

EDBTM POSTGRES



EDB PostgresTM Enterprise Manager Getting Started Guide Version 6.0

March 7, 2016

EDB Postgres Enterprise Manager Getting Started Guide
by EnterpriseDB® Corporation
Copyright © 2016 EnterpriseDB Corporation. All rights reserved.

EnterpriseDB Corporation, 34 Crosby Drive Suite 100, Bedford, MA 01730, USA
T +1 781 357 3390 **F** +1 978 589 5701 **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. © 2016.

Table of Contents

1	Introduction.....	5
1.1	Typographical Conventions Used in this Guide	7
2	Postgres Enterprise Manager - Overview	8
2.1	Why Postgres Enterprise Manager?.....	8
2.2	General Architecture.....	10
2.3	Supported Versions and Platforms.....	11
3	Installing Postgres Enterprise Manager	13
3.1	Starting and Stopping the PEM Server and Agents	14
3.2	The PEM Client - User Interface Basics.....	16
3.3	Using the PEM Web Client.....	17
3.4	Online Help and Documentation	20
3.5	Logging on to the PEM Server	21
3.6	Adding a Managed Server	23
3.7	Remotely Starting and Stopping Monitored Servers	33
4	General Database Administration.....	34
4.1	Editing a Server's Configuration	34
4.2	Managing Security	36
4.2.1	Login Roles.....	36
4.2.2	Group Roles	37
4.2.3	Using a Team Role.....	37
4.2.4	Object Permissions.....	38
4.3	Managing Storage	39
4.4	Creating and Maintaining Databases and Objects	40
4.5	SQL Development	42
5	Package Deployment	44
5.1	Installing a New Package.....	47
5.2	Upgrading an Installed Package.....	52
6	Performance Monitoring and Management	56
6.1	Using Dashboards to View Performance Information.....	56
6.2	Creating a Custom Dashboard	60
6.2.1	Creating an Ops Dashboard	66
6.3	Creating a Custom Chart.....	68

6.3.1	Controlling and Customizing Charts, Graphs and Tables	72
6.4	Customizing Probes	73
6.5	Customizing Alerts	75
6.5.1	Creating a Custom Alert	75
6.5.1.1	Audit Log Alerting.....	80
6.5.2	Using PEM with Nagios	81
6.5.3	Creating a Custom Alert Template	82
6.6	Viewing and Responding to Alerts	87
7	Capacity Manager	88
7.1	Performing Trend Analysis.....	88
7.2	Forecasting Future Trends	92
8	Audit Manager	93
8.1	Setting the Advanced Server Instance Service ID	94
8.2	Setting the EDB Audit Configuration Probe	95
8.3	Configuring Audit Logging with the Audit Manager.....	96
8.4	Viewing the Log with the Audit Log Dashboard.....	104
9	Log Manager	106
9.1	Reviewing the Server Log Analysis Dashboard	118
10	Postgres Log Analysis Expert.....	120
10.1	Reviewing the Postgres Log Analysis Expert Report.....	126
11	SQL Profiling and Analysis	127
11.1	Setup and Configuration	127
11.2	Creating a New SQL Trace.....	128
11.3	Analyzing a SQL Trace Output	130
11.4	Using the Index Advisor	132
12	Tuning Wizard	133
13	Postgres Expert - Best Practice Enforcement	141
13.1	Using the Postgres Expert Wizard	141
13.2	Reviewing Postgres Expert Recommendations	145
14	Configuring Streaming Replication	147
14.1	Monitoring Replication and Failover.....	160
14.2	Replacing a Master Node.....	162
15	Conclusion	163
15.1	About EnterpriseDB.....	163

1 Introduction

Notice: The names for EDB's products have changed.

The product formerly referred to as Postgres Plus Advanced Server is now referred to as EDB Postgres Advanced Server (Advanced Server).

The product formerly referred to as Postgres Enterprise Manager (PEM) is now referred to as EDB Postgres Enterprise Manager (EDB Enterprise Manager).

Until a new version of this documentation is published, wherever you see an earlier version of a product name, you may substitute it with the current name. Name changes in software and software outputs will be phased in over time.

This document provides an introduction to Postgres Enterprise Manager™ (PEM) and is written to acquaint you with the basics of the toolset and help you be successful in your database management activities. This guide is broken up into the following core sections and categories:

- **Postgres Enterprise Manager Overview** – [Chapter 2](#) provides information about PEM functionality, components, architecture, and supported platforms.
- **Installation, configuration and general setup** – [Chapter 3](#) provides an overview of PEM's installation and configuration steps.
- **General database administration** – [Chapter 4](#) will assist you in performing general database administration tasks, and SQL query and stored procedure development.
- **Package Deployment** – [Chapter 5](#) introduces you to the PEM Package Deployment wizard.
- **Performance monitoring** – [Chapter 6](#) provides an introduction to monitoring and troubleshooting performance on your database servers.
- **Capacity planning** – [Chapter 7](#) discusses how to perform trend analysis and forecasting from data collected in Postgres Enterprise Manager.
- **Audit Manager** – [Chapter 8](#) 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 9](#) 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 10](#) 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 11](#) describes how to capture a SQL workload and analyze the output to tune and fix poorly running SQL.
- **Tuning Wizard** – [Chapter 12](#) 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 13](#) discusses how to use Postgres Expert to uncover configuration or design issues that need correcting on your database servers.
- **Streaming Replication Wizard** – [Chapter 14](#) 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.

