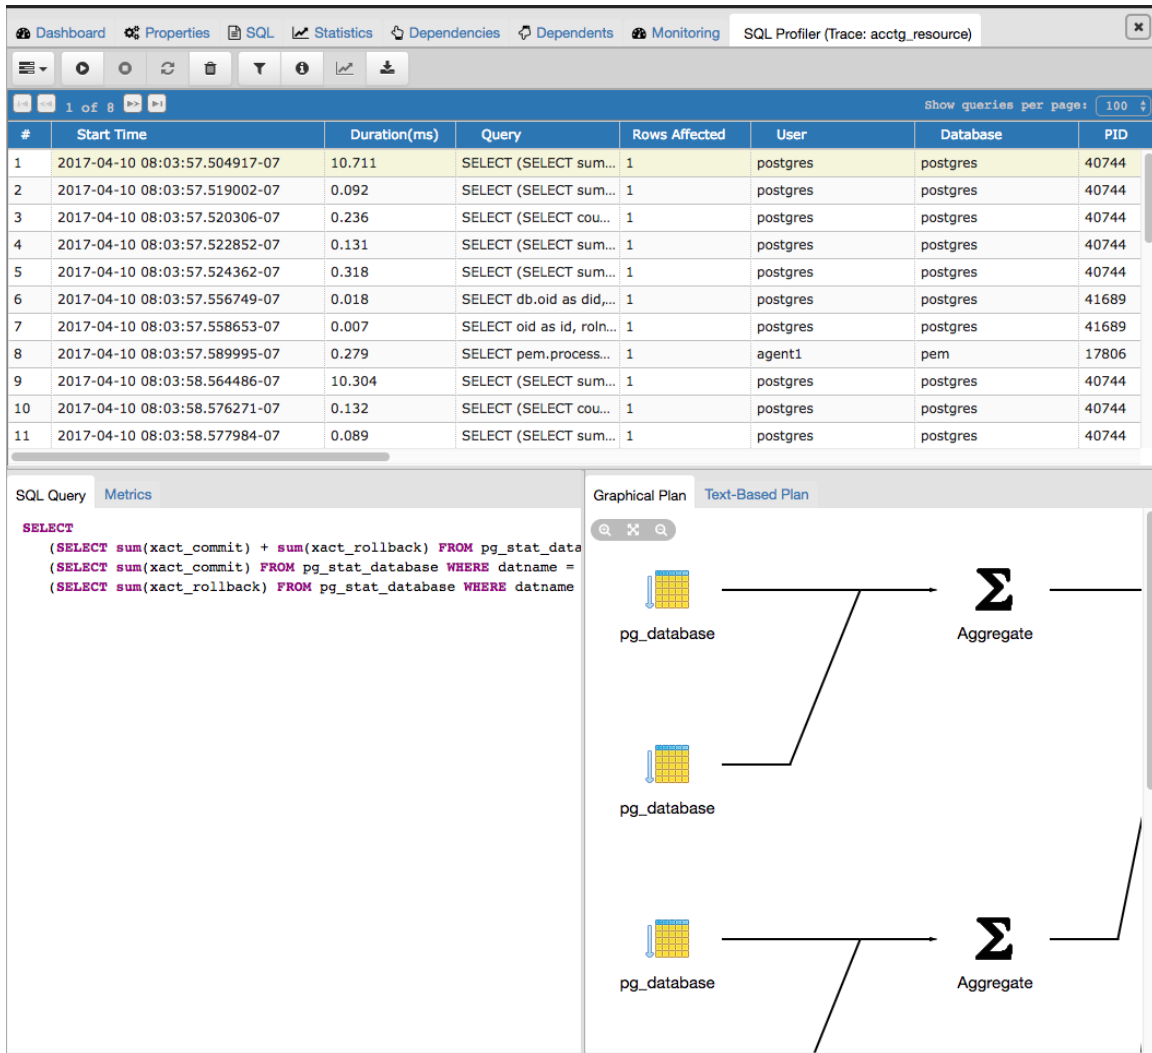


Figure 9.4 – The SQL Profiler tab, displaying the trace results.

If you elect to execute the trace immediately, the trace results will display in the PEM client (see Figure 9.4).

9.1.2 Opening an Existing Trace

To view a previous trace, highlight the name of the profiled database in the PEM client tree control; navigate through the **Management** menu to the **SQL Profiler** pull-aside menu, and select **Open trace....** You can also use the SQL Profiler toolbar menu to open a trace; select the **Open trace...** option. The **Open trace...** dialog opens as shown in Figure 9.5.

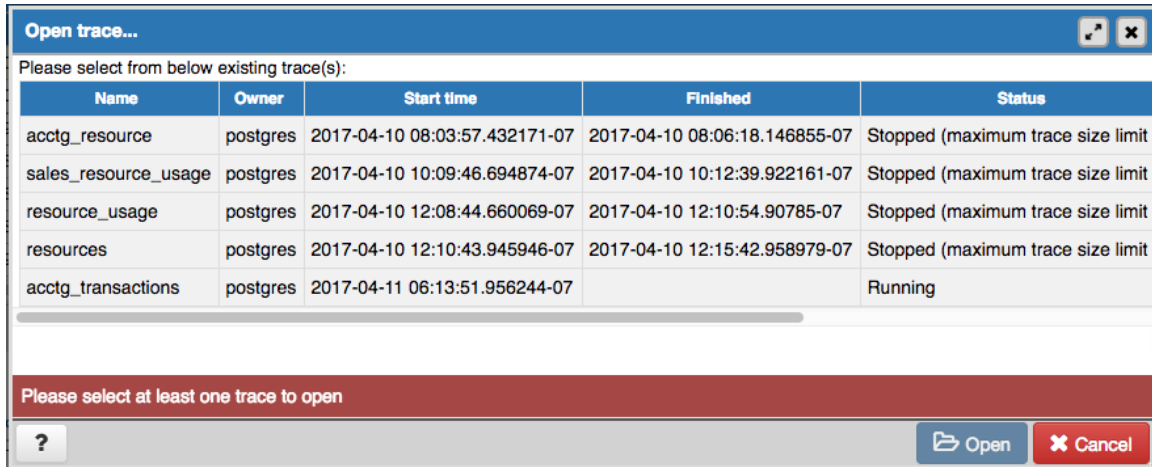


Figure 9.5 – Opening an existing trace.

Highlight an entry in the trace list and click **Open** to open the selected trace. The selected trace opens in the SQL Profiler tab.

9.1.3 Filtering a Trace

A filter is a named set of (one or more) rules, each of which can hide events from the trace view. When you apply a filter to a trace, the hidden events are not removed from the trace, but are merely excluded from the display.

Click the `Filter` icon to open the `Trace Filter` dialog and create a rule (or set of rules) that define a filter. Each rule will screen the events within the current trace based on the identity of the role that invoked the event, or the query type invoked during the event.

To open an existing filter, select the `Open` button; to define a new filter, click the `Add (+)` icon to add a row to the table displayed on the `General` tab and provide rule details:

- Use the `Type` drop-down list box to specify the trace field that the filter rule will apply to.
- Use the `Condition` drop-down list box to specify the type of operator that SQL Profiler will apply to the `Value` when it filters the trace:
 - Select `Matches` to filter events that contain the specified `Value`.
 - Select `Does not match` to filter events that do not contain the specified `Value`.
 - Select `Is equal to` to filter events that contain an exact match to the string specified in the `Value` field.
 - Select `Is not equal to` to filter events that do not contain an exact match to the string specified in the `Value` field.
 - Select `Starts with` to filter events that begin with the string specified in the `Value` field.
 - Select `Does not start with` to filter events that do not begin with the string specified in the `Value` field.
 - Select `Less than` to filter events that have a numeric value less than the number specified in the `Value` field.
 - Select `Greater than` to filter events that have a numeric value greater than the number specified in the `Value` field.

- Select `Less than or equal to` to filter events that have a numeric value less than or equal to the number specified in the `Value` field.
- Select `Greater than or equal to` to filter events that have a numeric value greater than or equal to the number specified in the `Value` field.
- Use the `Value` field to specify the string, number or regular expression that SQL Profiler will search for.

When you've finished defining a rule, click the `Add (+)` icon to add another rule to the filter. To delete a rule from a filter, highlight the rule and click the `Delete` icon.

Click the `Save` button to save the filter definition to a file without applying the filter; to apply the filter, click `OK`. Select `Cancel` to exit the dialog and discard any changes to the filter.

9.1.4 Deleting a Trace

To delete a trace, highlight the name of the profiled database in the PEM client tree control; navigate through the **Management** menu to the **SQL Profiler** pull-aside menu, and select **Delete trace(s) . . .**. You can also use the SQL Profiler toolbar menu to delete a trace; select the **Delete trace(s) . . .** option. The **Delete traces . . .** dialog opens as shown in Figure 9.6.

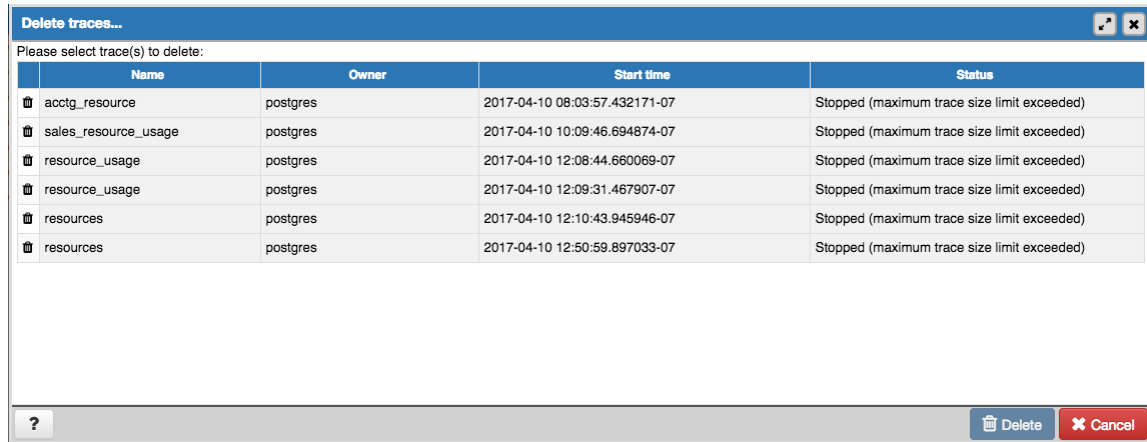


Figure 9.6 – The Delete traces... dialog.

Click the icon to the left of a trace name to mark one or more traces for deletion and click **Delete** (see Figure 9.7).

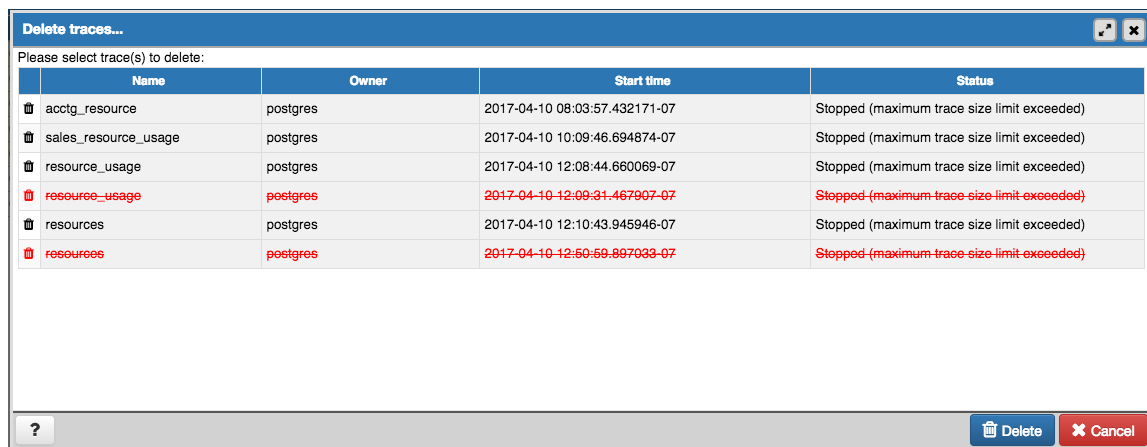


Figure 9.7 – Marking traces for deletion.

The PEM client will acknowledge that the selected traces have been deleted.

9.1.5 Viewing Scheduled Traces

To view a list of scheduled traces, highlight the name of the profiled database in the PEM client tree control; navigate through the **Management** menu to the **SQL Profiler** pull-side menu, and select **Scheduled traces...** You can also use the SQL Profiler toolbar menu to the list; select the **Scheduled traces...** option.

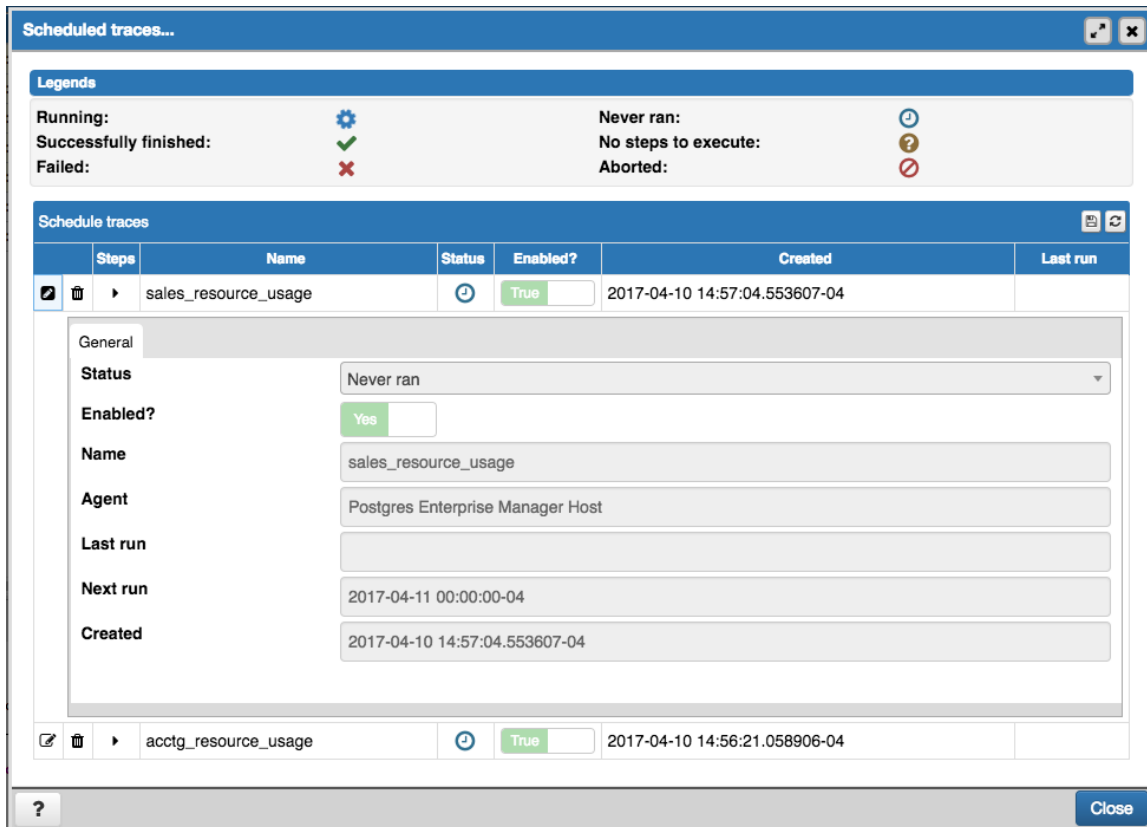


Figure 9.8 – Reviewing scheduled traces.

The **Scheduled traces...** dialog (see Figure 9.8) displays a list of the traces that are awaiting execution. Click the edit button to the left of a trace name to access detailed information about the trace:

- The **Status** field lists the status of the current trace.
- The **Enabled?** switch displays **Yes** if the trace is enabled; **No** if the trace is disabled.
- The **Name** field displays the name of the trace.

- The `Agent` field displays the name of the agent responsible for executing the trace.
- The `Last run` field displays the date and time of the last execution of the trace.
- The `Next run` field displays the date and time of the next scheduled execution of the trace.
- The `Created` field displays the date and time that the trace was defined.

9.2 Using the Index Advisor

Index Advisor is distributed with Advanced Server 9.0 and above. Index Advisor works with SQL Profiler by examining collected SQL statements and making indexing recommendations for any underlying tables to improve SQL response time. The Index Advisor works on all DML (INSERT, UPDATE, DELETE) and SELECT statements that are invoked by a superuser.

Diagnostic output from the Index Advisor includes:

- Forecasted performance benefits from any recommended indexes
- The predicted size of any recommended indexes
- DDL statements you can use to create the recommended indexes

Before using Index Advisor, you must:

1. Modify the postgresql.conf file on each Advanced Server host, adding the index_advisor library to the shared_preload_libraries parameter.
2. Install the Index Advisor contrib module. To install the module, use the psql client or PEM Query Tool to connect to the database, and invoke the following command:

```
\\i <complete_path>/share/contrib/index_advisor.sql
```

3. Restart the server for your changes to take effect.

Index Advisor can make indexing recommendations based on trace data captured by SQL Profiler. Simply highlight one or more queries in the SQL Profiler Trace Data pane, and click the Index Advisor toolbar button (or select Index Advisor from the View menu). For detailed usage information about Index Advisor, please see the EDB Postgres Advanced Server Guide

Please note: Index Advisor cannot analyze statements invoked by a non-superuser. If you attempt to analyze statements invoked by a non-superuser, the server log will include the following error:

```
ERROR: access to library "index_advisor" is not allowed
```

For more information about configuring and using Index Advisor, please see the EDB Postgres Advanced Server Guide, available from EnterpriseDB at:

<https://www.enterprisedb.com/resources/product-documentation>

10 Tuning Wizard

The Tuning Wizard reviews your PostgreSQL or Advanced Server installation, and recommends a set of configuration options that will help tune the installation to best suit its anticipated workload. Please note that benchmarking systems or systems with a high work load may require additional manual tuning to reach optimum performance.

Before using the Tuning Wizard, you must specify the name of the service in the `Service ID` field on the `Advanced` tab of the server's `Properties` dialog. PEM will use the service name when restarting the service after tuning.

The Tuning Wizard can only make recommendations for those servers that reside on the same server as their bound PEM agent. If you have specified a value of `Yes` in the `Remote monitoring` field when defining your server, the server will not be displayed in the Tuning Wizard tree control.

To open the Tuning Wizard, select `Tuning Wizard...` from the `Management` menu of the PEM client. The Tuning Wizard opens, welcoming you (see Figure 10.1).

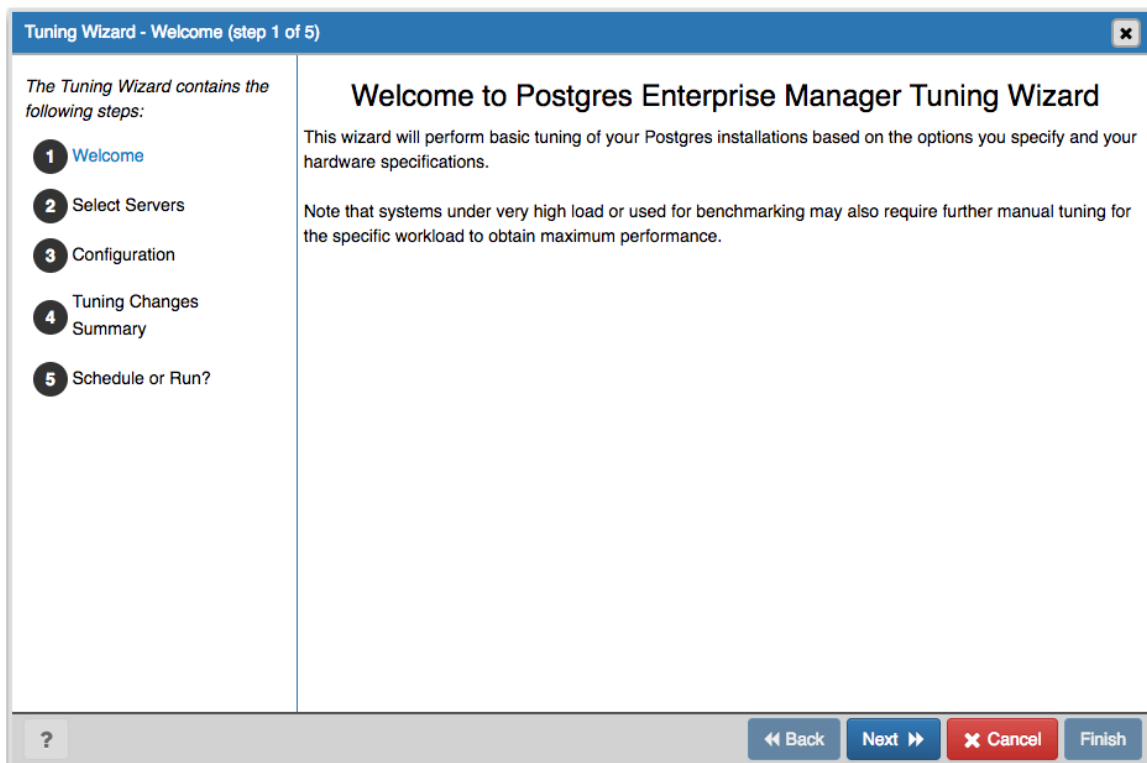


Figure 10.1 – The Tuning Wizard Welcome dialog.