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 *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 preceding term may be repeated. For example, [a | b] ... means that you may have the sequence, “b a a b a”.

2 Postgres Enterprise Manager - Overview

Postgres Enterprise Manager (PEM) is an enterprise management tool designed to assist database administrators, system architects, and performance analysts in administering, monitoring, and tuning PostgreSQL and EnterpriseDB Advanced Server database servers. PEM is architected to manage and monitor anywhere from a handful, to hundreds of servers from a single console, allowing complete and remote control over all aspects of your databases.

2.1 Why Postgres Enterprise Manager?

PEM provides a number of benefits not found in any other PostgreSQL management tool:

- **Management en Masse Design.** PEM is designed for enterprise database management, and is built to tackle the management of large numbers of servers across geographical boundaries. Global dashboards keep you up to date on the up/down/performance status of all your servers in an at-a-glance fashion.
- **Distributed Architecture.** PEM is architected in a way that maximizes its ability to gather statistical information and to perform operations remotely on machines regardless of operating system platform.
- **Graphical Administration.** All aspects of database administration can be carried out in the PEM client via a graphical interface. Server startup and shutdown, configuration management, storage and security control, object creation, performance management, and more can be handled from a single console.
- **Full SQL IDE.** PEM contains a robust SQL integrated development environment (IDE) that provides ad-hoc SQL querying, stored procedure/function development, and a graphical debugger.
- **Enterprise Performance Monitoring.** PEM provides enterprise-class performance monitoring for all managed database servers. Lightweight and efficient agents monitor all aspects of each database server's operations as well as each machine's underlying operating system and provide detailed statistics back to easily navigated performance pages within the interface.
- **Proactive Alert Management.** PEM ships out-of-the-box with the ability to create performance thresholds for each key metric (e.g. memory, storage, etc.) that are monitored around-the-clock. Any threshold violation results in an alert being sent to a centralized dashboard that communicates the nature of the problem

and what actions are necessary to prevent the situation from jeopardizing the overall performance of the server.

- **Simplified Capacity Planning.** All key performance-related statistics are automatically collected and retained for a specified period of time in PEM's repository. The Capacity Manager utility allows you to select various statistics and perform trend analysis over time to understand things such as peak load periods, storage consumption trends, and much more. A forecasting mechanism in the tool allows you to also forecast resource usage in the future and plan/budget accordingly.
- **Audit Manager.** The Audit Manager configures audit logging on Advanced Server instances. Activities such as connections to a database, disconnections from a database, and the SQL statements run against a database can be logged. The `Audit Log` dashboard can then be used to filter and view the log.
- **Log Manager.** The Log Manager wizard configures server logging parameters, with (optional) log collection into a central table. Use the wizard to specify your preference for logging behaviors such as log file rotation, log destination and error message severity. Use the `Server Log` dashboard to filter and review the collected server log entries.
- **SQL Workload Profiling.** PEM contains a SQL profiling utility that allows you to trace the SQL statements that are executed against one or more servers. SQL profiling can either be done in an ad-hoc or scheduled manner. Captured SQL statements can then be filtered so you can easily identify and tune poorly running SQL statements. SQL statements can also be fed into an Index Advisor on Advanced Server that analyzes each statement and makes recommendations on new indexes that should be created to help performance.
- **Expert Database Analysis.** PEM includes the Postgres Expert utility. Postgres Expert analyzes selected databases for best practice enforcement purposes. Areas such as general configuration, security setup, and much more are examined. Any deviations from recommended best practices are reported back to you, along with an explanation of each particular issue, and expert help on what to do about making things right.
- **Streaming Replication Configuration and Monitoring.** The Streaming Replication wizard simplifies the process of adding new servers to a Postgres streaming replication scenario or configuring existing servers to create a replication scenario. After configuring the replication scenario, you can monitor the scenario on the Streaming Replication dashboard or use options on the PEM client to promote a standby node to the master node.
- **Secure Client Connectivity.** PEM supports secure client connections through an encrypted SSH tunnel. The full-featured PEM client includes an SSH Tunnel

definition dialog that allows you to provide connection information for a secure connection.

- **Wide Platform Support.** PEM supports most major Linux and Windows platforms.

2.2 General Architecture

PEM is composed of three primary components (see Figure 2.1):

The PEM Server

The PEM server provides the functionality at the core of Postgres Enterprise Manager. The server is responsible for:

- Performing administration functions
- Processing information received from agents
- Maintaining information in its repository.

The PEM Agent

The PEM agent is responsible for performing tasks on each managed machine and collecting statistics for the database server and operating system.

The PEM Client

The PEM client is the user console from which all operations are carried out (e.g. database administration, viewing performance information, etc.).

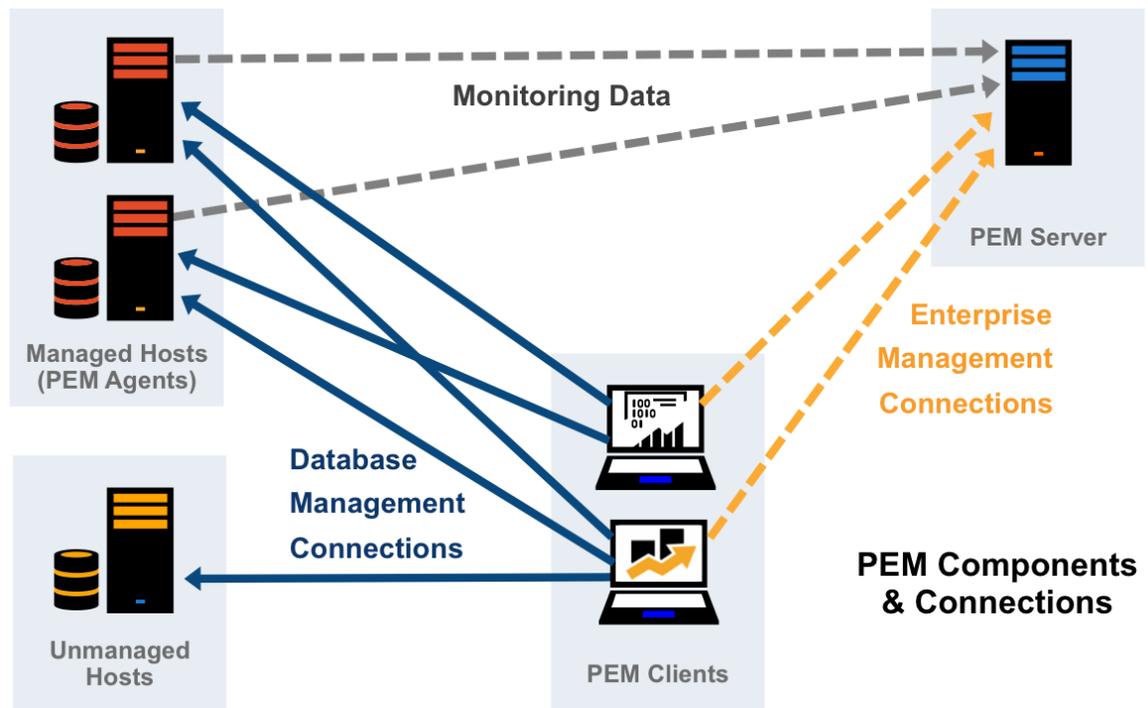


Figure 2.1 - The Postgres Enterprise Manager general architecture.

2.3 Supported Versions and Platforms

The PEM server and client are supported on:

64-bit Windows:

- Windows Server 2012 R2
- Windows Server 2008 R2 Server

64-bit Linux:

- CentOS 6.x and 7.x
- Debian 7.6
- OEL 6.x and 7.x
- Red Hat Enterprise Linux 6.x and 7.x
- SLES 11.x and 12.x
- Ubuntu 14.04

The PEM server is supported on these platforms, using Advanced Server or PostgreSQL backing databases (version 9.1 and higher).

The PEM agent is supported on any Linux or Windows platform on which Advanced Server or PostgreSQL version 9.1 or higher is supported.

The PEM web client is supported on the following browsers:

- Google Chrome
- Mozilla Firefox
- Internet Explorer
- Apple Safari

PostgreSQL Version Support

PEM can manage and monitor:

- PostgreSQL versions 9.1 and higher
- Advanced Server 9.1 and higher

SQL Profiler

SQL Profiler for Postgres Enterprise Manager is supported on Advanced Server EnterpriseDB distributions of PostgreSQL version 9.1 and above.

3 Installing Postgres Enterprise Manager

This section provides an overview of Postgres Enterprise Manager (PEM) installation and configuration; for detailed instructions, please consult the Postgres Enterprise Manager Installation Guide, available at:

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

The basic steps involved in the PEM installation process are:

1. Install the PEM server components. The PEM server software and backend database (named `pem`) may reside on the same host as the supporting ApacheHTTPD server, or may reside on a separate host. You may use an existing Postgres server to host the PEM server, or use the PostgreSQL installer bundled with the PEM server installer to create the backend database.

The PEM server installer also installs a PEM agent on the host of the PEM server.

2. Install a PEM client on the machine from which you will manage your Postgres servers.

You can use the PEM web client (installed by default with the PEM server installer) to manage your PEM installation, or install the full-featured PEM client. Unlike the lighter web-client, the full-featured PEM client allows you to execute SQL commands through a secure SSH tunnel.

3. Optionally, install a PEM agent on each additional physical or virtual machine that you would like to manage or monitor with PEM. PEM 6.0 supports remote monitoring of a server by an agent that does not reside on the same host as the monitored system.

If you do not wish to install an agent on each monitored system, specify a non-resident agent to monitor the system when you register the monitored database with the PEM server.

4. Install the SQL Profiler component into each Postgres instance on which you want to perform SQL capture and analysis. The SQL Profiler installer prompts you for the location of your Postgres installation, and places the required software into that directory. The SQL Profiler plugin is already installed on Advanced Server instances, and requires only configuration to enable profiling.