Click `Next` to continue to the server selection dialog (see Figure 10.2).

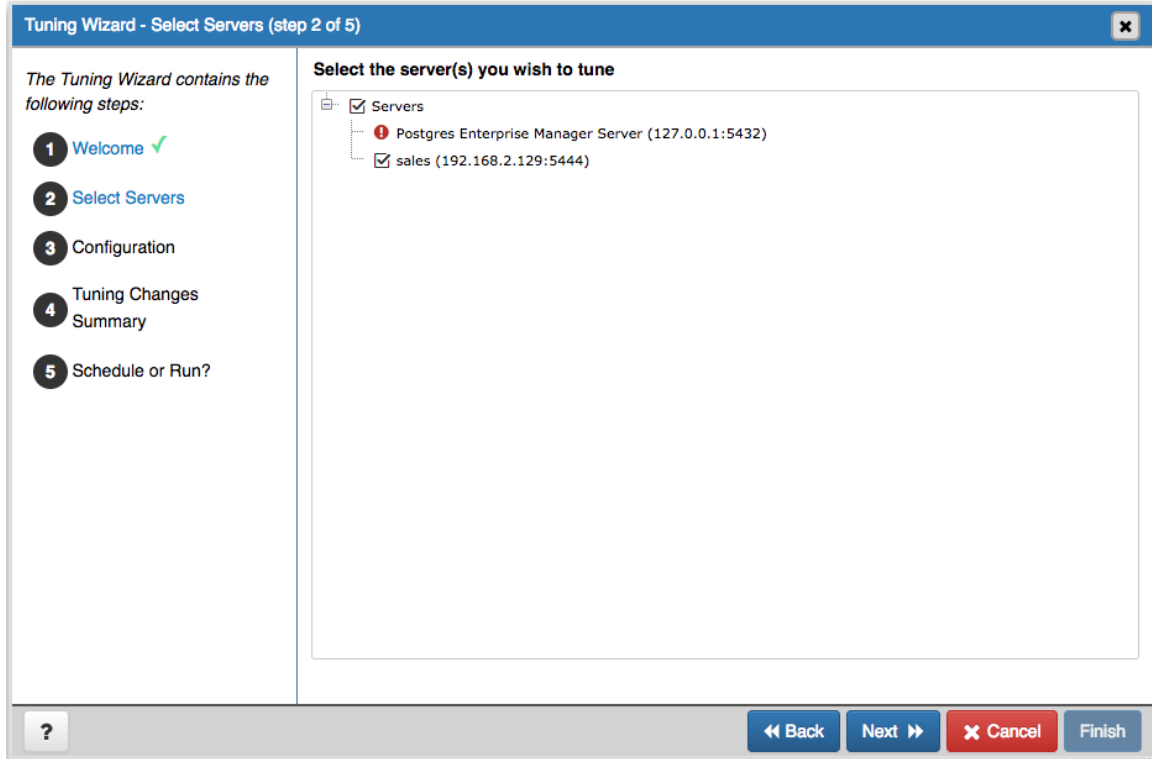


Figure 10.2 – The Select Servers dialog.

Expand the `Servers` node of the tree control to view a list of the servers that are currently monitored by PEM that are available for tuning. Check a box to the left of a server name to select the server for tuning.

Please note: the Tuning Wizard displays a red warning symbol to the left of a server name in the tree control if the service name for that server is not provided on the server's Properties dialog.

Click `Next` to continue to the `Configuration` dialog (see Figure 10.3).

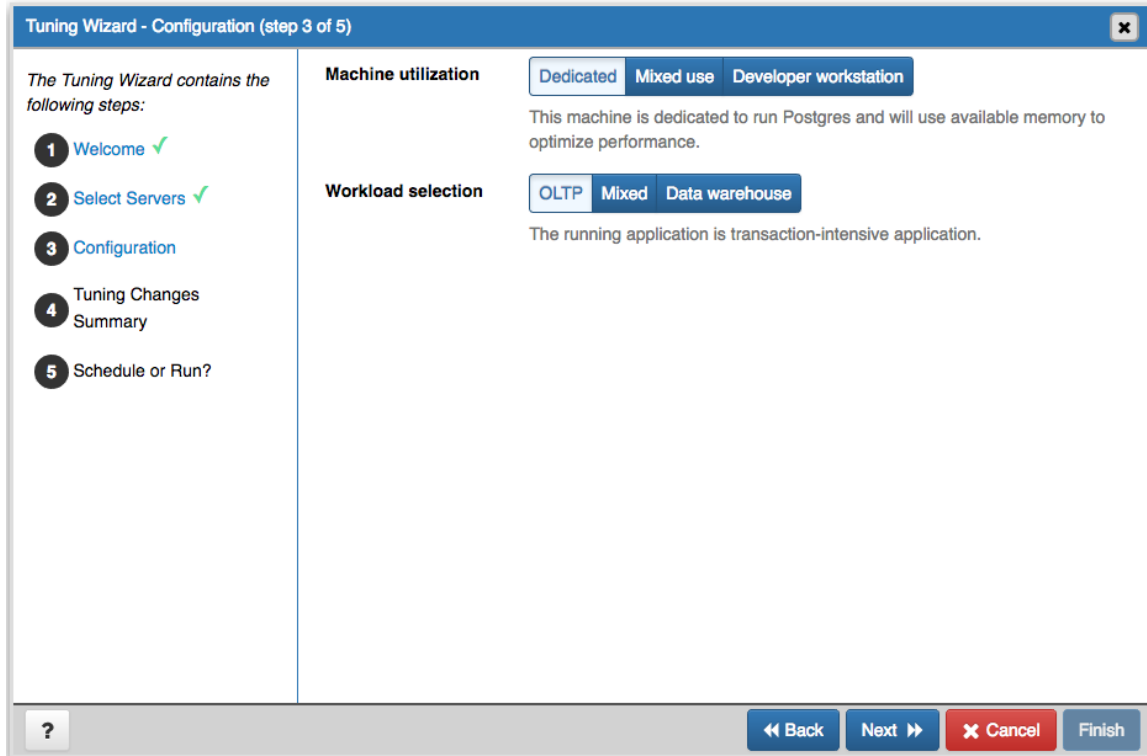


Figure 10.3 – The Configuration dialog.

Select an option in the `Machine utilization` field to specify the type of work performed by the selected servers. The type of work performed by the server determines how the tuning wizard will allocate system resources:

- Select `Dedicated` to dedicate the majority of the system resources to the database server.
- Select `Mixed use` to dedicate a moderate amount of system resources to the database server.
- Select `Developer workstation` to dedicate a relatively small amount of system resources to the database server.

Select an option in the `Workload Selection` field to specify the type of workload typically performed on the selected server:

- Select `OLTP` if the selected server is used primarily to process online transaction workloads.
- Select `Mixed` if the selected server provides a mix of transaction processing and data reporting.

- Select `Data warehouse` if the server is used for heavy data reporting.

Click `Next` to continue to the `Tuning Changes Summary` dialog (see Figure 10.4).

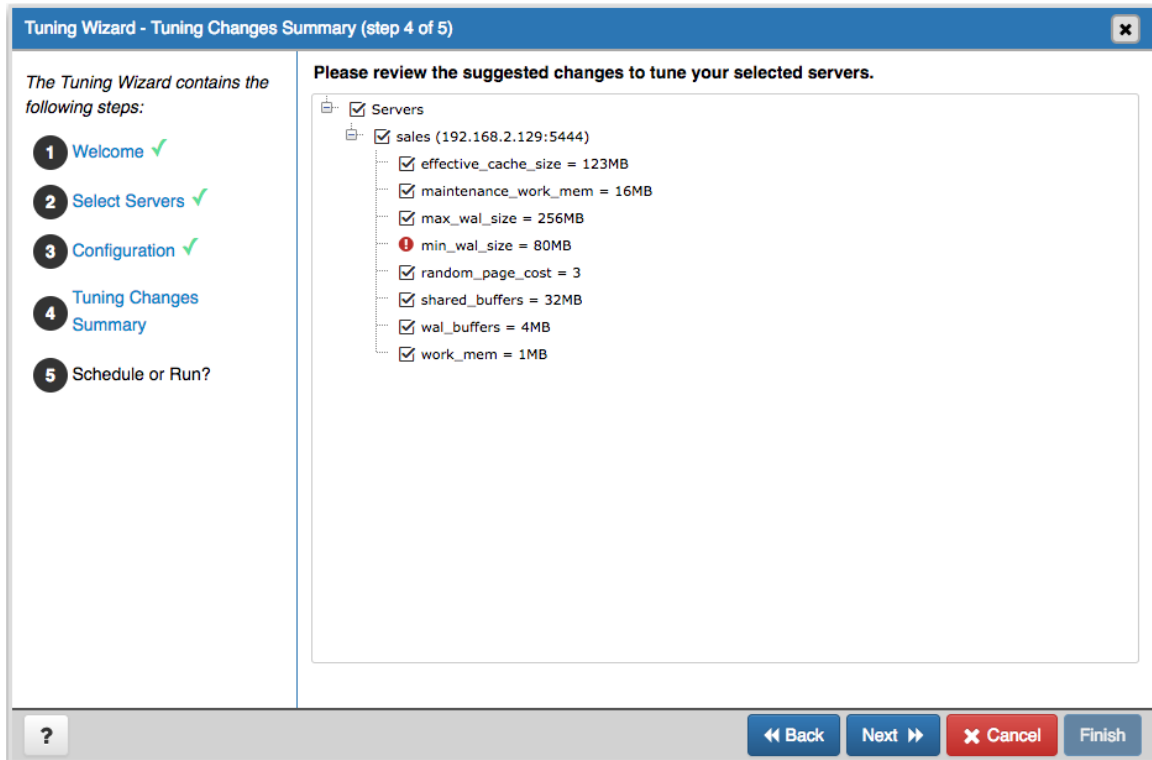


Figure 10.4 – The Tuning Changes Summary dialog.

The tree control on the `Tuning Changes Summary` dialog displays the parameter setting modifications recommended for each server analyzed by the Tuning Wizard. Use the checkboxes next to a server or parameter name to select the recommendations that tuning wizard will either include in a preview report or apply:

- A checked box to the left of a parameter name specifies that the Tuning Wizard will include the parameter setting.
- A checked box to the left of a server name specifies that the Tuning Wizard will include all parameter setting recommendations for the specified server.

Specify which Tuning Wizard recommendations you wish to include in a report or apply, and click `Next` to continue.

Use the `Schedule or Run?` dialog to either specify a time that PEM will apply the changes, or generate a report that details the recommended changes.

The selected actions will apply to all of the changes noted on the `Tuning Changes Summary`. If you opt to generate a report, PEM will create a report that contains a list of the current values and recommended modifications to the configuration parameters selected on the `Tuning Changes Summary` dialog. Note that to implement changes, you will need to invoke the `Tuning Wizard` a second time, specifying the parameters you wish to modify on the `Tuning Changes Summary` dialog.

Select `Schedule changes` to view and specify your scheduling options (see Figure 10.5).

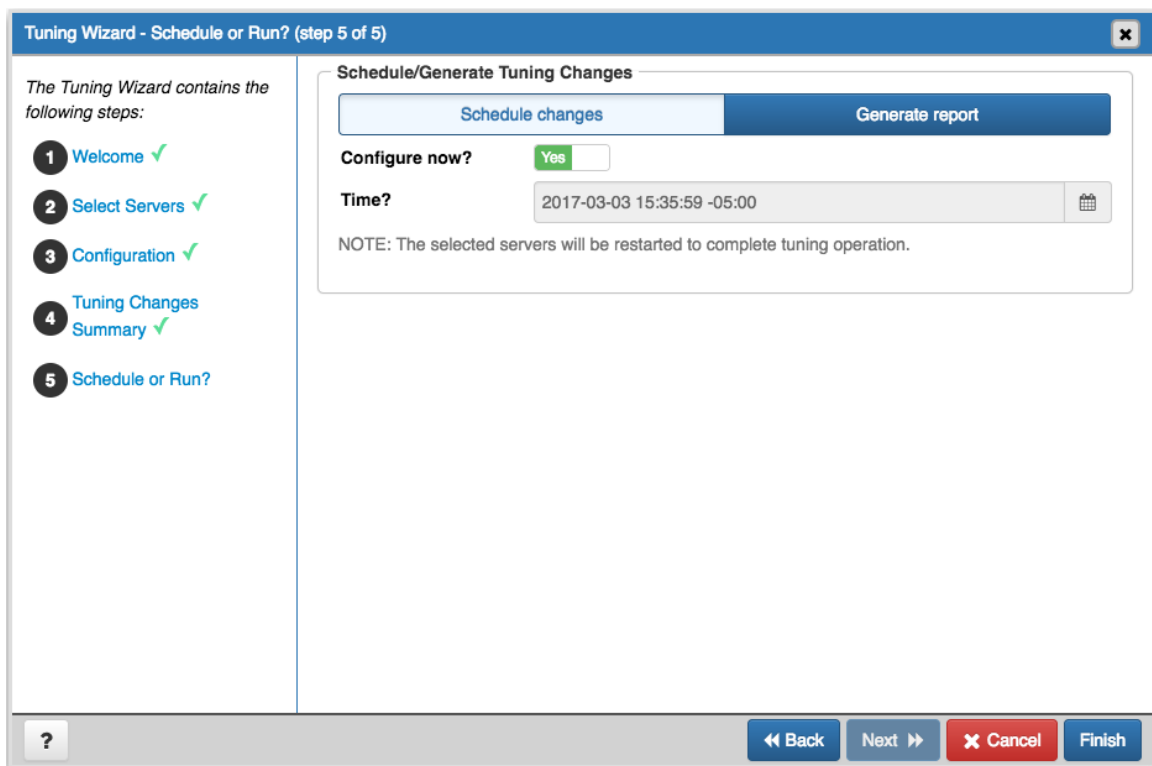


Figure 10.5 – The *Schedule or Run?* dialog.

You can:

- Set the `Configuration now?` slider to `Yes` to apply the tuning wizard's recommendations and restart the server now.
- Set the `Configuration now?` slider to `No` to enable the `Time?` field and use the calendar selector to specify a time for PEM to apply the tuning wizard's recommendations and restart the server. Note that if you schedule a time for the

changes to be applied, you will not be provided with a preview of the change recommendations.

Select `Generate report` to view your report options (see Figure 10.6).

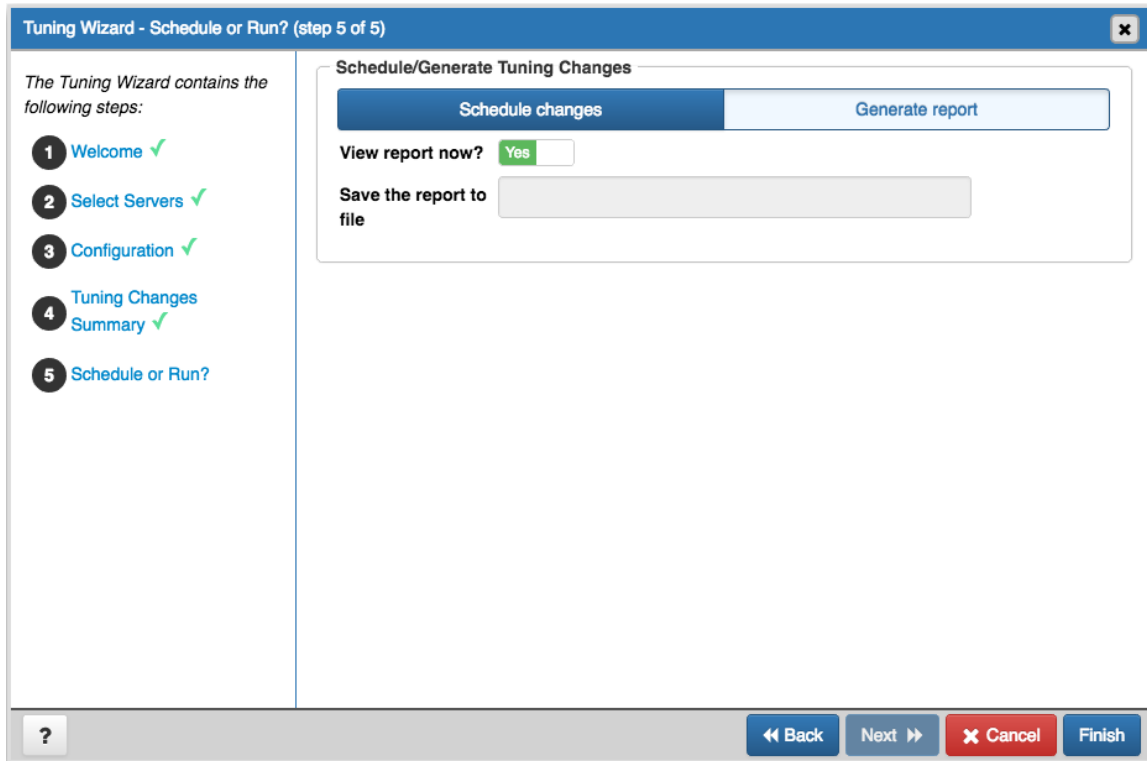


Figure 10.6 – The Schedule or Run? dialog.

You can:

- Set the `View report now?` slider to `Yes` to display the Tuning Wizard report onscreen.
- Set the `View report now?` slider to `No` to enable the `Save the report to file` field and use the calendar selector to specify a file name and location to which PEM will write the Tuning Wizard report.

Click the `Finish` button to either apply the Tuning Wizard's modifications or generate a report (see Figure 10.7) and exit the Tuning Wizard.

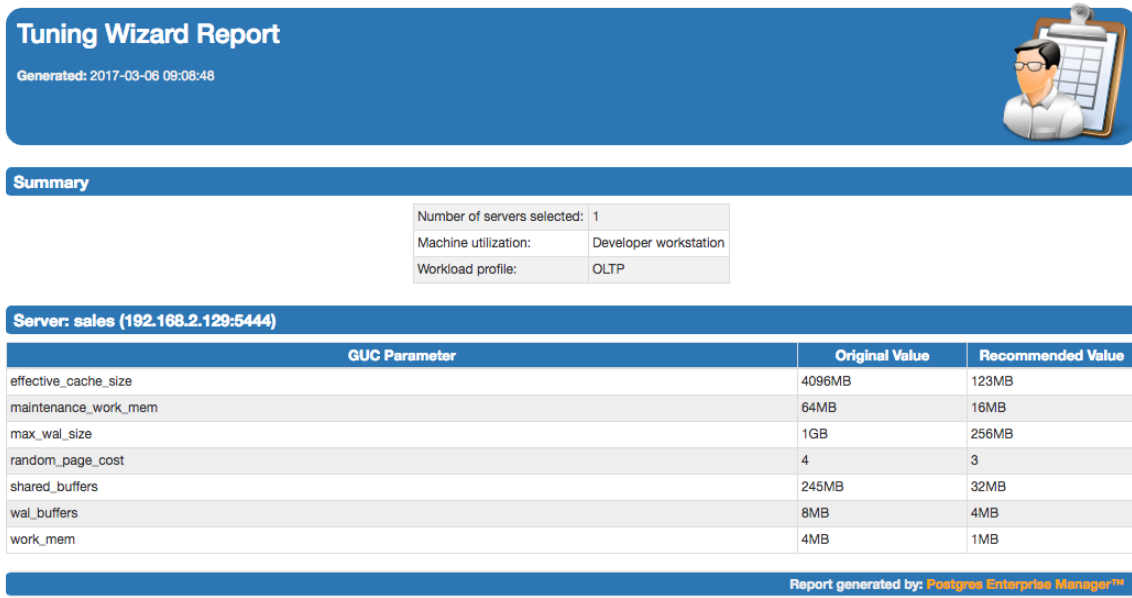


Figure 10.7 – The Tuning Wizard report.