Graphical installers will lead you through installing and configuring each component of PEM.

3.1 Starting and Stopping the PEM Server and Agents

The PEM Server starts, stops and restarts when the Postgres server instance on which it resides starts, stops or restarts; use the same commands to control the PEM server that you would use to control the Postgres server.

On Linux

On Linux platforms, PEM service scripts reside in the `/etc/init.d` directory. The default name of the service script that controls:

- A PEM server host on Advanced Server is `ppas-9.x`.
- A PEM server host on PostgreSQL is `postgresql-9.x`.
- A PEM agent is `pemagent`.

Where *x* indicates the server version number.

You can use the service script to control the service. Enter:

```
/etc/init.d/service_name action
```

Where:

service_name

service_name is the name of the service.

action

action specifies the action taken by the service command. Specify:

- `start` to start the service.
- `stop` to stop the service.
- `restart` to stop and then start the service.
- `status` to check the status of the service.

On Windows

The Windows operating system includes a graphical service controller that displays the server status, and offers point-and-click server control. The registered name of the service that controls:

- A PEM server host on PostgreSQL is `postgresql-9.x`.
- A PEM server host on Advanced Server is `ppas-9.x`.
- A PEM agent is `pemagent`.

The `Services` utility can be accessed through the Windows Control Panel. When the utility opens, use the scroll bar to navigate through the listed services to highlight the service name (see Figure 8.2).

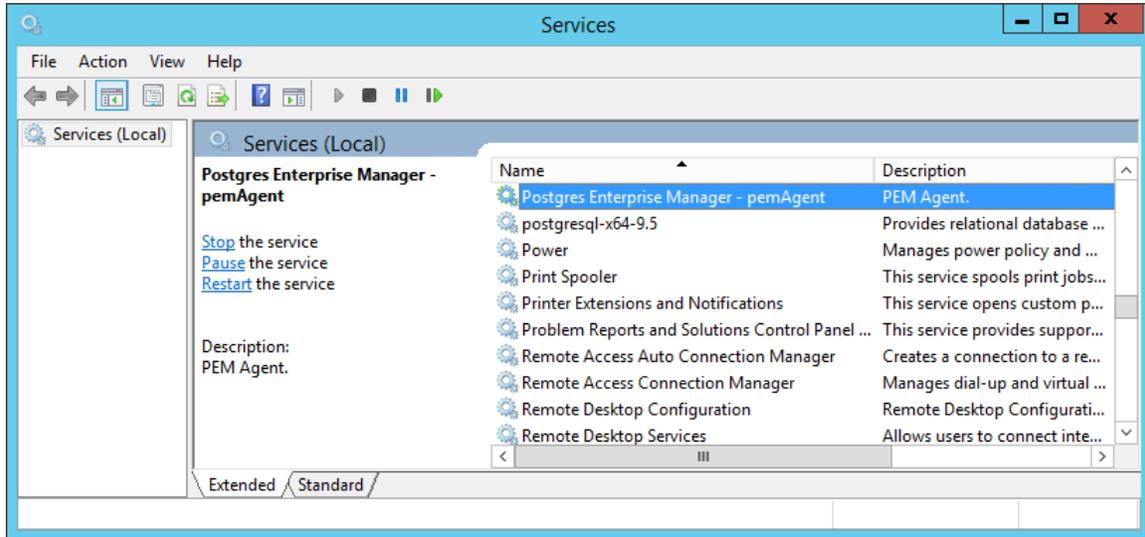


Figure 8.2 – The PEM service in the Windows Services window.

The `Stop the service` option stops the server instance. Any user (or client application) connected to the Postgres server will be abruptly disconnected if you stop the service.

Use the `Pause the service` option to instruct Postgres to reload the server configuration parameters. The `Pause the service` option is an effective way to reset the server parameters without disrupting user sessions for many of the configuration parameters.

Use the `Start the service` option to start the service.

3.2 The PEM Client - User Interface Basics

The main elements of the PEM client interface are the Tabbed Dashboard Browser, the Main Toolbar, and the Tree Control as shown in Figure 3.1.

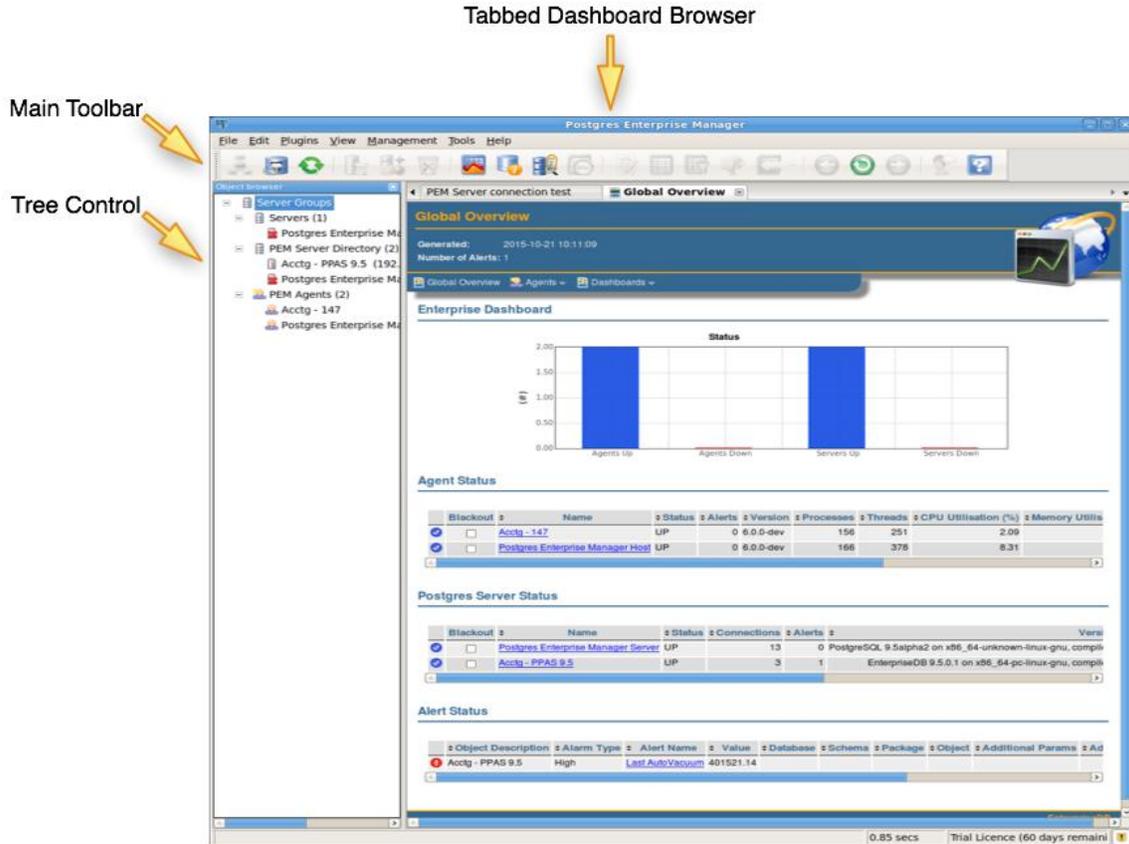


Figure 3.1 - The PEM client interface.

PEM uses a standard tree control to connect to and navigate through the contents of all managed servers. The menus provide easy access to various tasks, and are context sensitive so only those tasks that are appropriate for the selected object are active. The graphical toolbar provides quick access to the most commonly used tasks and utilities.

The right pane of the client interface allows you to use tabbed browsing to review details about selected objects in the tree control. Like most web browsers, you can open multiple tabs for different views, close selected tabs when you're finished reviewing the contents, and navigate through multiple reports on the same tab using back and forward toolbar buttons.

Note that some utilities and interfaces in PEM do not run within the client interface, but instead open in a separate interactive dialog. These include interfaces such as the SQL IDE and the SQL Profiler. This allows for greater flexibility when using these tools on systems with multiple monitors or virtual desktops.

3.3 Using the PEM Web Client

You can use the Postgres Enterprise Manager web client in your browser of choice to maintain your PostgreSQL and Advanced Server databases. The PEM web client can help preserve security on monitored servers by providing access to statistical and status information about objects that reside on a monitored server without requiring an actual connection to each server.

Please note that the PEM web client displays information gathered by the PEM agents, rather than displaying a current view from system catalogs on the monitored servers.

The web client is distributed with, and installed by the PEM server installer. To open the PEM web client, navigate through the Applications menu (on the Linux OS menu) or the Start menu (on the Windows OS menu) to the Postgres Enterprise Manager menu; select PEM Web Client from the pull-aside menu.

The Postgres Enterprise Manager Web Login window (shown in Figure 3.2) opens:

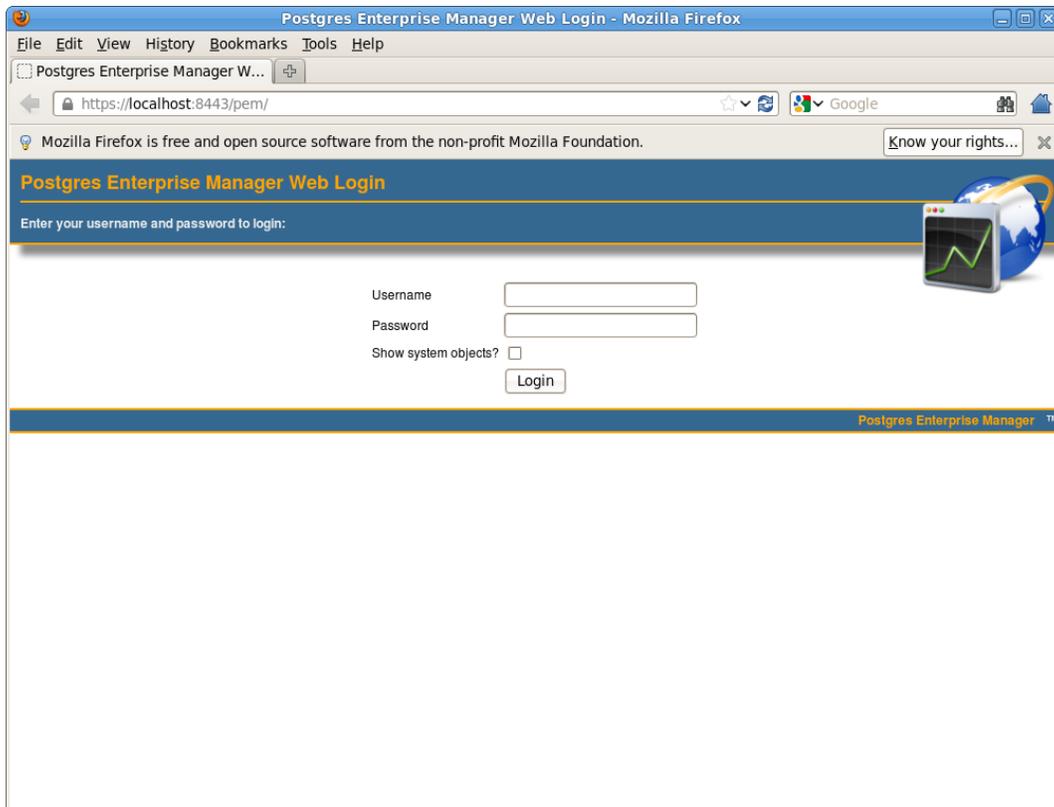


Figure 3.2 - The PEM Web Login page.

Use the fields on the PEM Web Login window to authenticate yourself with the PEM server:

- Provide the name of a PEM user in the `Username` field.
- Provide the password associated with the PEM user in the `Password` field.
- Check the box next to `Show system objects?` to instruct PEM to display the contents, properties and statistics of system databases (such as `template0`) and system schemas (such as `public`) in the tree control.

After providing your credentials, click `Login` to connect to PEM. The main window of the PEM web client opens, displaying the `Global Overview Dashboard` as shown in Figure 3.3.

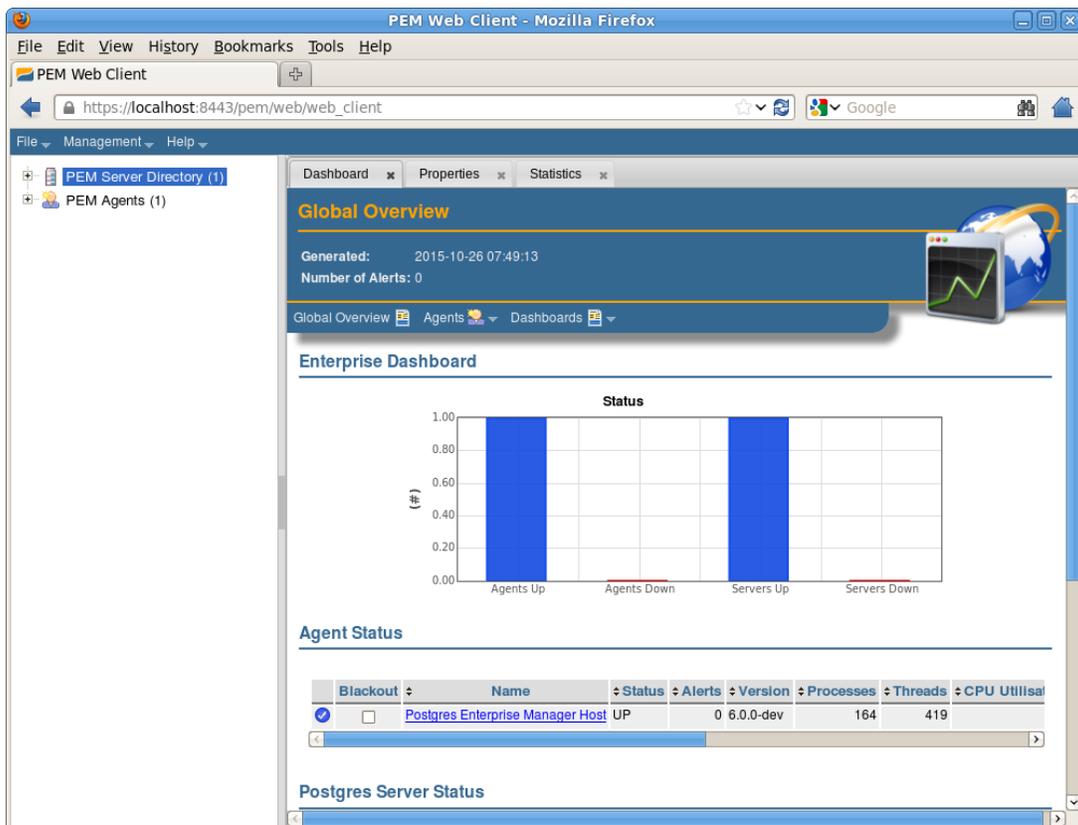


Figure 3.3 - The Global Overview dashboard, displayed in the web client.

Like the full-featured PEM client, the left panel of the web client displays a tree control that provides access to information about the database objects that reside on each server. The tree control expands to display a hierarchical view of the servers and objects that are monitored by the PEM server.