You can confirm that Tuning Wizard has implemented the recommended changes by reviewing the `postgresql.conf` file for the modified server. The Tuning Wizard adds a comment above each modified parameter in the `postgresql.conf` file when the change is applied (see Figure 10.8)

```

root@localhost:/opt/PostgresPlus/9.5AS/data
File Edit View Search Terminal Help
# The value for shared_buffers was changed by the Postgres Enterprise Manager Tuning Wizard at 2017-03-06 09:15:32
#shared_buffers = 196MB
# The value for shared_buffers was changed by the Postgres Enterprise Manager Tuning Wizard at 2017-03-06 09:18:23
shared_buffers = 196MB
# (change requires restart)
# on, off, or try
# (change requires restart)
#huge_pages = try
# temp_buffers = 8MB
# min 800kB
#max_prepared_transactions = 0
# zero disables the feature
# (change requires restart)
# Note: Increasing max_prepared_transactions costs ~600 bytes of shared memory
# per transaction slot, plus lock space (see max_locks_per_transaction).
# It is not advisable to set max_prepared_transactions nonzero unless you
# actively intend to use prepared transactions.
#work_mem = 4MB
# min 64kB
# The value for work_mem was changed by the Postgres Enterprise Manager Tuning Wizard at 2017-03-06 09:15:32
#work_mem = 3MB
# The value for work_mem was changed by the Postgres Enterprise Manager Tuning Wizard at 2017-03-06 09:18:23
work_mem = 3MB
#maintenance_work_mem = 64MB
# min 1MB
# The value for maintenance_work_mem was changed by the Postgres Enterprise Manager Tuning Wizard at 2017-03-06 09:15:32
:maintenance_work_mem = 25MB
--More-- (21%)

```

Figure 10.8 – Confirming a change in the `postgresql.conf` file.

You can also confirm a parameter value by querying the server. For example, to confirm the value of the `shared_buffers` parameter, open a SQL command line using either the Query Tool (accessed through the Tools menu) or the psql client, and issue the command:

```
SHOW shared_buffers;
```

The value returned by the server will confirm that the parameter has been modified.

11 Postgres Expert - Best Practice Enforcement

The Postgres Expert utility provides expert advice on how to best configure your Postgres servers for optimal performance, security, and more. Postgres Expert serves as a PostgreSQL 'DBA in a box' by analyzing your servers for deviations in best practices. Postgres Expert contains three specialized Experts:

- The Configuration Expert.
- The Schema Expert.
- The Security Expert.

You can select specific rules for each Expert to analyze, or accept all rules, and then review Postgres Expert reports detailing any best practice issues that require your attention.

11.1 Using the Postgres Expert Wizard

To use the Postgres Expert wizard select the `Postgres Expert` option from the `Management` menu in the PEM client. When the wizard's `Welcome` window opens (see Figure 11.1), click `Next` to continue.

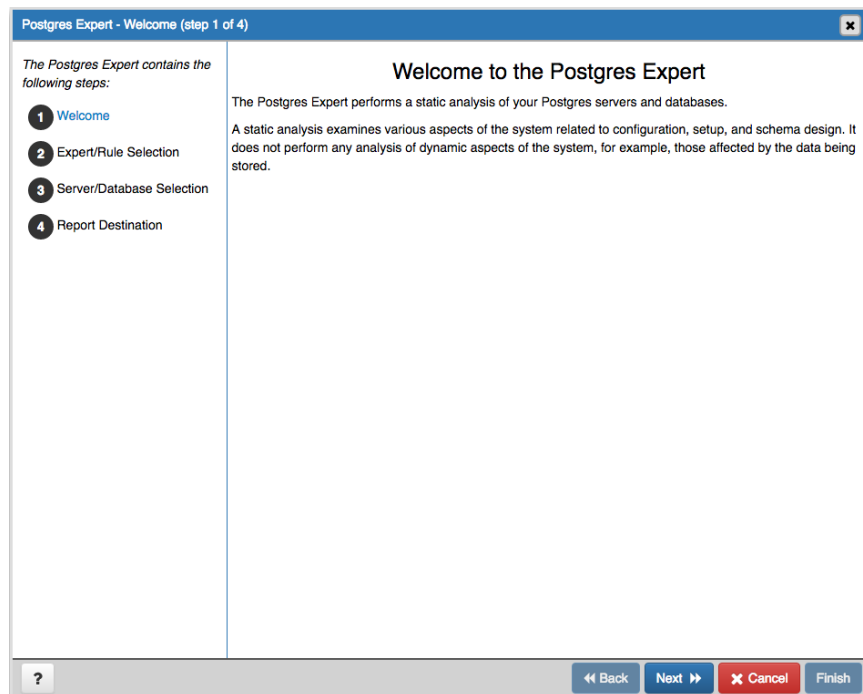


Figure 11.1 - The Postgres Expert Welcome dialog.

The wizard displays a tree control that allows you to choose the `Experts` and `Rules` with which Postgres Expert will evaluate the specified server or database (see Figure 11.2).

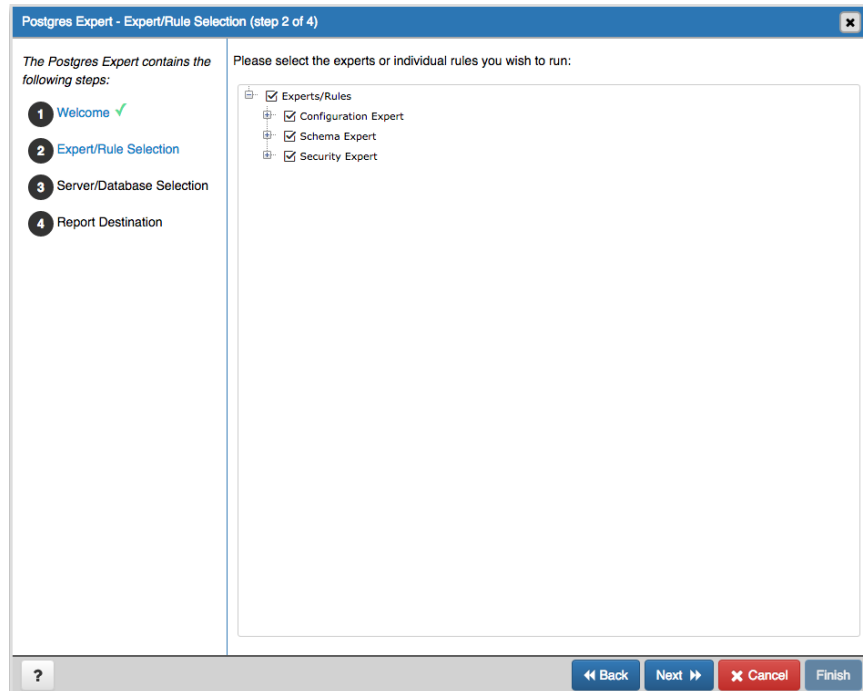


Figure 11.2 - The Experts/Rules dialog.

The tree control categorizes the rules under three `Expert` headings:

- Select from the `Configuration Expert` rules to analyze the parameter settings of the server or operating system to find any adjustments that might improve system performance.
- Select from the `Schema Expert` rules to analyze schema objects (locating missing primary keys, foreign keys without indexes, etc).
- Select from the `Security Expert` rules to review the system to find security vulnerabilities.

Use the checkmark indicator to the left of an expert or rule to indicate that the Postgres Expert should analyze the configuration of the selected servers for any best practice deviations related to the checked item.

You can:

- Check the box next to the name of an expert to select or deselect all of the configuration items listed under that node of the tree control.
- Check the box next to `Servers/Databases` to instruct Postgres Expert to review the selected server for all of the items in the tree control.
- Deselect the box next to `Servers/Databases` to un-check all of the rules; then, navigate through the tree control, specifying only the items that you wish Postgres Expert to evaluate.

After making your selections, click `Next` to continue to the `Server/Databases` tree control (see Figure 11.3).

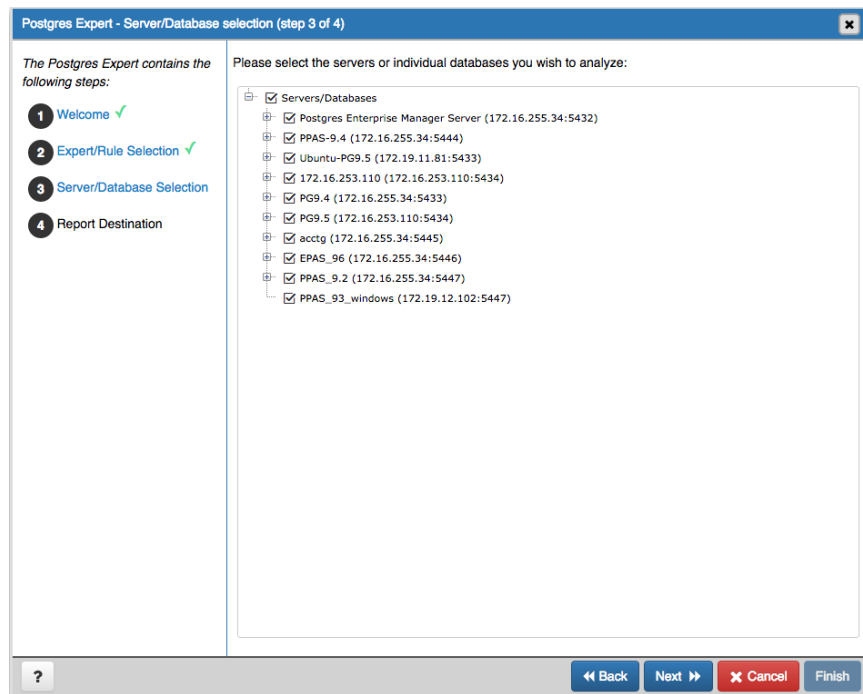


Figure 11.3 - The Servers/Databases dialog.

Select or de-select the servers and databases that you would like Postgres Expert to analyze. If you select multiple servers or databases, the resulting report will contain a separate analysis of each target. When you've finished, click `Next` to select a report destination (see Figure 11.4).

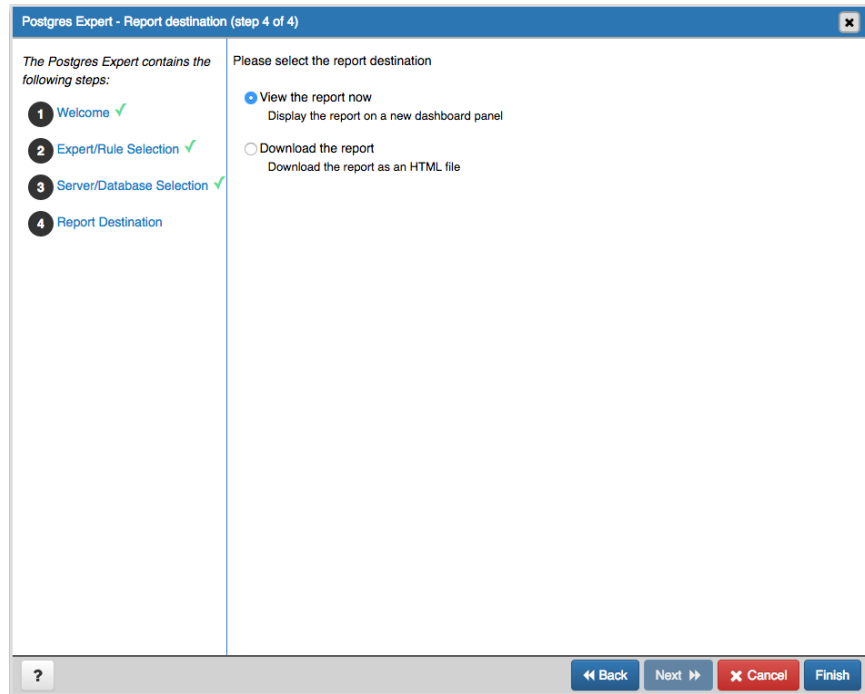


Figure 11.4 - Specify a report destination.

You can select the default option and click **Finish** to view an onscreen report from Postgres Expert, or check the box next to **Download the report** to save a copy of the report to an HTML file for later use. If you choose to save the report to a file, the download will begin immediately. The file will be saved in your default download directory.

11.2 Reviewing Postgres Expert Recommendations

Postgres Expert produces an easily navigated report that contains an analysis of the selected rules, categorized by high, medium, and low severities, for the selected servers (see Figure 11.5).

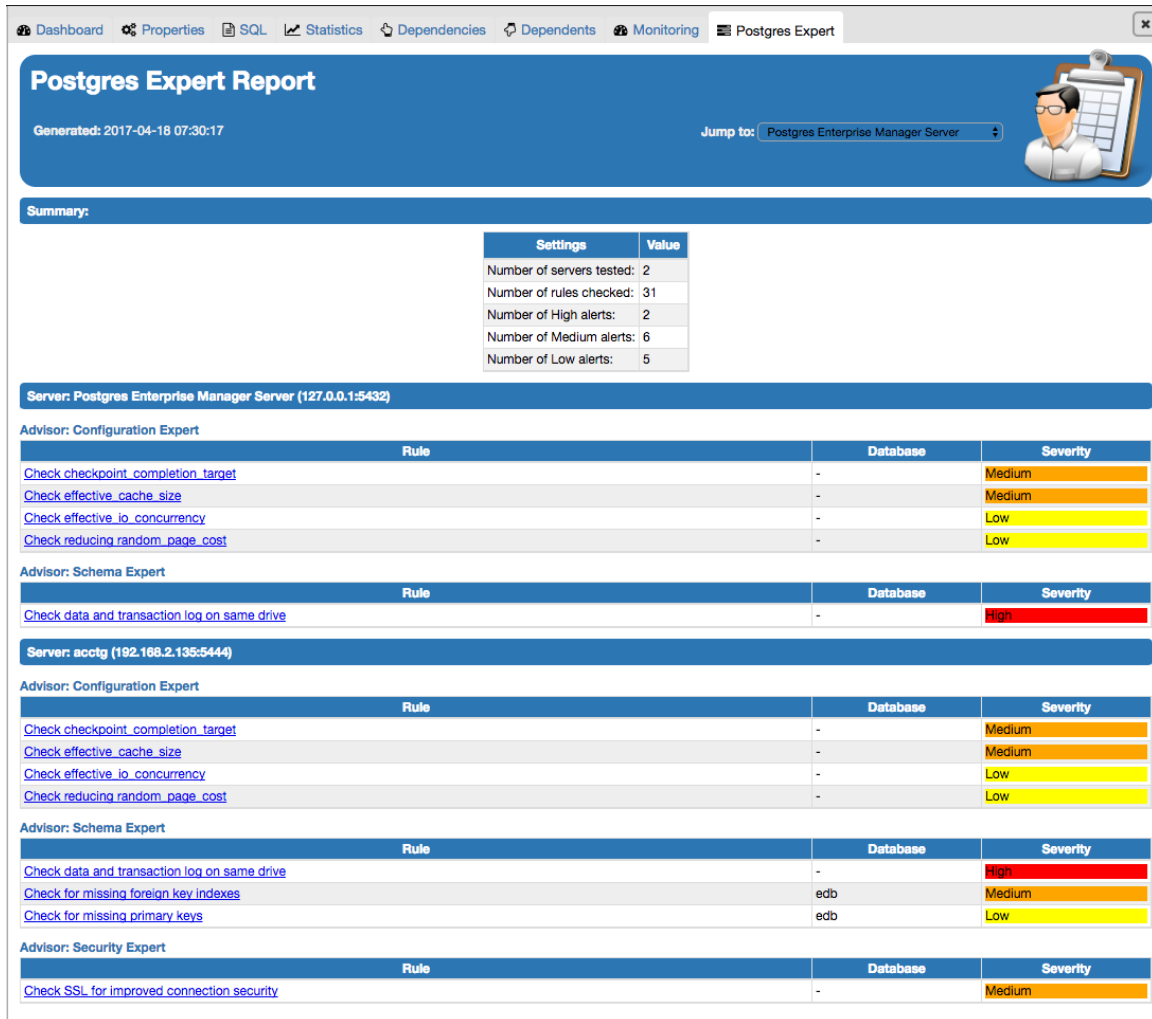


Figure 11.5 - The Postgres Expert report.

The report header contains a summary of the report, and includes the date and time that the report was generated, the number of rules analyzed, and the number of deviations from best practices found by Postgres Expert. Use the **Jump to** drop-down listbox to select a server to navigate to the section of the report that targets recommendations for that server.

The body of the report contains the detailed findings for each server selected for analysis. The findings are sorted by **Expert**; within each **Expert** heading, any rule violations are ranked by **Severity**.

Server: acctg (192.168.2.135:5444)

Advisor: Configuration Expert

Rule	Database	Severity
Check checkpoint_completion_target	-	Medium
Check effective_cache_size	-	Medium
Check effective_io_concurrency	-	Low
Check reducing_random_page_cost	-	Low

Advisor: Schema Expert

Rule	Database	Severity
Check data and transaction log on same drive	-	High

Trigger: data directory and transaction log directory share a device

Recommended Value: Avoid using the same storage device for the data directory and transaction logs

Description: The database server must write any changes to the data stored first to the transaction log, a sequential log of all changes to be made to the database files, and then to the database files themselves. On busy servers, significant performance gains may be seen when separating the data directory and transaction log directory onto different physical storage devices.

Current Values:

Setting	Value
data_dir	/opt/edb/as9.6/data
xlog_dir	/opt/edb/as9.6/data/pg_xlog

Check for missing foreign key indexes	edb	Medium
Check for missing primary keys	edb	Low

Figure 11.6 - The detailed recommendation for a rule.

Click on each rule in the Postgres Expert report to display details and recommendations for that rule (see Figure 11.6). Within each rule, section headings display:

- The **Advisor** section lists the name of the Postgres Expert advisor that prompted the recommendation.
- The **Trigger** section displays a description of the rule that raised the alert.
- The **Recommended Value** section displays the value to which Postgres Expert recommends setting the selected parameter.
- The **Description** section displays information and advice about the parameter that caused the alert.
- The **Current Values** section displays the current value(s) of any parameter(s) that influence the Postgres Expert's evaluation.

12 Configuring Streaming Replication

The PEM Streaming Replication Wizard walks you through the process of creating or modifying a streaming replication scenario. You can use the wizard to:

- Install new servers to act as master and standby nodes in a replication scenario.
- Configure existing servers in the roles of master and standby nodes in a replication scenario.
- Add new or existing standby servers to an existing replication scenario.

If you are configuring replication using an existing server as the master node or as a standby node within the replication scenario, the servers must have been installed with the graphical installer. The Streaming Replication wizard does not support pre-existing servers installed via RPM packages at this time.

The Streaming Replication wizard is supported by PEM agent version 6.0 (or later). Each node of a replication scenario must have a resident PEM agent; remote monitoring of master or standby nodes is not supported at this time. After installing the PEM agent, you must:

- on a Linux host, modify the PEM agent configuration file (`agent.cfg`) located in `/opt/edb/pem/agent/etc/agent.cfg` setting the following parameters to `true`:

```
allow_package_management
allow_server_restart
allow_streaming_replication
```

- on a Windows host, modify the Windows registry (`HKEY_LOCAL_MACHINE\Software\Wow6432Node\EnterpriseDB\PEM\agent`), setting the following entries to `true`:

```
AllowPackageManagement
AllowServerRestart
AllowStreamingReplication
```

After updating the configuration file or registry, restart the PEM agent service:

- on a Linux host, open a command line, assume superuser privileges and enter the command `/etc/init.d/pemagent restart` (on RHEL or CentOS 6.x) or `systemctl pemagent restart` (on RHEL or CentOS 7.x).