Menu selections displayed across the top of the tree control panel allow access to PEM features and functionality:

- Use options on the `File` menu to add and drop servers, drop PEM agents, change your PEM server password, or log out of the PEM web client.
- Use options on the `Management` menu to invoke PEM wizards and manage PEM features.
- Use options on the `Help` menu menu to access help text for PEM or PostgreSQL, or to review version information about PEM and supporting software.

The main panel of PEM web client displays a set of tabs; each tab displays a different collection of information about the object currently selected in the tree control.

- Open the `Properties` tab to display the properties of the item currently highlighted in the tree control.
- Open the `Statistics` tab to display usage statistics (if applicable) for the object currently highlighted in the tree control.
- Open the `Dashboard` tab to access information presented on PEM dashboards. Dashboards display statistical information about the objects monitored by the PEM server.

Navigation menus displayed in the dashboard header provide easy access to other dashboards. Menus are organised hierarchically; only those menus appropriate for the object currently highlighted in the tree control are available:

- Select `Global Overview` from any dashboard to return to the `Global Overview` dashboard.
- Select the name of an agent from the `Agents` menu to navigate to the `Operating System Analysis` dashboard for that agent.
- Select a server name from the `Servers` menu to navigate to the `Server Analysis` dashboard for that server.
- Select a database name from the `Databases` menu to navigate to the `Database Analysis` dashboard for that database.
- Use the `Dashboards` menu to navigate to informational dashboards at the global level, or for the selected agent, server or database.

3.4 Online Help and Documentation

PEM contains built-in help that you can reference for assistance in using the tool (see Figure 3.2). To access the full online help for PEM, select `Help` contents from the `Help` option on the main menu.

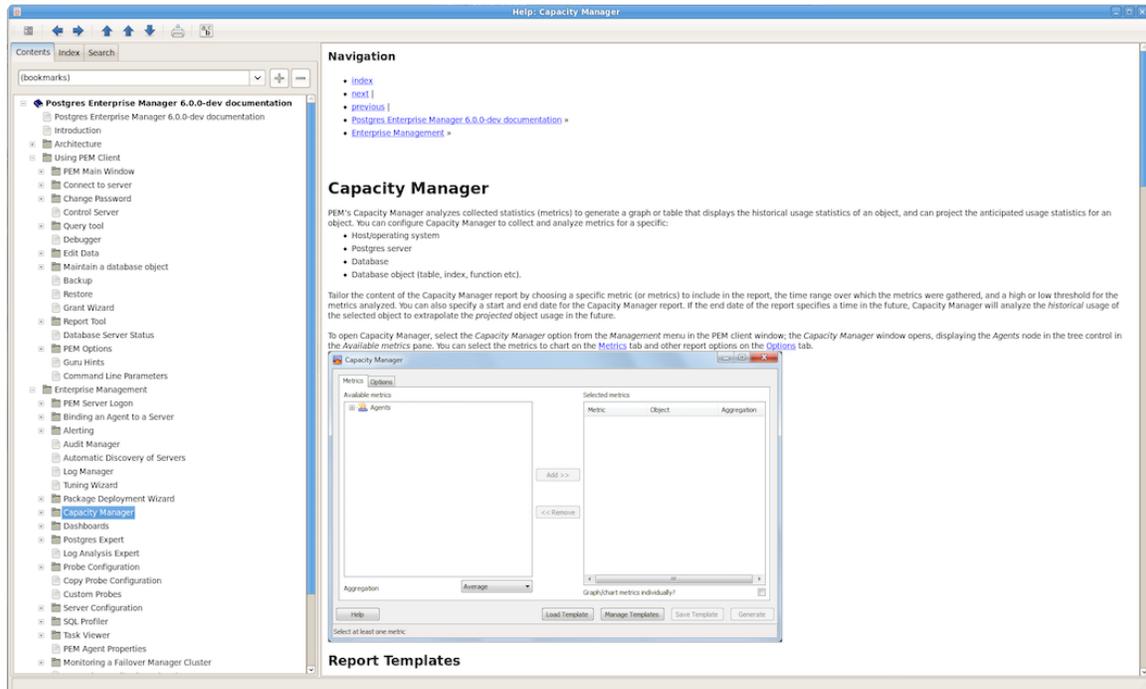


Figure 3.2 - PEM's Help dialog.

3.5 Logging on to the PEM Server

To logon to the PEM server, navigate through the `File` menu, and select the `Enterprise Manager Logon` menu option, or click the `Logon` toolbar button (as shown in Figure 3.3).



Figure 3.3 - PEM's Logon Button.

When the `PEM Server Logon` dialog opens, provide the PEM server's IP address or host name, and the user name and password specified during installation. A successful login will prompt the PEM client to display a new node in the control tree labeled `PEM Server Directory`, and a node called `PEM Agents (PEM agents)`.

After logging in to the PEM server, you can add new servers to manage.

Note: `ApacheHTTPD` must be running in order for the PEM client to connect to the PEM server.