- on a Windows host, use the `Services` applet to restart the `Postgres Enterprise Manager - pemAgent` service.

Then, to open the Streaming Replication wizard, select `Streaming Replication` from the `Management` menu. The Streaming Replication wizard welcomes you as shown in Figure 12.1.

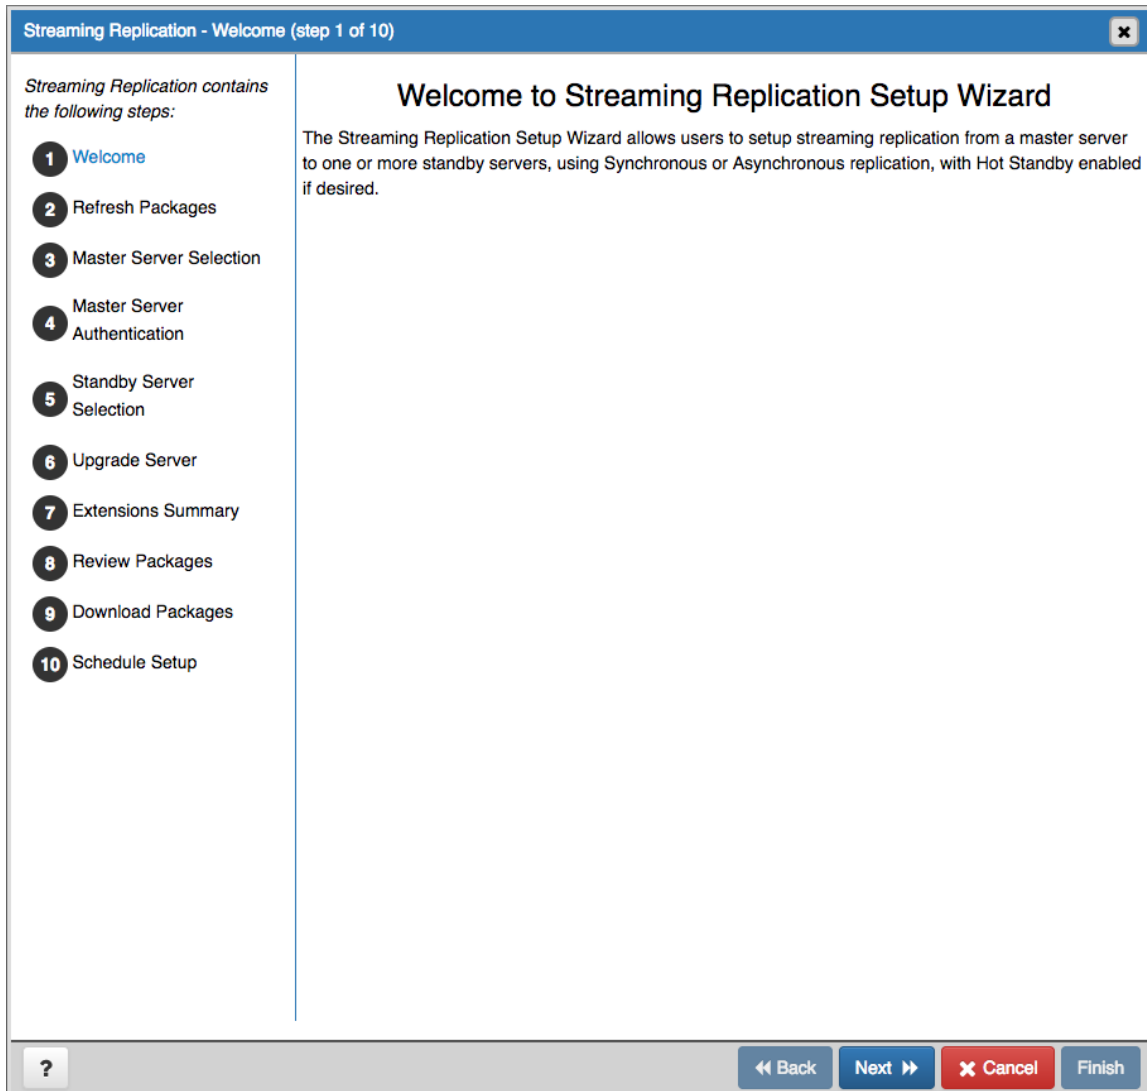


Figure 12.1– The Streaming Replication wizard's Welcome window.

Click `Next` to continue.

The `Refresh Packages` dialog opens (see Figure 12.2), offering you the option to refresh the package data that is stored on the PEM server about the currently installed packages.

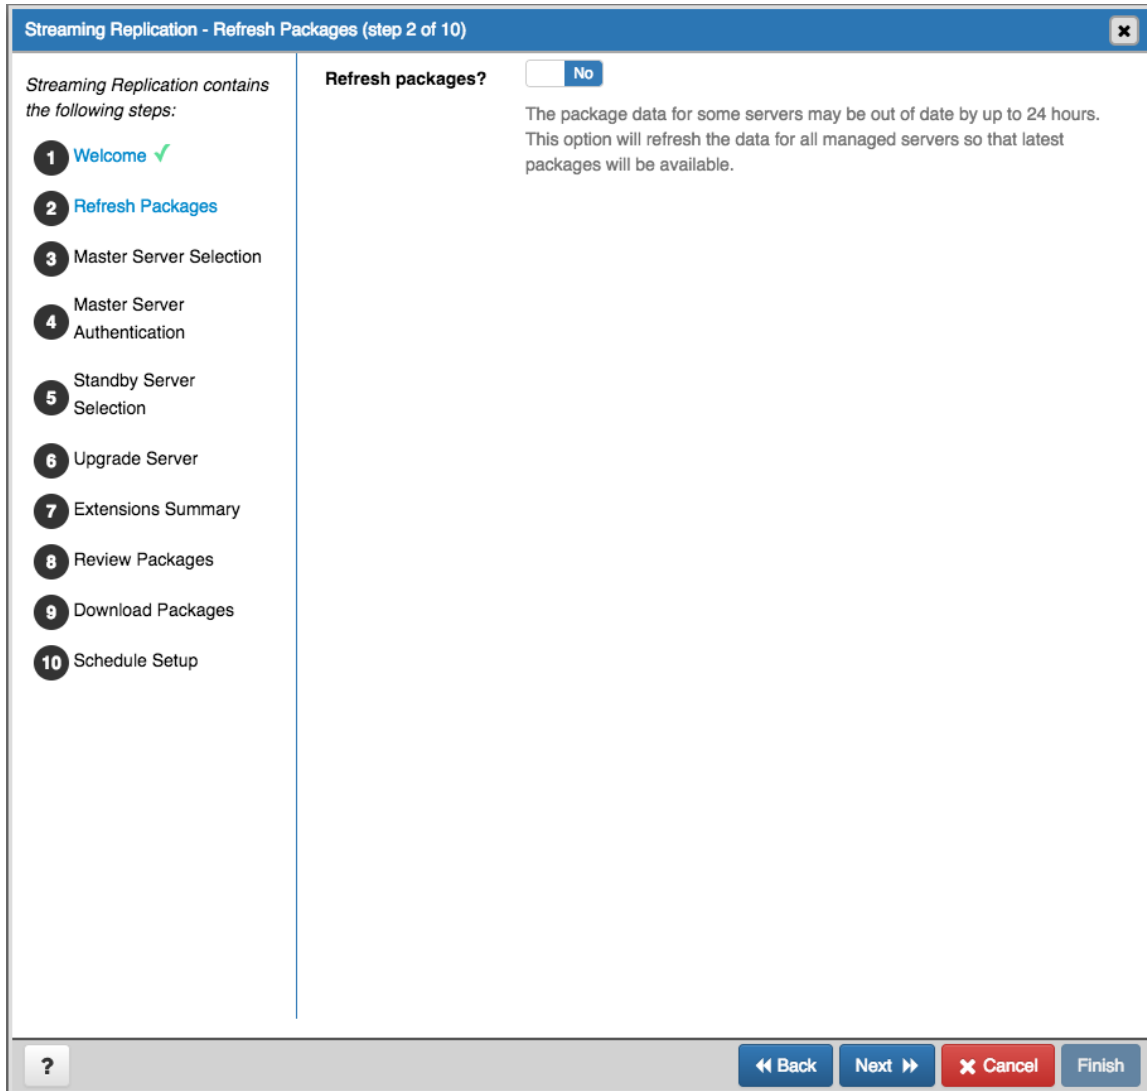


Figure 12.2 – Use the popup to refresh package data.

The `PEM installed_packages` probe retrieves information about the currently installed packages that reside on hosts that are monitored by PEM agents. Select `Yes` to invoke the probe and update the information that is stored on the PEM server. If you have not added servers to the monitored hosts since the last probe execution (by default, the `installed_packages` probe executes once every 24 hours), click `No` to continue without executing the probe.

Click `Next` to continue.

Streaming Replication - Master Server Selection (step 3 of 10)

Streaming Replication contains the following steps:

- 1 Welcome ✓
- 2 Refresh Packages ✓
- 3 Master Server Selection
- 4 Master Server Authentication
- 5 Standby Server Selection
- 6 Upgrade Server
- 7 Extensions Summary
- 8 Review Packages
- 9 Download Packages
- 10 Schedule Setup

Select the master host, and the desired database server which will be installed if required.

Host
 Select name of the PEM agent for the master node from the list of active agents.

Database server
 Select the database server of the master node.

Replication host address
 Specify the address of the host on which the master node will reside.

Replication slot name (optional)
 The PEM server will create the replication slot on the master node and add entries to the recovery.conf files on standby nodes. This option is supported only on server versions 9.4 and above.

Note: Hosts may not be available for selection if they do not have "allow_streaming_replication", "allow_package_management" and "allow_server_restart" enabled in their PEM Agent configuration file.

◀ Back Next ▶ Cancel Finish

Figure 12.3 - Specify information about the master node.

Fields on the master node selection dialog (see Figure 12.3) prompt you to provide information about the master node of the streaming replication scenario:

- Use the drop-down listbox in the `Host` field to select the name of the PEM agent that monitors the master node from the list of active agents. To be displayed in the listbox, the agent must be configured with `allow_streaming_replication`, `allow_package_management`, and `allow_server_restart` enabled (set to `true`) in the PEM Agent configuration file. Please note that each node of a replication scenario must have a resident agent; remote monitoring of replication nodes is not supported.
- Use the drop-down listbox in the `Database server` field to specify the server or server version of the master node. You can select:

- A previously installed server to act as the master node of the replication scenario; existing servers include the word (Installed) in their description. When you select an existing server, the `Validate` button will be enabled.
- The server version of the new master node that PEM will install when configuring the streaming replication scenario. To create a new server, select a server version that does not include the word (Installed) in the description.
- Use the drop-down listbox in the `Replication host address` field to select the address of the host on which the master node will reside.
- Optionally, provide a name for a replication slot in the `Replication slot name` field; the PEM server will create the replication slot on the master node, and add entries to the `recovery.conf` files on standby nodes. A replication slot name can contain lower-case letters, numbers, and the underscore character. This option is valid only for database server versions 9.4 and above.

For more information about replication slots, see the PostgreSQL Core documentation, available at:

<http://www.postgresql.org/docs/10/static/warm-standby.html#STREAMING-REPLICATION-SLOTS>

Click `Next` to continue.

Streaming Replication - Master Server Authentication (step 4 of 10)

Streaming Replication contains the following steps:

- 1 Welcome ✓
- 2 Refresh Packages ✓
- 3 Master Server Selection ✓
- 4 Master Server Authentication
- 5 Standby Server Selection
- 6 Upgrade Server
- 7 Extensions Summary
- 8 Review Packages
- 9 Download Packages
- 10 Schedule Setup

Replication user name
Specify the role that is either a database superuser or has REPLICATION privileges.

Replication password
Specify the password that will be associated with the replication user.

Database user name
Use the Database user name field to specify the name of a database superuser on the master node.

Database password
Use the Database password field to specify the password associated with the database superuser.

Validate

Note: Validate button will be enabled only when database server is already installed on Master node. Click the Validate button to authenticate replication/database server credential.

Navigation: ? Back Next Cancel Finish

Figure 12.4 - Specify information about the master node.

Use the Master Server Authentication dialog (see Figure 12.4) to provide authentication information for the master node.

- Use the `Replication user name` field to specify the name of an existing role that is either a database superuser or has `REPLICATION` privileges, or the name of a role that will be created by PEM for use during replication-related transactions. Please note that PEM will return an error if you specify the name of an existing user with insufficient privileges.
- Use the `Replication password` field to specify the password that will be associated with the replication user.

- Use the `Database user name` field to specify the name of a database superuser on the master node.
- Use the `Database password` field to specify the password associated with the database superuser.

If you are using an existing server as the master node of the replication scenario, you must use the `Validate` button to confirm that the connection information provided. When you press the `Validate` button, the server will attempt to connect with the credentials supplied.

Click `Next` to continue.

Streaming Replication - Standby Server Selection (step 5 of 10)

Streaming Replication contains the following steps:

- Welcome ✓
- Refresh Packages ✓
- Master Server Selection ✓
- Master Server Authentication ✓
- Standby Server Selection
- Upgrade Server
- Extensions Summary
- Review Packages
- Download Packages
- Schedule Setup

Specify hosts on which standby servers will be configured.

Standby Servers		
Agent	IP address	Hot standby?
No standby server(s) found.		



Note: Hosts may not be available for selection as Standby server if the operating system differs from that of the master. If standby server already has a database server installed then in order to synchronize with the primary server, a new data directory will be created by PEM and will be populated with a copy of the primary database. This is done to protect the existing data on the selected standby node.

?
◀ Back
Next ▶▶
✖ Cancel
Finish

Figure 12.5 - Select the standby servers.

Use the table on the Standby Server Selection dialog (see Figure 12.5) to provide properties of one or more standby nodes. Click the add icon (+) to add an entry to the table.

Specify hosts on which standby servers will be configured.

Standby Servers			
	Agent	IP address	Hot standby?
 	Standby Agent	192.168.2.164	<input type="checkbox"/> No

Options

Agent Standby Agent

Use the Agent drop-down to select the name of the agent that will monitor a standby node in the replication scenario. Please note that you will not be able to edit the properties of a standby node that is already part of a replication scenario.

IP address 192.168.2.164

Use the IP address drop-down listbox to select the IP address of the standby node.

Hot standby? ☐ No

Use Hot standby if the standby node should be used for read-only queries while acting as a standby node in the replication scenario.

Synchronous? ☐ No

Use Synchronous to enable synchronous replication; streaming replication is asynchronous by default. If a standby node is specified as Synchronous, a transaction will not be committed until it is written to the transaction log of both the master node and standby node. Data loss is less-likely in the event of a server failure of a node of a synchronous replication scenario, but will increase the processing time of each transaction.

Priority 1

Use the Priority to specify the order in which the standby nodes will be listed in the postgresql.conf file of the master node. For example, select 1 to indicate that in the standby should be listed first, 2 to indicate that the node should be listed second, etc.

Replication state

Application name

Note: Hosts may not be available for selection as Standby server if the operating system differs from that of the master. If standby server already has a database server installed then in order to synchronize with the primary server, a new data directory will be created by PEM and will be populated with a copy of the primary database. This is done to protect the existing data on the selected standby node.

Figure 12.6 - Select the standby servers.

Use the Standby Server Options form (see Figure 12.6) to provide information about the standby node:

- Use the Agent drop-down listbox to select the name of the agent that resides on the standby node in the replication scenario. Please note that you will not be able to edit the properties of a standby node that is already part of a replication scenario.

- Use the `IP address` drop-down listbox to select the IP address of the standby node.
- Set the `Hot standby` field to `Yes` if the standby node should be used for read-only queries while acting as a standby node in the replication scenario.
- Set the `Synchronous?` field to `Yes` to enable synchronous replication; streaming replication is asynchronous by default. If a standby node is specified as `Synchronous`, a transaction will not be committed until it is written to the transaction log of both the master node and standby node.

Data loss is less-likely in a synchronous replication scenario should a failover occur, but using synchronous replication increases the processing time of each transaction.

- Use the `Priority` drop-down listbox to specify the order in which the standby nodes will be listed in the `postgresql.conf` file of the master node. For example, select `1` to indicate that in the standby should be listed first, `2` to indicate that the node should be listed second, etc.

If you are adding the standby to an existing replication scenario, PEM will display the identity of the replication master in the `Replication state` field, and the name of the application (from the `pg_stat_replication` table) in the `Application name` field. These values are not user-modifiable.

Click the `Add/Change` button to add another standby node to the list of servers, or the `Edit` icon to modify the values associated with a server in the list. Use the `Delete` icon to remove a standby definition from the table. When you've finished defining the standby nodes, click `Next`.

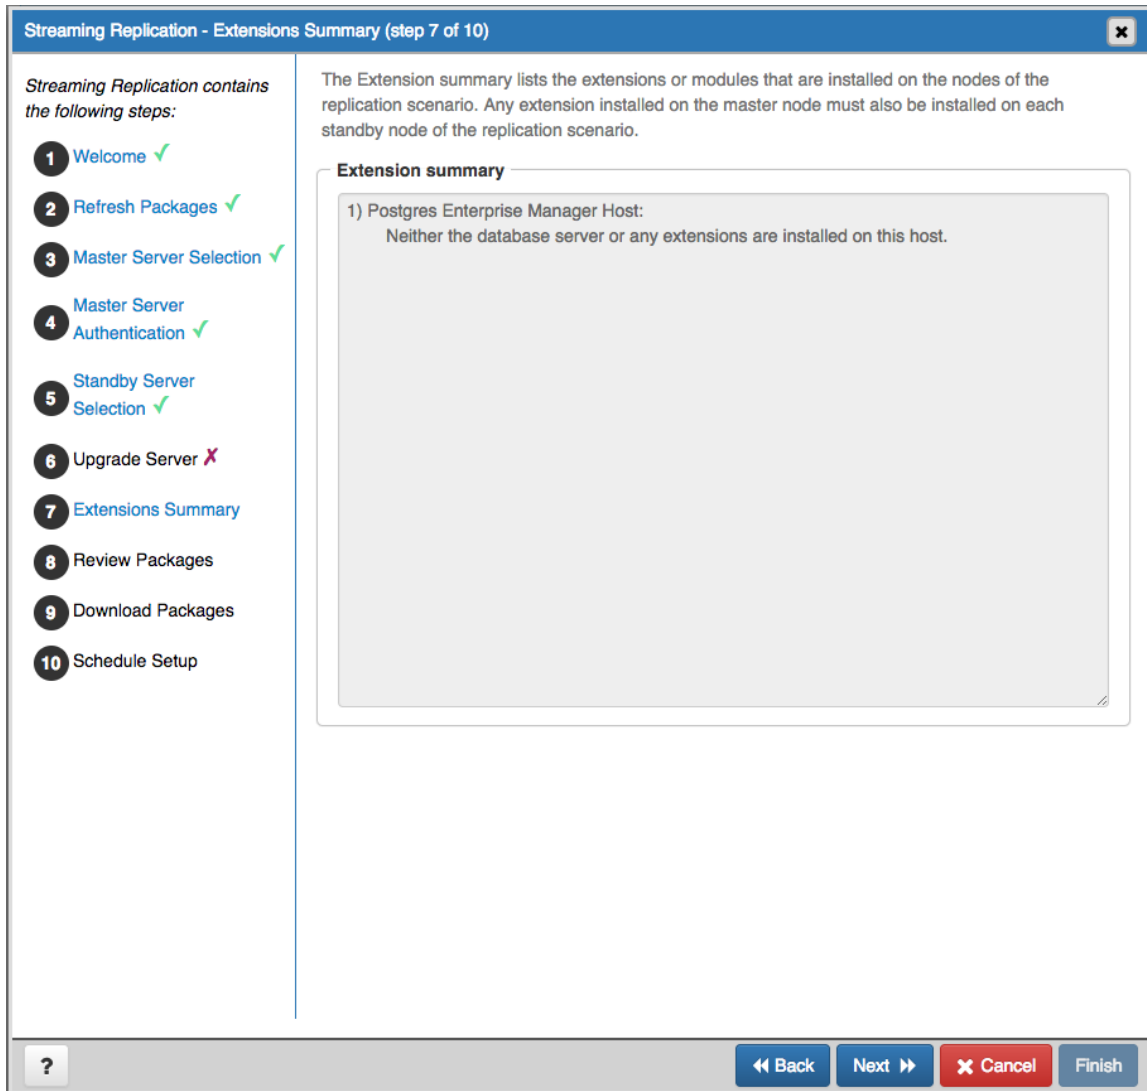


Figure 12.7 - The wizard's upgrade and extension window.

The `Extension Summary` panel (see Figure 12.7) lists the extensions or modules that are installed on the nodes of the replication scenario. Any extension installed on the master node must also be installed on each standby node of the replication scenario.

Streaming Replication - Review Packages (step 8 of 10)

Streaming Replication contains the following steps:

- 1 Welcome ✓
- 2 Refresh Packages ✓
- 3 Master Server Selection ✓
- 4 Master Server Authentication ✓
- 5 Standby Server Selection ✓
- 6 Upgrade Server ✗
- 7 Extensions Summary ✓
- 8 Review Packages
- 9 Download Packages
- 10 Schedule Setup

Review your selections and provide installer options for each selected package.

Master Node

- Postgres Enterprise Manager Host
 - Advanced Server v9.6.2.7-1
- Standby Nodes
 - Standby Agent
 - Advanced Server v9.6.2.7-1

Note: Options shown in bold are required.

Option name	Option value
Installation Directory	
Data Directory	
Write-Ahead Log (WAL) Directory	
Superaccount	enterprisedb
Operating System User Account	enterprisedb
Password	
Configuration Mode	oracle
Port	5444
Language selection	en
Extract Only	No
Unattended Mode UI	minimal

Installation Directory is required for package "Advanced Server v9.6.2.7-1" of agent "Postgres Enterprise Manager Host"

◀ Back Next ▶ ✗ Cancel Finish

Figure 12.8 - Specify installation options for the master and standbys.

If PEM is installing new servers, the Streaming Replication wizard opens to a tree control (see Figure 12.8) that provides an overview of the master and standby nodes and allows you to specify installation properties for each server. To review or modify the installation properties, highlight the name of a node in the tree control; provide values for the selected server in the `Option value` fields. Prompts on the taskbar will notify you of each required field.

When updating the installation properties, you should confirm that the user name and password specified match the name and password provided on the `Master Server Selection` dialog. You should also confirm that the specified port is not already in use on the host of the master or standby.

The `data` directory for the cluster may be created in a non-default location. If you move the `data` directory from the path specified during the installation, you must update the path specified in the `/etc/postgres-reg.ini` file.

After providing installation options, click `Next` to continue.

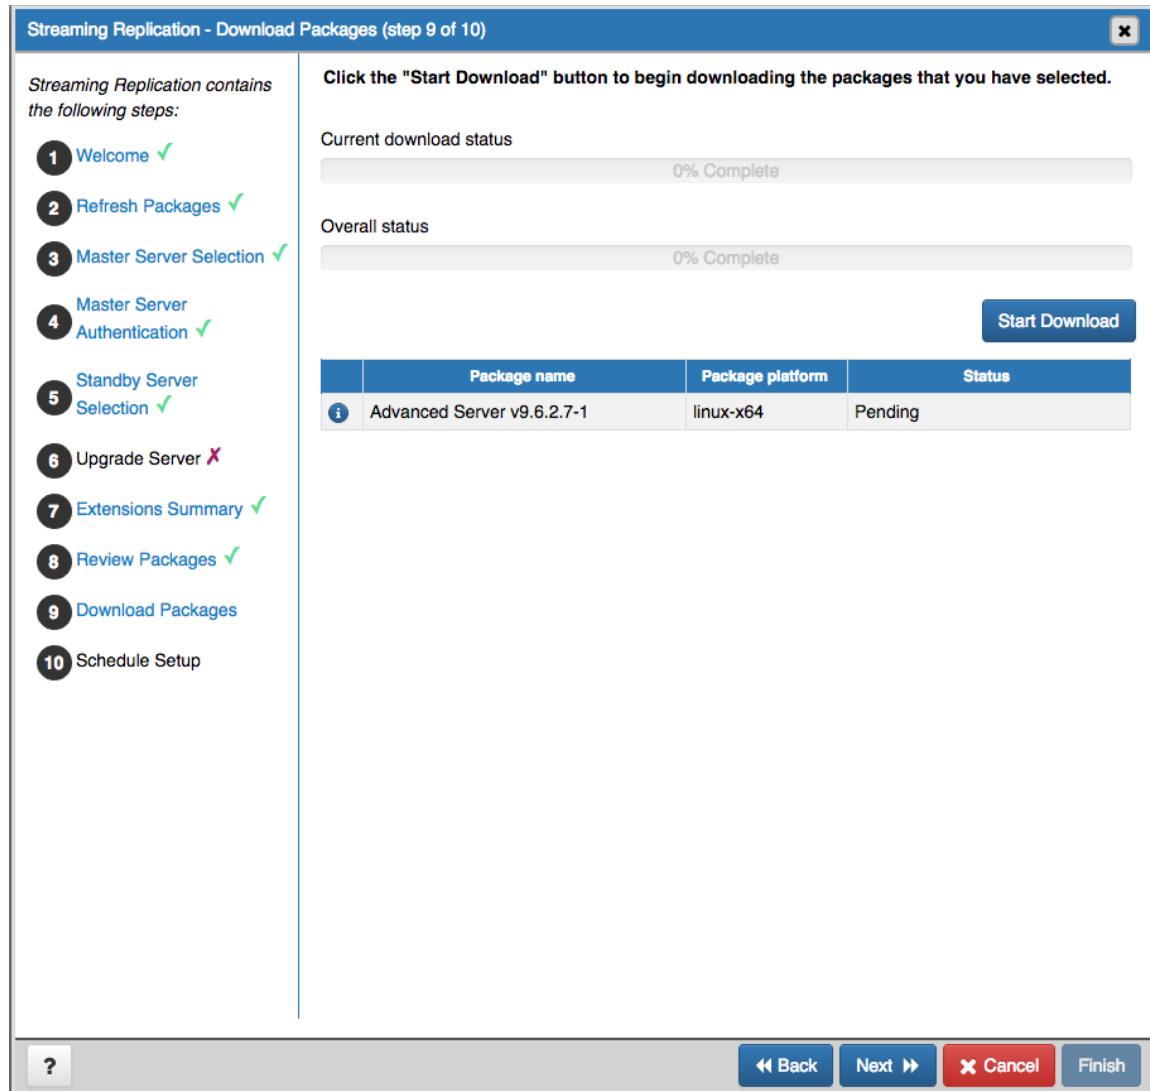


Figure 12.9 - The Download Packages dialog.

The `Download Packages` dialog (see Figure 12.9) displays a list of the packages that will be required to install the configured replication scenario. Click `Back` to return to a previous screen and modify the selections, or click `Start Download` to begin downloading the packages that will be used for the installation.

When the download completes, click `Next` to continue; the streaming replication wizard will open a dialog that allows you to schedule streaming replication setup.

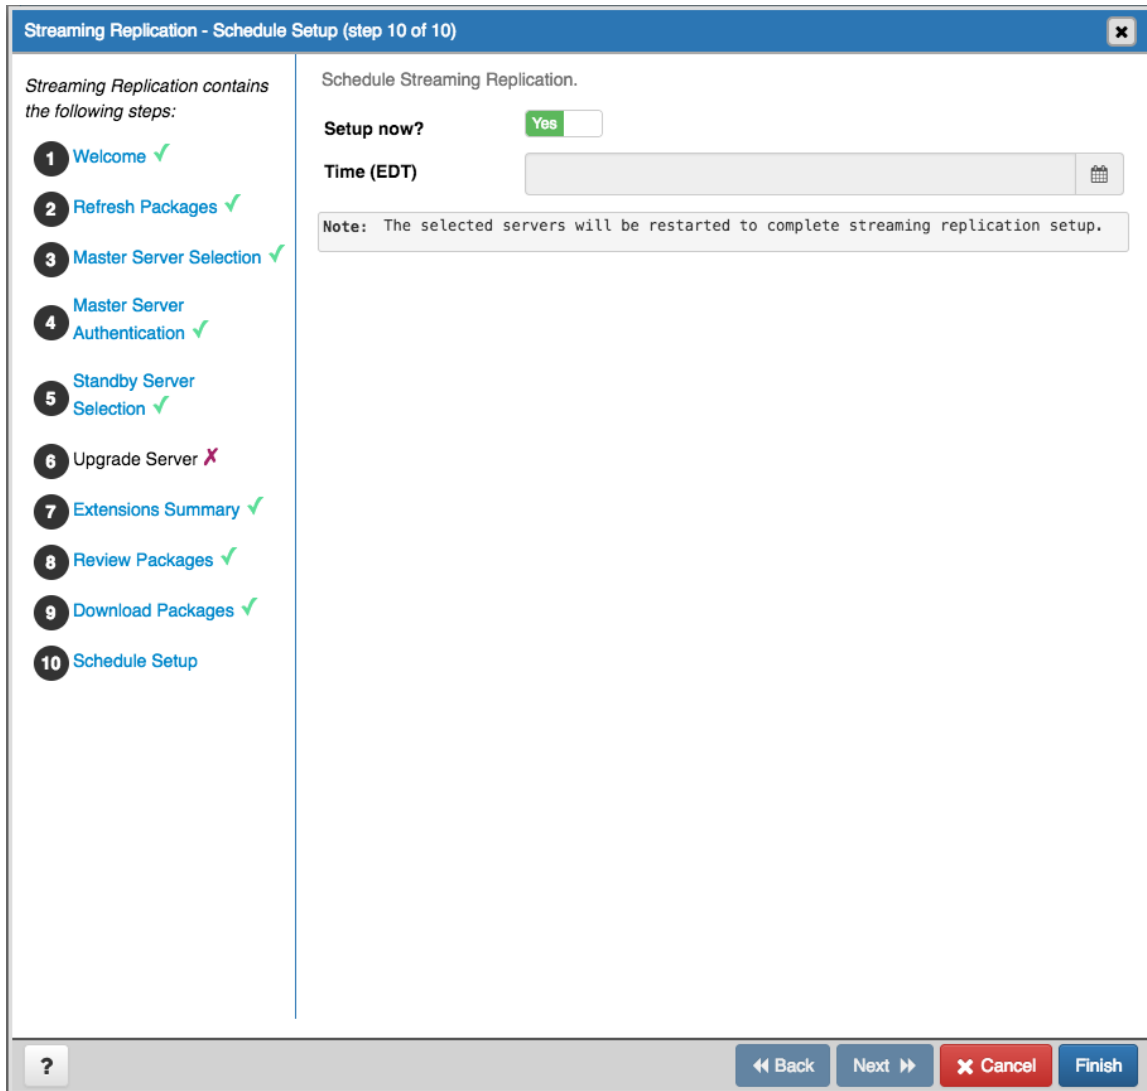


Figure 12.10 - Select a time to configure replication.

Use fields on the `Schedule Setup` dialog (see Figure 12.10) to specify the most convenient time for the server to configure the replication scenario:

- Set `Setup now?` to `Yes` to instruct PEM that it should install and configure streaming replication immediately.
- Set `Setup now?` to `No` to enable the date and time selectors; use the selectors to specify when you would like PEM to (optionally) perform installations and configure streaming replication.

Click `Finish` to save your choice and exit the wizard; PEM will either begin the installation and configuration process or schedule the installation and configuration for the specified time. You can review the job schedule and job progress on the `Scheduled`

`Tasks` tab; to open the Scheduled Tasks tab, highlight the name of the PEM agent for which you wish to review the job queue and select `Scheduled Tasks...` from the `Management` menu.

When the Streaming Replication installation wizard completes, you can register the servers and monitor the replication scenario on the Streaming Replication dashboard.

Please note that the Streaming Replication wizard only modifies the `pg_hba.conf` file on replication nodes to allow connections by the replication user; before defining a server connection in the PEM client, you may need to modify the `pg_hba.conf` file on each node to allow the connection from the client.

12.1 Monitoring Streaming Replication and Failover Manager

You can use the Streaming Replication Analysis dashboard to monitor Streaming Replication and Failover Manager high-availability scenarios. To view the Streaming Replication Analysis dashboard, you must enable probes that monitor replication-related activity; to change a probe configuration, highlight the server name in the PEM client tree control, and select `Manage Probes...` from the `Management` dashboard.

To view the Streaming Replication Analysis dashboard for the master node of a replication scenario, enable the following probes:

```
Streaming Replication
WAL Archive Status
```

To view the Streaming Replication Analysis dashboard for the standby node of a replication scenario, you must enable the following probe:

```
Streaming Replication Lag Time
```

Then, to open the Streaming Replication Analysis dashboard, navigate to the `Monitoring` tab, and:

1. Select the name of the agent that monitors the node from the `Agents` drop-down menu.
2. Select the name of the monitored server from the `Servers` drop-down menu.
3. Select `Streaming Replication Analysis` from the `Dashboards` drop-down menu.
4. Then, to open the Streaming Replication dashboard (see Figure 12.11), right click on the name of the master or standby node of the replication scenario in the `Object Browser` tree control, and select `Streaming Replication Analysis` from the `Dashboards` context menu.



Figure 12.11 The Streaming Replication dashboard.

12.1.1 Configuring High-Availability for PEM

Replication ensures that data written to the Master node of a cluster is preserved on a Standby node; if a problem occurs on the Master node (such as hardware failure), a Standby node can easily be promoted to replace the failed node. The behavior of a replication cluster can be described as:

active/active – In an active/active cluster, the master node or nodes manages write transactions while the standby nodes are available for read requests. Streaming Replication *with* hot standby or EDB Postgres Multi-Master Replication manage active/active clusters. If you are using an active/active replication scenario, you probably have a PEM agent on each node of the cluster, and are actively monitoring each node with PEM.

active/passive – In an active/passive cluster, only the master node is used for read and write transactions. Standby nodes ensure that in the event of a failure of the master node, a complete backup is available to replace the master node. RHCS (Red Hat Cluster Suite), Veritas Replicator, or Streaming Replication *without* hot standby manage active/passive clusters. If you are using an active/passive replication scenario, you are probably only actively monitoring the master node of your replication scenario with PEM.

If PEM is configured to monitor the master and standby nodes, and a standby is promoted, the PEM agents that are currently monitoring the nodes will continue to monitor the newly promoted master unless the node on which the agent resides fails.

If PEM is configured to only monitor the master node and the master node fails, the new master node will be unmonitored unless you configure a standby PEM agent to takeover monitoring of the new master. To create an agent hierarchy that allows an agent to assume monitoring a newly promoted Master, you should install a PEM agent on the Master node, and on any Standby node that might be promoted to master; each agent should be bound to the PEM server.

To configure PEM to promote an agent on a Standby node to monitor the newly promoted Master, you must:

- Enable takeover on any Standby node that might be promoted.
- Add a line to your failover script that creates a flag directory, and instructs the agent to assume monitoring.

To enable takeover of a server, right-click on the name of a server in the PEM client tree control, and select **Properties** from the context menu. When the **Properties** dialog opens, check the box next to **Allow takeover?** on the **PEM Agent** tab of the **Server**

`Properties` dialog. Each server that will potentially be promoted to the role of Master should be configured to allow takeover by another agent.

After allowing takeover of the server, add configuration steps to your failover script that instruct the PEM agent on a promoted Standby node (the new Master node) to assume monitoring the database server.

To instruct the agent to takeover the monitoring of a server, the failover process must create a file in a special *flag* directory which will instruct the agent to take responsibility for the specified server. By default, the flag directory used by the agent is:

```
$TMPDIR/pem/agent-AGENTID
```

Where `$TMPDIR` is the temporary directory for the user account under which the agent runs.

The user account is usually `root` on a Linux system or `Administrator` on Windows. You can override the directory path by specifying a value for the `AgentFlagDir` configuration option in the registry on Windows, or in the `agent_flag_dir` parameter in the agent configuration file on Linux.

For example, you might add the following command to a failover script on a Linux server:

```
touch /tmp/pem/agent-agent_id/takeover-server-server_id
```

where

agent_id is the numeric identifier of the agent that should takeover the monitoring of the server

server_id is the numeric identifier of the server that will be taken over.

To find the *agent_id* and *server_id*, log into the PEM client, and highlight the name of the agent or server; the numeric identifier will be displayed in the `ID` row on the `Properties` pane of the PEM client.

13 Monitoring Failover Manager

If you are using EDB Failover Manager to monitor your replication scenario, you must manually install and configure Failover Manager. For detailed information about installing Failover Manager, visit the EnterpriseDB website at:

<https://www.enterprisedb.com/products/edb-postgres-platform/edb-postgres-failover-manager>

To monitor the status of a Failover Manager cluster on the Streaming Replication dashboard, you must provide the following information on the **Advanced** tab of the **server Properties** dialog for each node of the cluster:

- Use the `EFMClusterName` field to specify the name of the Failover Manager cluster. The cluster name is the prefix of the name of the cluster properties file. For example, if your cluster properties file is named `efm.properties`, your cluster name is `efm`.
- Use the `EFMInstallationPath` field to specify the location of the Failover Manager binary file. By default, the Failover Manager binary file is installed in `/usr/efm-2.1/bin`.

After registering your servers, the Streaming Replication Analysis dashboard will display status information about your EFM cluster near the bottom of the dashboard (see Figure 13.1).

Failover Manager Cluster Status

Failover Manager Cluster Information

Properties	Values
Cluster Name	efm
Failover Manager Agent Running Status	UP
Allowed Node List	192.168.2.31, 192.168.2.39
Standby Priority List	192.168.2.39
Cluster Status Message	

Failover Manager Node Status

Agent Type	Address	Agent	DB	XLog Location	Status Information	XLog Information	VIP	VIP Status
Master	192.168.2.31	UP	UP	0/4E01ED90	null		192.168.2.160	True
Standby	192.168.2.39	UP	UP	0/4E01ED90	null		192.168.2.160	False

Figure 13.1 - The Failover Manager cluster status report.

The **Failover Manager Cluster Status** section of the Streaming Replication Analysis dashboard displays information about the monitored cluster:

The **Failover Manager Cluster Information** table provides information about the Failover Manager cluster:

- The **Properties** column displays the name of the cluster property.

- The `Values` column displays the current value of the property.

The `FailoverManagerNodeStatus` table displays information about each node of the Failover Manager cluster:

- The `Agent Type` column displays the type of agent that resides on the node; the possible values are Master, Standby, Witness, Idle, and Promoting.
- The `Address` column displays the IP address of the node.
- The `Agent` column displays the status of the agent that resides on the node.
- The `DB` column displays the status of the database that resides on the node.
- The `XLog Location` column displays the transaction log location of the database.
- The `Status Information` column displays any error-related information about the node.
- The `XLog Information` column displays any error-related information about the transaction log.
- The `VIP` column displays the VIP address that is associated with the node.
- The `VIP Status` column displays `True` if the VIP is active for the node, `False` if the VIP is not.

13.1.1 Replacing a Master Node

You can use the PEM client to replace the Master node of a Failover Manager cluster with a standby node. To initiate the failover process, select `Replace Cluster Master` from the `Management` menu. A dialog opens, asking you to confirm that you wish to replace the current master node (see Figure 13.2).

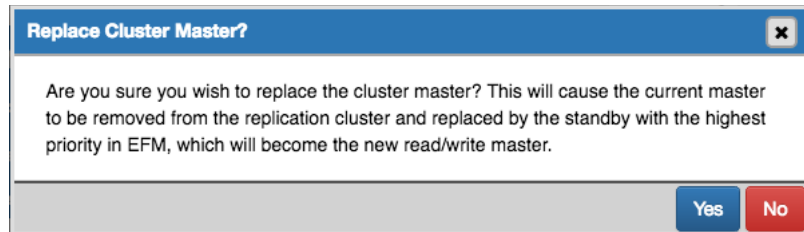


Figure 13.2 – Replacing the Master node of a cluster.

Select `Yes` to remove the current master node from the Failover Manager cluster and promote a standby node to the role of read/write master node within a Failover Manager cluster. The node with the highest promotion priority (defined in Failover Manager) will become the new master node. PEM will display a dialog, reporting the job status (see Figure 13.3).

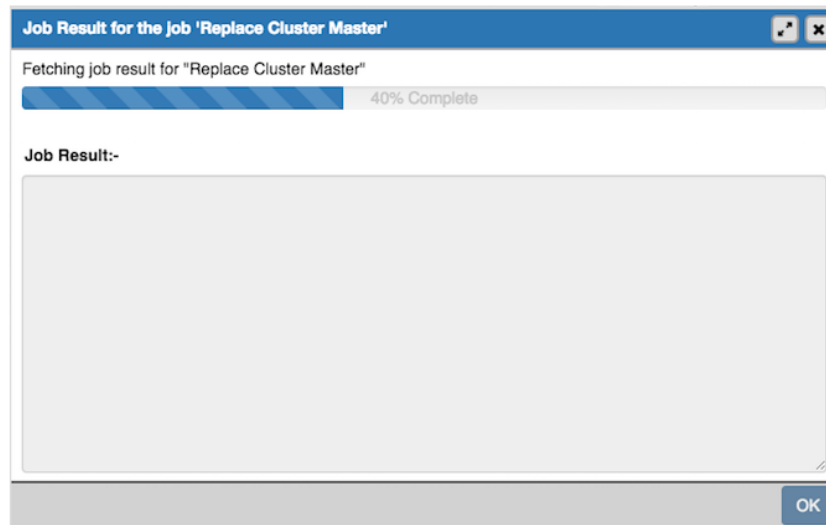


Figure 13.3 – Confirmation of the promotion.

When the job completes and the Streaming Replication Analysis dashboard refreshes, you can review the `Failover Manager Node Status` table to confirm that a standby node has been promoted to the role of master within the Failover Manager cluster.

14 Monitoring an xDB Replication Cluster

Before configuring PEM to retrieve statistics from an Advanced Server or PostgreSQL database that is part of an xDB replication scenario, you must manually install and configure xDB Replication. For more information about xDB replication solutions and documentation, please visit the EnterpriseDB website at:

<http://www.enterprisedb.com/products-services-training/products-overview/xdB-replication-server-multi-master>

The PEM xDB Replication probe monitors lag data for clusters that use xDB multi-master or single-master replication that have a publication database that is an EDB Postgres Advanced Server or PostgreSQL database. Please note that if you have configured replication between other proprietary database hosts (i.e. Oracle or SQL Server) and Advanced Server or PostgreSQL, the probe cannot return lag information.

Dashboard Properties SQL Statistics Dependencies Dependents Monitoring Manage Probes									
Description Manage Custom Probes: PEM uses probes to retrieve statistics from a monitored server, database, operating system or agent. You can view, reconfigure, delete or create your own custom probes. Copy Probes: PEM allows copying of probes from any of chosen object recursively down through the object hierarchy. Click on "Copy Probes" to quickly copy the displayed probe configurations to a selected target.									
Quick Links Manage Custom Probes Copy Probes Help									
Probes									
Probe name	Execution Frequency			Enabled?			Data Retention		
	Default	Minutes	Seconds	Default	Probe Enable?		Default	Days	
Database Frozen XID	Yes	720	0	Yes	Yes		Yes	180	
Database Size	Yes	30	0	Yes	Yes		Yes	180	
Database Statistics	Yes	30	0	Yes	Yes		Yes	90	
Function Statistics	Yes	30	0	Yes	Yes		Yes	90	
Index Size	Yes	30	0	Yes	Yes		Yes	180	
Index Statistics	Yes	30	0	Yes	Yes		Yes	90	
Materialized View Bloat	Yes	30	0	Yes	Yes		Yes	180	
Materialized View Frozen XID	Yes	720	0	Yes	Yes		Yes	180	
Materialized View Size	Yes	30	0	Yes	Yes		Yes	180	
Object Catalog: Database	Yes	5	0	Yes	Yes		Yes	180	
Object Catalog: Extension	Yes	1440	0	Yes	Yes		Yes	30	
Object Catalog: Schema	Yes	5	0	Yes	Yes		Yes	180	
Slony Replication	Yes	5	0	Yes	Yes		Yes	180	
SQL/Protect	Yes	5	0	Yes	No		Yes	180	
Streaming Replication Database Conflicts	Yes	5	0	Yes	No		Yes	180	
Table Bloat	Yes	30	0	Yes	Yes		Yes	180	
Table Frozen XID	Yes	720	0	Yes	Yes		Yes	180	
Table Size	Yes	30	0	Yes	Yes		Yes	180	
Table Statistics	Yes	30	0	Yes	Yes		Yes	90	
xdB Replication	Yes	5	0	Yes	No		Yes	180	

Figure 14.1 – The Manage Probes tab.

By default, the probe is xDB Replication probe is disabled (see Figure 14.1). To enable the xDB Replication probe, right click on the name of the server, and select **Connect** from the context menu; if prompted, provide authentication information. After

connecting, expand the server node of the tree control, and highlight the name of the replicated database. Then, select `Manage Probes...` from the `Management` menu.

Use fields on the `Manage Probes` tab to configure the xDB Replication probe:

- Move the `Default` slider to `No` to modify the `Minutes` and `Seconds` between probe executions.
- Use the `Enabled?` slider to instruct PEM to execute the xDB Replication probe.
- Set the `Default` slider in the `Data Retention` field to `No` to modify the number of days that PEM will store the information retrieved by the probe.

After enabling the probe, you can use the metrics returned to create custom charts and dashboards in the PEM client.

15 Reference

The following sections are provided for reference; please note that the items referred to in the following tables are subject to change.

15.1 PEM Server Configuration Parameters - Reference

You can use global configuration options to modify aspects of the PEM Server's behavior. For detailed information about using the Server Configuration dialog to modify server parameters, see Section **Error! Reference source not found.**

The server configuration parameters used are listed below. Please note that the list of configuration parameters is subject to change.

Parameter name	Value and Unit	Description
audit_log_retention_time	30 days	Specifies the number of days that an audit log will be retained on the PEM server.
auto_create_agent_alerts	true	Specifies whether to create default agent level alerts automatically when an agent is registered.
auto_create_server_alerts	true	Specifies whether to create default server level alerts automatically when a server is bound to an agent.
chart_disable_bullets	false	Enable/disable bullets on line charts on dashboards and Capacity Manager reports.
cm_data_points_per_report	50	Specifies the number of data points to plot on charts on Capacity Manager reports.
cm_max_end_date_in_years	5 years	Specifies the maximum amount of time that the Capacity Manager will extrapolate data for. Ensures that threshold-based end dates of on reports do not get extrapolated indefinitely.
dash_alerts_timeout	60 seconds	Specifies the number of seconds after which the components of the Alerts dashboard are auto-refreshed.
dash_db_comrol_span	7 days	Specifies the number of days worth of data to plot on the Commit/Rollback Analysis chart on the Database Analysis dashboard and Server Analysis dashboard.
dash_db_comrol_timeout	1800 seconds	Specifies the number of seconds after which the Commits/Rollbacks line chart is auto-refreshed on the Database Analysis dashboard and Server Analysis dashboard.
dash_db_connovervw_timeout	300 seconds	Specifies the number of seconds after which the Connection Overview pie chart is auto-refreshed in the Database Analysis dashboard.
dash_db_eventlag_span	7 days	Specifies the number of days worth of data to plot on the Number of Events Lag chart for slony replication on the Database Analysis

		dashboard.
dash_db_eventlag_timeout	1800 seconds	Specifies the number of seconds after which the Number of Events Lag line chart for slony replication is auto-refreshed on the Database Analysis dashboard.
dash_db_hottable_rows	25 rows	Specifies the number of rows to show on the HOT Table Analysis table on the Database Analysis dashboard.
dash_db_hottable_timeout	300 seconds	Specifies the number of seconds after which the Hot Tables table is auto-refreshed in the Database Analysis dashboard.
dash_db_io_span	7 days	Specifies the number of days worth of data to plot on the Database I/O Analysis chart on the Database Analysis dashboard and I/O Analysis dashboard.
dash_db_io_timeout	1800 seconds	Specifies the number of seconds after which the Database I/O line chart is auto-refreshed on the Database Analysis dashboard and I/O Analysis dashboard.
dash_db_rowact_span	7 days	Specifies the number of days worth of data to plot on the Row Activity Analysis chart on the Database Analysis dashboard, the I/O Analysis dashboard, and the Server Analysis dashboard.
dash_db_rowact_timeout	1800 seconds	Specifies the number of seconds after which the Row Activity line chart is auto-refreshed on the Database Analysis dashboard, the I/O Analysis dashboard, and the Server Analysis dashboard.
dash_db_storage_timeout	300 seconds	Specifies the number of seconds after which the Storage bar chart is auto-refreshed in the Database Analysis dashboard.
dash_db_timelag_span	7 days	Specifies the number of days worth of data to plot on the Time Lag chart for Slony replication on the Database Analysis dashboard.
dash_db_timelag_timeout	1800 seconds	Specifies the number of seconds after which the Time Lag line chart for Slony replication is auto-refreshed on the Database Analysis dashboard.
dash_db_useract_span	7 days	Specifies the number of days worth of data to plot on the User Activity Analysis chart on the Database Analysis dashboard.
dash_db_useract_timeout	1800 seconds	Specifies the number of seconds after which the User Activity line chart is auto-refreshed in the Database Analysis dashboard.
dash_efm_timeout	300 seconds	Specifies the number of seconds after which the Failover Manager Node Status and Failover Manager Cluster Info line chart is auto-refreshed on the Streaming Replication dashboard.
dash_global_overview_timeout	30 seconds	Specifies the number of seconds after which the components of the Global Overview dashboard are auto-refreshed.
dash_header_timeout	60 seconds	Specifies the number of seconds after which the information on the header of all the

		dashboards are auto-refreshed.
dash_io_chkpt_span	7 days	Specifies the number of days worth of data to plot on the Checkpoints chart on the I/O Analysis dashboard.
dash_io_chkpt_timeout	1800 seconds	Specifies the number of seconds after which the Checkpoints line chart is auto-refreshed on the I/O Analysis dashboard.
dash_io_hotindx_timeout	300 seconds	Specifies the number of seconds after which the Hot Indexes bar chart is auto-refreshed on the I/O Analysis dashboard.
dash_io_hottbl_timeout	300 seconds	Specifies the number of seconds after which the Hot Tables bar chart is auto-refreshed on the I/O Analysis dashboard.
dash_io_index_objectio_rows	25 rows	Specifies the number of rows displayed on the Index Activity table on the I/O Analysis dashboard and the Object Activity Analysis dashboard.
dash_io_index_objectio_timeout	60 seconds	Specifies the number of seconds after which the Index Activity table is auto-refreshed on the I/O Analysis dashboard and the Object Activity Analysis dashboard.
dash_io_objectio_rows	25 rows	Specifies the number of rows displayed in the Object I/O Details table on the I/O Analysis dashboard and Object Activity Analysis dashboard.
dash_io_objectio_timeout	300 seconds	Specifies the number of seconds after which the Object I/O Details table is auto-refreshed on the I/O Analysis dashboard and Object Activity Analysis dashboard.
dash_memory_hostmemact_span	7 days	Specifies the number of days worth of data to plot on the Host Memory Activity Analysis chart on the Memory Analysis dashboard.
dash_memory_hostmemact_timeout	1800 seconds	Specifies the number of seconds after which the Host Memory Activity line chart is auto-refreshed on the Memory Analysis dashboard.
dash_memory_hostmemconf_timeout	300 seconds	Specifies the number of seconds after which the Host Memory Configuration pie chart
		is auto-refreshed on the Memory Analysis dashboard and Server Analysis dashboard.
dash_memory_servmemact_span	7 days	Specifies the number of days worth of data to plot on the server Memory Activity Analysis chart on the Memory Analysis dashboard.
dash_memory_servmemact_timeout	1800 seconds	Specifies the number of seconds after which the Server Memory Activity line chart is auto-refreshed on the Memory Analysis dashboard.
dash_memory_servmemconf_timeout	300 seconds	Specifies the number of seconds after which the Server Memory Configuration pie chart is auto-refreshed on the Memory Analysis dashboard.
dash_objectact_objstorage_rows	15 rows	Specifies the number of rows to show on the Object Storage table on the Object Activity Analysis dashboard.
dash_objectact_objstorage_timeout	300 seconds	Specifies the number of seconds after which

timeout		the Object Storage table is auto-refreshed in the Object Activity Analysis dashboard.
dash_objectact_objtopindexes_timeout	300 seconds	Specifies the number of seconds after which the Top 5 Largest Indexes bar chart is auto-refreshed in the Object Activity Analysis dashboard.
dash_objectact_objtoptables_timeout	300 seconds	Specifies the number of seconds after which the Top 5 Largest Tables bar chart is auto-refreshed in the Object Activity Analysis dashboard.
dash_os_cpu_span	7 days	Specifies the number of days worth of data to plot on the CPU chart on the Operating System Analysis dashboard.
dash_os_cpu_timeout	1800 seconds	Specifies the number of seconds after which the CPU line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_data_span	7 days	Specifies the number of days worth of data to plot on the I/O line chart on the Operating System Analysis dashboard.
dash_os_disk_span	7 days	Specifies the number of days worth of data to plot on the Utilisation chart on the Operating System Analysis dashboard.
dash_os_hostfs_timeout	1800 seconds	Specifies the number of seconds after which the Host File System Details table is auto-refreshed on the Operating System Analysis dashboard.
dash_os_io_timeout	1800 seconds	Specifies the number of seconds after which the I/O line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_memory_span	7 days	Specifies the number of days worth of data to plot on the Memory chart on the Operating System Analysis dashboard.
dash_os_memory_timeout	1800 seconds	Specifies the number of seconds after which the Memory line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_packet_span	7 days	Specifies the number of days worth of data to plot on the Packet chart on the Operating System Analysis dashboard.
dash_os_packet_timeout	1800 seconds	Specifies the number of seconds after which the Network Packets line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_process_span	7 days	Specifies the number of days worth of data to plot on the Process chart on the Operating System Analysis dashboard.
dash_os_process_timeout	1800 seconds	Specifies the number of seconds after which the Process line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_storage_timeout	1800 seconds	Specifies the number of seconds after which the Storage pie chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_traffic_span	7 days	Specifies the number of days worth of data to plot on the Traffic chart on the Operating System Analysis dashboard.

dash_os_traffic_timeout	1800 seconds	Specifies the number of seconds after which the Traffic line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_os_util_timeout	1800 seconds	Specifies the number of seconds after which the Utilisation line chart is auto-refreshed on the Operating System Analysis dashboard.
dash_probe_log_timeout	300 seconds	Specifies the number of seconds after which the Probe Log table is auto-refreshed on
dash_replication_archivist_at_span	7 days	Specifies the number of days worth of data to plot on the WAL Archive Status chart on the Streaming Replication Analysis dashboard.
dash_replication_archivist_at_timeout	1800 seconds	Specifies the number of seconds after which the WAL Archive Status line chart is auto-refreshed on the Streaming Replication dashboard.
dash_replication_pagelag_span	7 days	Specifies the number of days worth of data to plot on the WAL Lag Pages chart on the Streaming Replication dashboard.
dash_replication_pagelag_timeout	1800 seconds	Specifies the number of seconds after which the WAL Lag Pages line chart is auto-refreshed on the Streaming Replication dashboard.
dash_replication_segmentlag_span	7 days	Specifies the number of days worth of data to plot on the WAL Lag Segments chart on the Streaming Replication dashboard.
dash_replication_segmentlag_timeout	1800 seconds	Specifies the number of seconds after which the WAL Lag Segments line chart is auto-refreshed on the Streaming Replication dashboard.
dash_replication_timelag_span	7 days	Specifies the number of days worth of data to plot on the Replication Lag Time chart on the Streaming Replication dashboard.
dash_replication_timelag_timeout	1800 seconds	Specifies the number of seconds after which the Replication Lag Time line chart is auto-refreshed on the Streaming Replication dashboard.
dash_server_buffers_writen	168 hours	Specifies the number of days worth of data to plot on the Background Writer Statistics chart on the Server Analysis dashboard.
dash_server_buffers_writen_timeout	300 seconds	Specifies the number of seconds after which the Background Writer Statistics line chart is auto-refreshed on the Server Analysis dashboard.
dash_server_connoverview_timeout	300 seconds	Specifies the number of seconds after which the Connection Overview pie chart is auto-refreshed in the Server Analysis dashboard.
dash_server_database_timeout	300 seconds	Specifies the number of seconds after which the Databases table is auto-refreshed in the Server Analysis dashboard.
dash_server_dbsize_span	7 days	Specifies the number of days worth of data to plot on the Database Size Analysis on the Server Analysis dashboard.
dash_server_dbsize_timeout	1800 seconds	Specifies the number of seconds after which the Database Size line chart is auto-refreshed in the Server Analysis dashboard.

dash_server_disk_timeout	1800 seconds	Specifies the number of seconds after which the Disk line chart is auto-refreshed in the Server Analysis dashboard.
dash_server_global_span	7 days	Specifies the number of days worth of data to plot on the Disk line chart on the Server Analysis dashboard.
dash_server_sharedbuff_span	7 days	Specifies the number of days worth of data to plot on the Shared Buffer chart on the Server Analysis dashboard.
dash_server_sharedbuff_timeout	1800 seconds	Specifies the number of seconds after which the Shared Buffers line chart is auto-refreshed in the Server Analysis dashboard.
dash_server_tabspacesize_span	7 days	Specifies the number of days worth of data to plot on the Tablespace Size chart on the Server Analysis dashboard.
dash_server_tabspacesize_timeout	1800 seconds	Specifies the number of seconds after which the Tablespace Size line chart is auto-refreshed in the Server Analysis dashboard.
dash_server_useract_span	7 days	Specifies the number of days worth of data to plot on the User Activity chart on the Server Analysis dashboard.
dash_server_useract_timeout	1800 seconds	Specifies the number of seconds after which the User Activity line chart is auto-refreshed in the Server Analysis dashboard.
dash_sessact_lockact_timeout	300 seconds	Specifies the number of seconds after which the Session Lock Activity table is auto-refreshed in the Session Activity Analysis dashboard.
dash_sessact_workload_timeout	300 seconds	Specifies the number of seconds after which the Session Workload table is auto-refreshed in the Session Activity Analysis dashboard.
dash_sess_waits_nowaits_timeout	300 seconds	Specifies the number of seconds after which the Session Waits By Number Of Waits pie
dash_sess_waits_timewait_timeout	300 seconds	Specifies the number of seconds after which the Session Waits By Time Waited pie chart is auto-refreshed in the Session Waits Analysis dashboard.
dash_sess_waits_waitdtl_timeout	300 seconds	Specifies the number of seconds after which the Session Waits Details table is auto-refreshed in the Session Waits Analysis dashboard.
dash_storage_dbdtls_timeout	300 seconds	Specifies the number of seconds after which the Database Details table is auto-refreshed in the Storage Analysis dashboard.
dash_storage_dbovervw_timeout	300 seconds	Specifies the number of seconds after which the Database Overview pie chart is auto-refreshed in the Storage Analysis dashboard.
dash_storage_hostdtls_timeout	300 seconds	Specifies the number of seconds after which the Host Details table is auto-refreshed
dash_storage_hostovervw_timeout	300 seconds	Specifies the number of seconds after which the Host Overview pie chart is auto-refreshed in the Storage Analysis dashboard.
dash_storage_tblspcdtls_timeout	300 seconds	Specifies the number of seconds after which the Tablespace Details table is auto-refreshed in

meout		the Storage Analysis dashboard.
dash_storage_tblspcovervw_timeout	300 seconds	Specifies the number of seconds after which the Tablespace Overview pie chart is auto-refreshed in the Storage Analysis dashboard.
dash_sys_waits_nowaits_timeout	300 seconds	Specifies the number of seconds after which the System Waits By Number Of Waits pie chart is auto-refreshed in the System Waits Analysis dashboard.
dash_sys_waits_timewait_timeout	300 seconds	Specifies the number of seconds after which the System Waits By Time Waited pie chart is auto-refreshed in the System Waits Analysis dashboard.
dash_sys_waits_waitdtl_timeout	300 seconds	Specifies the number of seconds after which the System Waits Details table is auto-refreshed in the System Waits Analysis dashboard.
deleted_charts_retention_time	7 days	Specifies the number of days that a custom chart (displayed on a user-defined dashboard) is stored.
deleted_probes_retention_time	7 days	Specifies the number of days that a custom probe (displayed on a user-defined dashboard) is stored.
download_chart_format	jpeg	Specifies the format in which a downloaded chart will be stored. May be jpeg or png.
flapping_detection_state_change	3	Specifies the number of state changes detected within a specified interval to define a given alert as flapping.
job_retention_time	30 days	Specifies the number of days that non-recurring scheduled tasks and their associated
long_running_transaction_minutes	5 minutes	Specifies the number of minutes a query executes for before being considered long running.
nagios_cmd_file_name	<file_name>	Specifies nagios command file to which passive service check result will be sent.
nagios_enabled	t	Specifies whether alert notification will be submitted to nagios or not.
nagios_medium_alert_as_critical	f	Specifies whether medium level PEM alert will be considered as critical in nagios.
nagios_spool_retention_time	7 days	Specifies the number of days to retain nagios messages in the spool table before they are discarded.
package_catalog_xml	<address>	Specifies path of the catalog file which will be read by package catalog probe to get the list of all the supported packages.
package_download_chunk_size	1024 bytes	Specify the size(in Bytes) to be read from network while downloading packages. By default, through PEM 6, 1KB, for PEM 7 and forward, 1MB.
probe_log_retention_time	30 days	Specifies the number of days that probe log records are retained.
proxy_server	127.0.0.1	Use this parameter to specify the IP Address of a proxy server.
proxy_server_authentication	f	Set this parameter to t (true) if your proxy server requires authentication.

proxy_server_enabled	f	If you use a proxy server on a client machine (when connecting with internet), enable this setting to read the manifest file and download packages.
proxy_server_password		If your proxy server requires authentication, use this parameter to provide the password that will be used for authentication.
proxy_server_port	80	Use this parameter to specify the port for a proxy server.
proxy_server_username		If your proxy server requires authentication, use this parameter to provide the
reminder_notification_interval	24 hours	Specifies the number of hours after which a reminder email is sent in case an alert has not been cleared.
server_log_retention_time	30 days	Specifies the number of days that the server log is retained on the PEM server.
show_data_tab_on_graph	false	If 'true', a Data tab is added to each graph. Select the Data tab to review the data that is plotted on the graph.
smtp_authentication	false	Specifies whether to enable/disable authentication over SMTP.
smtp_enabled	true	Specifies whether to enable/disable sending of emails.
smtp_encryption	false	Specifies whether to send SMTP email using an encrypted connection.
smtp_password		Specifies the password to be used to connect to the SMTP server.
smtp_port	25	Specifies the SMTP server port to be used for sending email.
smtp_server	127.0.0.1	Specifies the SMTP server host address to be used for sending email.
smtp_spool_retention_time	7 days	Specifies the number of days to retain sent email messages in the spool table before they are discarded.
smtp_username		Specifies the username to be used to connect to SMTP server.
snmp_community	public	Specifies the SNMP community used when sending traps. Used only with SNMPv1 and SNMPv2.
snmp_enabled	true	Specifies whether to enable/disable sending SNMP traps.
snmp_port	162	Specifies the SNMP server port to be used for sending SNMP traps.
snmp_server	127.0.0.1	Specifies the SNMP server host address to be used for sending SNMP traps.
snmp_spool_retention_time	7 days	Specifies the number of days to retain sent traps in the spool table before they are discarded.
webclient_help_pg	EnterpriseDB hosted documentation	Specifies the location of the online PostgreSQL core documentation.
web_client_product_key		Specifies the product key of the PEM web client. This parameter may be changed by users with Administrative privileges.

15.2 Capacity Manager Metrics - Reference

Please Note that the Capacity Manager metrics available will vary by platform, and are subject to change. The available metrics may include the metrics described in the table below.

Metric Name	Description
# Dead Tuples	The number of dead tuples in the selected table.
# Dead Tuples+	The cumulative number of dead tuples in the selected table.
# Heap Tuples Fetched by Index Scans	The number of heap tuples fetched by index scans.
# Heap Tuples Fetched by Index Scans	The cumulative number of heap tuples fetched by index scans.
# Idle Backends+	The cumulative number of currently idle backend clients.
# Index Scans	The number of index scans performed on the specified object.
# Index Scans+	The cumulative number of index scans performed on the specified object.
# Index Tuples Read	The number of index tuples read.
# Index Tuples Read+	The cumulative number of index tuples read.
# Live Tuples	The number of tuples visible to transactions.
# Live Tuples+	The cumulative number of tuples visible to transactions.
# Pages Estimated by ANALYZE	The number of pages estimated by ANALYZE.
# Pages Estimated by ANALYZE+	The cumulative number of pages estimated by ANALYZE.
# Sequential Scans	The number of sequential scans performed on the specific table.
# Sequential Scans+	The cumulative number of sequential scans performed on the specific table.
# Sequential Scan Tuples	The number of tuples sequentially scanned in the specific table.
# Sequential Scan Tuples+	The cumulative number of tuples sequentially scanned in the specific table.
# Tuples Deleted	The number of tuples deleted.
# Tuples Deleted+	The cumulative number of tuples deleted.
# Tuples Estimated by ANALYZE	The number of live (visible) tuples estimated by ANALYZE.
# Tuples Estimated by ANALYZE+	The cumulative number of live tuples estimated by ANALYZE.
# Tuples HOT Updated	The number of tuples HOT updated. In a HOT update, the new tuple resides in the same block as the original tuple and the tuples share an index entry.
# Tuples HOT Updated+	The cumulative number of tuples HOT updated.
# Tuples Inserted	The number of tuples inserted into the specified table.
# Tuples Inserted+	The cumulative number of tuples inserted into the specified table.
# Tuples Updated	The number of tuples updated in the selected table.
# Tuples Updated+	The cumulative number of tuples updated in the selected table.
Blocks Hit	The number of blocks found in the cache.
Blocks Hit+	The cumulative number of blocks found in the cache.
Blocks Read	The number of blocks read.
Blocks Read+	The cumulative number of blocks read.
Blocks Read from InfiniteCache	The number of blocks read from InfiniteCache.
Blocks Read from InfiniteCache+	The cumulative number of blocks read from InfiniteCache.
Blocks Written	The number of blocks written.

Blocks Written+	The cumulative number of blocks written.
Buffers Allocated	The number of buffers allocated.
Buffers Allocated+	The cumulative number of buffers allocated.
Buffers Written - Backends	The number of buffer blocks written to disk by server processes (processes connected to a client application).
Buffers Written - Backends+	The cumulative number of buffer blocks written to disk by server processes.
Buffers Written - Checkpoint	The number of blocks written to disk by the checkpoint process.
Buffers Written - Checkpoint+	The cumulative number of blocks written to disk by the checkpoint process.
Buffers Written - Cleaning Scan	The number of blocks written to disk by the autovacuum process.
Buffers Written - Cleaning Scan+	The cumulative number of blocks written to disk by the autovacuum process.
Bytes Received (KB)	The number of bytes received from the client (in kilobytes).
Bytes Received (KB)+	The cumulative number of bytes received (in kilobytes).
Bytes Sent (KB)	The number of bytes sent to the client (in kilobytes).
Bytes Sent (KB)+	The cumulative number of bytes sent (in kilobytes).
Checkpoints - Timed	The number of checkpoint operations triggered by the checkpoint interval.
Checkpoints - Timed+	The cumulative number of checkpoint operations triggered by the checkpoint interval.
Checkpoints - Untimed	The number of checkpoint operations triggered by checkpoint size.
Checkpoints - Untimed+	The cumulative number of checkpoint operations triggered by checkpoint size.
Database Size (MB)	The size of the specified database (in megabytes).
Free RAM Memory	The amount of free RAM memory (in megabytes).
Free Swap Memory	The amount of free swap space on disk (in megabytes).
Heap Blocks Hit	The number of heap blocks found in the cache.
Heap Blocks Hit+	The cumulative number of heap blocks found in the cache.
Heap Blocks Read	The number of heap blocks read.
Heap Blocks Read+	The cumulative number of heap blocks read.
Index Blocks Hit	The number of index blocks found in the cache.
Index Blocks Hit+	The cumulative number of index blocks found in the cache.
Index Blocks Read	The number of index blocks read.
Index Blocks Read+	The cumulative number of index blocks read.
Index Size (MB)	The size of the specified index (in megabytes).
In Packets Discards	The number of inbound packets discarded.
In Packets Discards+	The cumulative number of inbound packets discarded.
In Packets Errors	The number of inbound packets that contain errors.
In Packets Errors+	The cumulative number of inbound packets that contain errors.
Link Bandwidth (Mbit/s)	The speed of the network adapter (in megabits per second).
Load Average - 15 Minute	CPU saturation (in percent) - 15 minute sampling average.
Load Average - 1 Minute	CPU saturation (in percent) - 1 minute sampling average.
Load Average - 5 Minute	CPU saturation (in percent) - 5 minute sampling average.
Load Percentage	CPU saturation in percent.
Number of Prepared Transactions+	The cumulative number of prepared transactions.
Number of WAL Files+	The cumulative number of write-ahead log files.
Out Packets Discards	The number of outbound packets discarded.
Out Packets Discards+	The cumulative number of outbound packets discarded.
Out Packets Errors	The number of outbound packets that contain errors.
Out Packets Errors+	The cumulative number of outbound packets that contain errors.

Packets Received	The number of packets received.
Packets Received+	The cumulative number of packets received.
Packets Sent	The number of packets sent.
Packets Sent+	The cumulative number of packets sent.
Size (MB)	The total size of the disk (in megabytes).
Size of Indexes (MB)	The size of indexes on the specified table (in megabytes).
Space Available (MB)	The current disk space available (in megabytes).
Space Used (MB)	The current disk space used (in megabytes).
Table Size (MB)	The size of the specified table (in megabytes).
Tablespace Size (MB)	The size of the specified tablespace (in megabytes).
Temp Buffers (MB)	The size of temporary buffers (in megabytes).
Toast Blocks Hit	The number of TOAST blocks found in the cache.
Toast Blocks Hit+	The cumulative number of TOAST blocks found in the cache.
Toast Blocks Read	The number of TOAST blocks read.
Toast Blocks Read+	The cumulative number of TOAST blocks read.
Total RAM Memory	The total amount of RAM memory on the system (in megabytes).
Total Swap Memory	The total amount of swap space on the system (in megabytes).
Total Table Size w/Indexes and Toast	The total size of the specified table (including indexes and associated oversized attributes).
Transactions Aborted	The number of aborted transactions.
Transactions Aborted+	The cumulative number of aborted transactions.
Transactions Committed	The number of committed transactions.
Transactions Committed+	The cumulative number of committed transactions.
Tuples Deleted	The number of tuples deleted from the specified table.
Tuples Deleted+	The cumulative number of tuples deleted from the specified table.
Tuples Estimated by ANALYZE	The number of visible tuples in the specified table.
Tuples Estimated by ANALYZE+	The cumulative number of visible tuples in the specified table.
Tuples Fetched	The number of tuples fetched from the specified table.
Tuples Fetched+	The cumulative number of tuples fetched from the specified table.
Tuples HOT Updated	The number of tuples HOT updated. In a HOT update, the new tuple resides in the same block as the original tuple and the tuples share an index entry.
Tuples HOT Updated+	The cumulative number of tuples HOT updated. In a HOT update, the new tuple resides in the same block as the original tuple and the tuples share an index entry.
Tuples Inserted	The number of tuples inserted into the specified table.
Tuples Inserted+	The cumulative number of tuples inserted into the specified table.
Tuples Returned	The number of tuples returned in result sets.
Tuples Returned+	The cumulative number of tuples returned in result sets.
Tuples Updated	The number of tuples updated in the specified table.
Tuples Updated+	The cumulative number of tuples updated in the specified table.
WAL Segment Size (MB)	The segment size of the write-ahead log (in megabytes).

****Note:**** The '+' following the name of a metric signifies that the data for the metric is gathered cumulatively; those metrics that are not followed by the '+' sign are collected as a 'point-in-time' value.

15.3 PEM Probes – Reference

A probe is a scheduled task that retrieves information about the database objects that are being monitored by the PEM agent. PEM uses the collected information to build the graphs displayed on each dashboard. The `Manage Probes` tab (accessed via the `Management` menu) allows you to modify the data collection schedule and the length of time that PEM will retain information returned by a specific probe.

Probe Name	Information Monitored by Probe	Level
Background Writer Statistics	This probe monitors information about the background writer. The information includes: The number of timed checkpoints The number of requested checkpoints The number of buffers written (by checkpoint) The number of buffers written (by background writer) The number of background writer cycles The number of background buffers written The number of buffers allocated	Server
Blocked Session Information	This probe provides information about blocked sessions.	Server
CPU Usage	This probe monitors CPU Usage information.	Agent
Data and Log File Analysis	This probe monitors information about log files. The information includes: The name of the log file The directory in which the log file resides	Server
Database Statistics	This probe monitors database statistics. The information includes: The number of backends The number of transactions committed The number of transactions rolled back The number of blocks read The number of blocks hit The number of rows returned The number of rows fetched The number of rows inserted The number of rows updated The number of rows deleted	Server
Disk Busy Info	This probe monitors information about disk activity. Note: This probe is not supported on Mac OS X, Solaris or HP-UX	Agent
Disk Space	This probe monitors information about disk space usage. The information includes: The amount of disk space used The amount of disk space available	Agent
EDB Audit Configuration	This probe monitors the audit logging configuration of	Server

	EDB Postgres Advanced Server.	
Failover Manager Cluster Info	This probe monitors a Failover Manager cluster, returning information about the cluster. This probe is disabled unless a cluster name and path of the Failover Manager binary is provided on the Server Properties dialog.	Server
Failover Manager Node Status	This probe monitors a Failover Manager cluster, returning detailed about each node within the cluster. This probe is disabled unless a cluster name and path of the Failover Manager binary is provided on the Server Properties dialog.	Server
Function Statistics	This probe monitors a database, retrieving information about functions. The information includes: Function names Argument types Return values	Database
Index Size	This probe monitors a database, retrieving information about indexes. The information includes: The name of the index The time the data was gathered The size of the index (in MB's)	Database
Index Statistics	This probe monitors index statistics. The information includes: The number of index scans The number of rows read The number of rows fetched The number of blocks read The number of blocks hit	Database
Installed Packages	This probe monitors the packages that are currently installed. The information gathered includes: The name of the installed package The version of the installed package The date and time that the probe executed	Agent
IO Analysis	This probe monitors disk I/O information in. The information includes: The number of blocks read The number of blocks written The date and time that the probe executed Note: This probe is not supported on Mac OS X	Agent
Load Average	This probe monitors CPU load averages. The information includes: The 1-minute load average The 5-minute load average The 15-minute load average Note: This probe is not supported on Windows	Agent
Lock Information	This probe monitors lock information. The information includes:	Server

	The database name The lock type The lock mode The process holding the lock	
Memory Usage	This probe monitors information about system memory usage.	Agent
Network Statistics	This probe monitors network statistics. The information includes: The interface IP address The number of packets sent The number of packets received The number of bytes sent The number of bytes received The link speed (in MB/second)	Agent
Number of Prepared Transactions	This probe stores the number of prepared transactions.	Server
Number of WAL Files	This probe monitors the number of WAL files.	Server
Object Catalog: Database	This probe monitors a list of databases and their properties. The information includes: The database name The database encoding type If the database allows user connections or system connections	Server
Object Catalog: Foreign Key	This probe monitors a list of foreign keys and their properties. The information includes: The name of the table that contains the foreign key The name of the table that the foreign key references The name of the database in which the table resides The name of the schema in which the table resides	Schema
Object Catalog: Function	This probe monitors a list of functions and their properties. The information includes: The name of the function The name of the schema in which the function resides The name of the database in which the function resides	Schema
Object Catalog: Index	This probe monitors a list of indexes and their properties. The information includes: The name of the index The name of the table that the index is associated with The name of the database in which the indexed table resides	Schema
Object Catalog: Schema	This probe monitors a list of schemas and their associated databases and servers.	Database
Object Catalog: Sequence	This probe monitors a list of sequences and their properties.	Schema
Object Catalog: Table	This probe monitors a list of table information. The information includes:	Schema

	<p>The table name</p> <p>The name of the schema in which the table resides</p> <p>The name of the database in which the schema resides</p> <p>A Boolean indicator that indicates if the table has a primary key</p>	
Object Catalog: Tablespace	This probe monitors a list of tablespaces.	Server
Operating System Information	This probe monitors the operating system details and boot time.	Agent
Package Catalog	<p>This probe monitors the packages that are currently available for installation. The information gathered includes:</p> <p>The package name</p> <p>The package version</p>	Agent
PG HBA Conf	This probe monitors authentication configuration information from the <code>pg_hba.conf</code> file.	Server
Server Information	This probe monitors server information.	Server
Session Information	<p>This probe monitors session information. The information includes:</p> <p>The name of the session user</p> <p>The date and time that the session connected to the server</p> <p>The status of the session at the time that the information was gathered (idle, waiting, etc)</p> <p>The client address and port number</p>	Server
Settings	This probe monitors the values currently assigned to GUC variables.	Server
SQL Protect	This probe monitors a server, retrieving information about SQL injection attacks.	Server
Slony Replication	This probe monitors lag data for clusters replicated using Slony.	Database
Streaming Replication	<p>This probe monitors a cluster that is using streaming replication, retrieving information about:</p> <p>The sent Xlog location (in bytes)</p> <p>The write Xlog location (in bytes)</p> <p>The flush Xlog location (in bytes)</p> <p>The replay Xlog location (in bytes)</p> <p>The Xlog lag (in segments)</p> <p>The Xlog lag (in pages)</p>	Server
Streaming Replication Lag Time	<p>This probe monitors a cluster that is using streaming replication, retrieving lag information about:</p> <p>Replication lag time (in seconds)</p> <p>Current status of replication (running/paused)</p>	Server
Streaming Replication Database Conflicts	<p>This probe monitors a database that is using streaming replication, retrieving information about any conflicts that arise. This includes information about queries that have been canceled due to:</p> <p>The # of drop tablespace conflicts</p> <p>The # of lock timeout conflicts</p> <p>The # of old snapshot conflicts</p>	Server

	<p>The # of pinned buffer conflicts</p> <p>The # of deadlock conflicts</p>	
Table Bloat	<p>This probe monitors information about the current table bloat. The information includes:</p> <p>The name of the table</p> <p>The name of the schema in which the table resides</p> <p>The estimated number of pages</p> <p>The estimated number of wasted pages</p> <p>The estimated number of bytes per row</p>	Database
Table Frozen XID	This probe monitors the frozen XID of each table.	Schema
Table Size	<p>This probe monitors table statistics. The information includes:</p> <p>The number of sequential scans</p> <p>The number of sequential scan rows</p> <p>The number of index scans</p> <p>The number of index scan rows</p> <p>The number of rows inserted</p> <p>The number of rows updated</p> <p>The number of rows deleted</p> <p>The number of live rows</p> <p>The number of dead rows</p> <p>The last VACUUM</p> <p>The last auto-vacuum</p> <p>The last ANALYZE</p> <p>The last auto-analyze</p> <p>The number of pages estimated by ANALYZE</p> <p>The number of rows estimated by ANALYZE</p>	Database
Table Statistics	This probe monitors a list of tablespaces and their sizes.	Server
Tablespace Size	This probe monitors a list of tablespaces and their sizes.	Server
User Information	<p>This probe monitors a list of the current users. The stored information includes:</p> <p>The user name</p> <p>The user type (superuser vs. non-superuser)</p> <p>The server to which the user is connected</p>	Server
WAL Archive Status	<p>This probe monitors the status of the WAL archive. The stored information includes:</p> <p>The # of WAL archives done</p> <p>The # of WAL archives pending</p> <p>The last archive time</p> <p>The # of WAL archives failed</p> <p>The time of the last failure</p>	Server
xDB Replication	This probe monitors lag data for clusters replicated using xDB replication.	Database

15.4 PEM Pre-defined Alert Templates – Reference

An alert definition contains a system-defined or user-defined set of conditions that PEM compares to the system statistics; if the statistics deviate from the boundaries specified for that statistic, the alert triggers, and the PEM client displays a warning on the *Alerts Overview* page, and optionally sends a notification to a monitoring user.

The tables that follow list the system-defined alert templates that you can use to create an alert; please note that this list is subject to change, and may vary by system:

15.4.1 Templates applicable on Agent

Template Name	Description
Load Average (1 minute)	1-minute system load average.
Load Average (5 minutes)	5-minute system load average.
Load Average (15 minutes)	15-minute system load average.
Load Average per CPU Core (1 minutes)	1-minute system load average per CPU core.
Load Average per CPU Core (5 minutes)	5-minute system load average per CPU core.
Load Average per CPU Core (15 minutes)	15-minute system load average per CPU core.
CPU utilization	Average CPU consumption.
Number of CPUs running higher than a	Number of CPUs running at greater than K% utilization threshold
Free memory percentage	Free memory as a percent of total system memory.
Memory used percentage	Percentage of memory used.
Swap consumption	Swap space consumed (in megabytes).
Swap consumption percentage	Percentage of swap area consumed.
Disk Consumption	Disk space consumed (in megabytes).
Disk consumption percentage	Percentage of disk consumed.
Disk Available	Disk space available (in megabytes).
Disk busy percentage	Percentage of disk busy.
Most used disk percentage	Percentage used of the most utilized disk on the system.
Total table bloat on host	The total space wasted by tables on a host, in MB.
Highest table bloat on host	The most space wasted by a table on a host, in MB.
Average table bloat on host	The average space wasted by tables on host, in MB.
Table size on host	The size of tables on host, in MB.
Database size on host	The size of databases on host, in MB.
Number of ERRORS in the logfile on agent N in last X hours.	The number of ERRORS in the logfile on agent N in last X hours
Number of WARNINGS in the logfile on agent N in last X hours	The number of WARNINGS in the logfile on agent N in last X hours.
Number of WARNINGS or ERRORS in the logfile on agent N in last X hours	The number of WARNINGS or ERRORS in the logfile on agent N in last X hours.

Package version mismatch	Check for package version mismatch as per catalog.
Total materialized view bloat on host	The total space wasted by materialized views on a host, in MB.
Highest materialized view bloat on host	The most space wasted by a materialized view on a host, in MB.
Average materialized view bloat on host	The average space wasted by materialized views on host, in MB.
Materialized view size on host	The size of materialized views on host, in MB.
Agent Down	Specified agent is currently down.

15.4.2 Templates applicable on Server

Template Name	Description
Total table bloat in server	The total space wasted by tables in server, in MB.
Largest table (by multiple of unbloated size)	Largest table in server, calculated as a multiple of its own estimated unbloated size; exclude tables smaller than N MB.
Highest table bloat in server	The most space wasted by a table in server, in MB.
Average table bloat in server	The average space wasted by tables in server, in MB.
Table size in server	The size of tables in server, in MB.
Database size in server	The size of databases in server, in MB.
Number of WAL files	Total number of Write Ahead Log files.
Number of prepared transactions	Number of transactions in prepared state.
Total connections	Total number of connections in the server.
Total connections as percentage of	Total number of connections in the server as a percentage of maximum
max_connections	connections allowed on server, settings.
Unused, non-superuser connections	Number of unused, non-superuser connections on the server, user_info, settings.
Unused, non-superuser connections as percentage of max_connections	Number of unused, non-superuser connections on the server as a percentage of max_connections of max_connections, user_info, settings.
Ungranted locks	Number of ungranted locks in server.
Percentage of buffers written by backends	The percentage of buffers written by backends vs. the total buffers written.
Percentage of buffers written by checkpoint	The percentage of buffers written by the checkpoints vs. the total buffers written.
Buffers written per second	Number of buffers written per second, over the last two probe cycles.
Buffers allocated per second	Number of buffers allocated per second, over the last two probe cycles.
Connections in idle state	Number of connections in server that are in idle state.
Connections in idle-in-transaction state	Number of connections in server that are in idle-in-transaction state.
Connections in idle-in-transaction state, as percentage of max_connections	Number of connections in server that are in idle-in-transaction state, as a percentage of maximum connections allowed on server, settings.
Long-running idle connections	Number of connections in the server that have been idle for more than N seconds.
Long-running idle connections	Number of connections in the server that have been idle or

and idle transactions	transactions idle-in-transaction for more than N seconds.
Long-running idle transactions	Number of connections in the server that have been idle in transaction for more than N seconds.
Long-running transactions	Number of transactions in server that have been running for more than N seconds.
Long-running queries	Number of queries in server that have been running for more than N seconds.
Long-running vacuums	Number of vacuum operations in server that have been running for more than N seconds.
Long-running autovacuum	Number of autovacuum operations in server that have been running for more than N seconds.
Committed transactions percentage	Percentage of transactions in the server that committed vs. that rolled-back over last N minutes.
Shared buffers hit percentage	Percentage of block read requests in the server that were satisfied by shared buffers, over last N minutes.
Tuples inserted	Tuples inserted into server over last N minutes.
InfiniteCache buffers hit percentage	Percentage of block read requests in the server that were satisfied by InfiniteCache, over last N minutes.
Tuples fetched	Tuples fetched from server over last N minutes.
Tuples returned	Tuples returned from server over last N minutes.
Dead Tuples	Number of estimated dead tuples in server.
Tuples updated	Tuples updated in server over last N minutes.
Tuples deleted	Tuples deleted from server over last N minutes.
Tuples hot updated	Tuples hot updated in server, over last N minutes.
Sequential Scans	Number of full table scans in server, over last N minutes.
Index Scans	Number of index scans in server, over last N minutes.
Hot update percentage	Percentage of hot updates in the server over last N minutes.
Live Tuples	Number of estimated live tuples in server.
Dead tuples percentage	Percentage of estimated dead tuples in server.
Last Vacuum	Hours since last vacuum on the server.
Last AutoVacuum	Hours since last autovacuum on the server.
Last Analyze	Hours since last analyze on the server.
Last AutoAnalyze	Hours since last autoanalyze on the server.
Percentage of buffers written by backends over the last N minutes	The percentage of buffers written by backends vs. the total buffers backends over last N
Table Count	Total number of tables in server.
Function Count	Total number of functions in server.
Sequence Count	Total number of sequences in server.
A user expires in N days	Number of days before a user's validity expires.
Index size as a percentage of table size	Size of the indexes in server, as a percentage of their tables' size.
Largest index by table-size percentage oc_index, table_size.	Largest index in server, calculated as percentage of its table's size.
Number of ERRORS in the logfile on server M in the last X hours	The number of ERRORS in the logfile on server M in last X hours.
Number of WARNINGS in the logfile on server M in the last X hours	The number of WARNINGS in logfile on server M in the last X hours.
Number of WARNINGS or ERRORS in the logfile on server	The number of WARNINGS or ERRORS in the logfile on server M in the last X hours.

M in the last X hours	
Number of attacks detected in the last N minutes	The number of SQL injection attacks occurred in the last N minutes.
Number of attacks detected in the last N minutes by username	The number of SQL injection attacks occurred in the last N minutes by username.
Number of standby servers lag behind the master by write location	Streaming Replication: number of standby servers lag behind the master by write location.
Number of standby servers lag behind the master by flush location	Streaming Replication: number of standby servers lag behind the master by flush location.
Number of standby servers lag behind the master by replay location	Streaming Replication: number of standby servers lag behind the master by replay location.
Standby server lag behind the master by write location	Streaming Replication: standby server lag behind the master by write location in MB.
Standby server lag behind the master by flush location	Streaming Replication: standby server lag behind the master by flush location in MB.
Standby server lag behind the master by replay location	Streaming Replication: standby server lag behind the master by replay location in MB.
Standby server lag behind the master by size (MB)	Streaming Replication: standby server lag behind the master by size in MB.
Standby server lag behind the master by WAL segments	Streaming Replication: standby server lag behind the master by WAL segments.
Standby server lag behind the master by WAL pages	Streaming Replication: standby server lag behind the master by WAL pages.
Total materialized view bloat in server	The total space wasted by materialized views in server, in MB.
Largest materialized view (by multiple of unbloated size)	Largest materialized view in server, calculated as a multiple of its own estimated unbloated size; exclude materialized views smaller than N MB.
Highest materialized view bloat in server	The most space wasted by a materialized view in server, in MB.
Average materialized view bloat in server	The average space wasted by materialized views in server, in MB.
Materialized view size in server	The size of materialized view in server, in MB.
View Count	Total number of views in server.
Materialized View Count	Total number of materialized views in server.
Audit config mismatch	Check for audit config parameter mismatch
Server Down	Specified server is currently inaccessible.
Number of WAL archives pending	Streaming Replication: number of WAL files pending to be replayed at standby.
Number of minutes lag of standby server from master server	Streaming Replication: number of minutes standby node is lagging behind the master node.
Log config mismatch	Check for log config parameter mismatch.

15.4.3 Templates applicable on Database

Template Name	Description
Total table bloat in database	The total space wasted by tables in database, in MB.
Largest table (by multiple of unbloated size)	Largest table in database, calculated as a multiple of its own estimated unbloated size; exclude tables smaller than N MB.
Highest table bloat in database	The most space wasted by a table in database, in MB.
Average table bloat in database	The average space wasted by tables in database, in MB.
Table size in database	The size of tables in database, in MB.
Database size	The size of the database, in MB.
Total connections	Total number of connections in the database.
Total connections as percentage of max_connections	Total number of connections in the database as a percentage of maximum connections allowed on server, settings.
Ungranted locks	Number of ungranted locks in database.
Connections in idle state	Number of connections in database that are in idle state.
Connections in idle-in-transaction state	Number of connections in database that are in idle-in-transaction state
Connections in idle-in-transaction state, as percentage of max_connections	Number of connections in database that are in idle-in-transaction state, as a percentage of maximum connections allowed on server, settings.
Long-running idle connections	Number of connections in the database that have been idle for more than N seconds.
Long-running idle connections and idle transactions	Number of connections in the database that have been idle or idle-in-transaction for more than N seconds.
Long-running idle transactions	Number of connections in the database that have been idle in transaction for more than N seconds.
Long-running transactions	Number of transactions in database that have been running for more than N seconds.
Long-running queries	Number of queries in database that have been running for more than N seconds.
Long-running vacuums	Number of vacuum operations in database that have been running for more than N seconds.
Long-running autovacums	Number of autovacuum operations in database that have been running for more than N seconds.
Committed transactions percentage	Percentage of transactions in the database that committed vs. that rolled-back over last N minutes.
Shared buffers hit percentage	Percentage of block read requests in the database that were satisfied by shared buffers, over last N minutes.
InfiniteCache buffers hit percentage	Percentage of block read requests in the database that were satisfied by InfiniteCache, over last N minutes.
Tuples fetched	Tuples fetched from database over last N minutes.
Tuples returned	Tuples returned from database over last N minutes.
Tuples inserted	Tuples inserted into database over last N minutes.
Tuples updated	Tuples updated in database over last N minutes.
Tuples deleted	Tuples deleted from database over last N minutes.
Tuples hot updated	Tuples hot updated in database, over last N minutes.
Sequential Scans	Number of full table scans in database, over last N minutes.

Index Scans	Number of index scans in database, over last N minutes.
Hot update percentage	Percentage of hot updates in the database over last N minutes.
Live Tuples	Number of estimated live tuples in database.
Dead Tuples	Number of estimated dead tuples in database.
Dead tuples percentage	Percentage of estimated dead tuples in database.
Last Vacuum	Hours since last vacuum on the database.
Last AutoVacuum	Hours since last autovacuum on the database.
Last Analyze	Hours since last analyze on the database.
Last AutoAnalyze	Hours since last autoanalyze on the database.
Table Count	Total number of tables in database.
Function Count	Total number of functions in database.
Sequence Count	Total number of sequences in database.
Index size as a percentage of table size	Size of the indexes in database, as a percentage of their tables' size.
Largest index by table-size percentage	Largest index in database, calculated as percentage of its table's size, <code>oc_index</code> , <code>table_size</code> .
Database Frozen XID	The age (in transactions before the current transaction) of the database's frozen transaction ID.
Number of attacks detected in the	The number of SQL injection attacks occurred in the last N minutes.
Number of attacks detected in the	The number of SQL injection attacks occurred in the last N minutes by last N minutes by username.
Queries that have been cancelled due to dropped tablespaces	Streaming Replication: number of queries that have been cancelled due to dropped tablespaces.
Queries that have been cancelled due to lock timeouts	Streaming Replication: number of queries that have been cancelled due to lock timeouts.
Queries that have been cancelled due to old snapshots	Streaming Replication: number of queries that have been cancelled due to old snapshots.
Queries that have been cancelled due to pinned buffers	Streaming Replication: number of queries that have been cancelled due to pinned buffers.
Queries that have been cancelled due to deadlocks	Streaming Replication: number of queries that have been cancelled due to deadlocks.
Total events lagging in all slony clusters	Slony Replication: total events lagging in all slony clusters.
Events lagging in one slony cluster	Slony Replication: events lagging in one slony cluster.
Lag time (minutes) in one slony cluster	Slony Replication: lag time (minutes) in one slony cluster.
Total rows lagging in xdb single master replication	xDB Replication: Total rows lagging in xdb single master replication
Total rows lagging in xdb multi master replication	xDB Replication: Total rows lagging in xdb multi master replication.
Total materialized view bloat in database	The total space wasted by materialized views in database, in MB.
Largest materialized view (by multiple of unbloated size)	Largest materialized view in database, calculated as a multiple of its estimated unbloated size; exclude materialized views smaller than N MB.
Highest materialized view bloat in database	The most space wasted by a materialized view in database, in MB.

Average materialized view bloat in database	The average space wasted by materialized views in database, in MB.
Materialized view size in database	The size of materialized view in database, in MB.
View Count	Total number of views in database.
Materialized View Count	Total number of materialized views in database.

15.4.4 Templates applicable on Schema

Template Name	Description
Total table bloat in schema	The total space wasted by tables in schema, in MB.
Largest table (by multiple of unbloated size)	Largest table in schema, calculated as a multiple of its own estimated unbloated size; exclude tables smaller than N MB.
Highest table bloat in schema	The most space wasted by a table in schema, in MB.
Average table bloat in schema	The average space wasted by tables in schema, in MB.
Table size in schema	The size of tables in schema, in MB.
Tuples inserted	Tuples inserted in schema over last N minutes.
Tuples updated	Tuples updated in schema over last N minutes.
Tuples deleted	Tuples deleted from schema over last N minutes.
Tuples hot updated	Tuples hot updated in schema, over last N minutes.
Sequential Scans	Number of full table scans in schema, over last N minutes.
Index Scans	Number of index scans in schema, over last N minutes.
Hot update percentage	Percentage of hot updates in the schema over last N minutes.
Live Tuples	Number of estimated live tuples in schema.
Dead Tuples	Number of estimated dead tuples in schema.
Dead tuples percentage	Percentage of estimated dead tuples in schema.
Last Vacuum	Hours since last vacuum on the schema.
Last AutoVacuum	Hours since last autovacuum on the schema.
Last Analyze	Hours since last analyze on the schema.
Last AutoAnalyze	Hours since last autoanalyze on the schema.
Table Count	Total number of tables in schema.
Function Count	Total number of functions in schema.
Sequence Count	Total number of sequences in schema.
Index size as a percentage of table size	Size of the indexes in schema, as a percentage of their table's size.
Largest index by table-size percentage	Largest index in schema, calculated as percentage of its table's size, <code>oc_index</code> , <code>table_size</code>
Materialized View bloat	Space wasted by the materialized view, in MB.
Total materialized view bloat in schema	The total space wasted by materialized views in schema, in MB.
Materialized view size as a multiple of unbloated size	Size of the materialized view as a multiple of estimated unbloated size.
Largest materialized view (by multiple of unbloated size)	Largest materialized view in schema, calculated as a multiple of its own estimated unbloated size; exclude materialized view smaller than N MB.
Highest materialized view bloat in schema	The most space wasted by a materialized view in schema, in MB.
Average materialized view bloat in schema	The average space wasted by materialized views in schema, in MB.
Materialized view size	The size of materialized view, in MB.
Materialized view size in schema	The size of materialized views in schema, in MB.
View Count	Total number of views in schema.
Materialized View Count	Total number of materialized views in schema.
Materialized View Frozen XID	The age (in transactions before the current transaction) of the materialized view's frozen transaction ID.

15.4.5 Templates applicable on Table

Template Name	Description
Table bloat	Space wasted by the table, in MB.
Table size	The size of table, in MB.
Table size as a multiple of ubloated size	Size of the table as a multiple of estimated ubloated size.
Tuples inserted	Tuples inserted in table over last N minutes.
Tuples updated	Tuples updated in table over last N minutes.
Tuples deleted	Tuples deleted from table over last N minutes.
Tuples hot updated	Tuples hot updated in table, over last N minutes.
Sequential Scans	Number of full table scans on table, over last N minutes.
Index Scans	Number of index scans on table, over last N minutes.
Hot update percentage	Percentage of hot updates in the table over last N minutes.
Live Tuples	Number of estimated live tuples in table.
Dead Tuples	Number of estimated dead tuples in table.
Dead tuples percentage	Percentage of estimated dead tuples in table.
Last Vacuum	Hours since last vacuum on the table.
Last AutoVacuum	Hours since last autovacuum on the table.
Last Analyze	Hours since last analyze on the table.
Last AutoAnalyze	Hours since last autoanalyze on the table.
Row Count	Estimated number of rows in a table.
Index size as a percentage of table size	Size of the indexes on table, as a percentage of table's size.
Table Frozen XID	The age (in transactions before the current transaction) of the table's frozen transaction ID.

15.4.6 Global Templates

Template Name	Description
Agents Down	Number of agents that haven't reported in recently.
Servers Down	Number of servers that are currently inaccessible.
Alert Errors	Number of alerts in an error state.

EnterpriseDB, EDB Postgres, Postgres Plus, Postgres Enterprise Manager, and DynaTune are trademarks of EnterpriseDB Corporation. Other names may be trademarks of their respective owners. © 2018.

16 Conclusion

The goal of Postgres Enterprise Manager is provide you with a solution that allows you to intelligently manage all your database servers across your enterprise with a single console. To meet this objective, PEM supplies you with all the core features and functionality needed for visual database administration, as well as a number of advanced components that assist you in managing the performance and design of your database servers.

For more information about Postgres Enterprise Manager, please visit the EnterpriseDB Web site (<http://www.enterprisedb.com>) where you will find PEM's online documentation, as well as other tutorials and educational aids.

EnterpriseDB is the enterprise PostgreSQL company, providing products and services worldwide that are based on and support PostgreSQL, the world's most advanced open source database. EDB's products are ideally suited for transaction-intensive applications requiring superior performance, massive scalability, and compatibility with proprietary database products. EDB's products provide an economical open source alternative or complement to proprietary databases without sacrificing features or quality.

If you would like to discuss training, consulting, or enterprise support options, please contact EnterpriseDB. [EnterpriseDB](http://www.enterprisedb.com) has offices in North America, Europe, and Asia. EnterpriseDB was founded in 2004 and is headquartered in Bedford, MA. For more information, please visit <http://www.enterprisedb.com>.

Sales Inquiries:

sales-us@enterprisedb.com (US)
sales-intl@enterprisedb.com (Intl)
+1-781-357-3390 or 1-877-377-4352 (US Only)

General Inquiries:

info@enterprisedb.com
info.asiapacific@enterprisedb.com (APAC)
info.emea@enterprisedb.com (EMEA)