On Linux, you can confirm the status of the `Apache` service by opening a command line, and entering the following command:

```
ps -ef | grep apache
```

If Linux responds with an answer that is similar to the following example, `ApacheHTTPD` is not running:

```
user      13321 13267  0 07:37 pts/1    00:00:00 grep apache
```

You can use the following command to start `ApacheHTTPD`:

```
/etc/init.d/EnterpriseDBApacheHTTPD start
```

If `ApacheHTTPD` starts properly, the `ps -ef` command will result in output similar to the following:

```
$ # ps -ef | grep apache
root      24924      1  0 07:50 ?                00:00:00 /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/bin/httpd -k start -f /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/conf/httpd.conf
```

```

daemon 24925 24924 0 07:50 ?        00:00:00 /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/bin/httpd -k start -f /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/conf/httpd.conf
daemon 24926 24924 0 07:50 ?        00:00:00 /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/bin/httpd -k start -f /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/conf/httpd.conf
daemon 24927 24924 0 07:50 ?        00:00:00 /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/bin/httpd -k start -f /opt/PostgreSQL/EnterpriseDB-
ApacheHTTPD/apache/conf/httpd.conf
root    25237 22923 0 07:52 pts/0    00:00:00 grep apache
    
```

On Windows, you can use the Services applet to check the status of the ApacheHTTPD service.

To confirm the status of the ApacheHTTPD service, navigate through the Control Panel, to System and Security; select Administrative Tools, and then Services. The ApacheHTTPD server runs as a service named EnterpriseDB-ApacheHTTPD (see Figure 8.3).

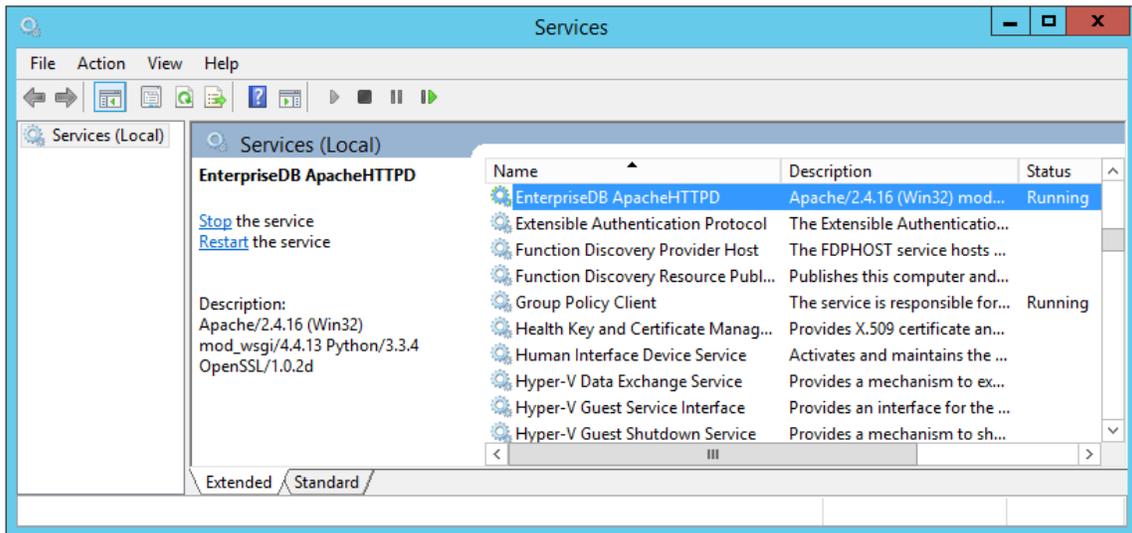


Figure 3.4 - The EnterpriseDB ApacheHTTPD Windows service.

The Status column displays the current state of the ApacheHTTPD server. Click the Start link to start ApacheHTTPD if the service is not running.

3.6 Adding a Managed Server

A server definition may be configured locally or in the PEM Server Directory:

- Local servers typically reside on the same machine as the PEM client, and are considered to be unmanaged by PEM.
- Servers residing in the PEM Server Directory are considered to be managed, and are available to all PEM users. Servers in the PEM directory will normally be bound to an agent to enable monitoring and other PEM functionality.

The process of configuring a server, either managed or unmanaged, is referred to as adding or *registering* a server.

Please note: You must ensure the `pg_hba.conf` file of the server that you are registering allows connections from the host of the PEM client.

To open the New Server Registration dialog, click on the Add Server button on the client toolbar, or select Add Server from the File menu.

The screenshot shows the 'New Server Registration' dialog box with the 'Properties' tab selected. The dialog contains the following fields and options:

- Name:** An empty text input field.
- Host:** An empty text input field.
- Port:** A text input field containing the value '5432'.
- Service:** An empty text input field.
- Maintenance DB:** A dropdown menu with 'postgres' selected.
- Username:** A text input field containing the value 'postgres'.
- Password:** An empty password input field.
- Store password:** A checked checkbox.
- Store on PEM Server:** A checked checkbox.
- Colour:** A color selection field with a small color swatch.
- Group:** A dropdown menu with 'PEM Server Directory' selected.
- Team:** An empty text input field.

At the bottom of the dialog are three buttons: 'Help', 'OK', and 'Cancel'.

Figure 3.5 - The Properties tab of the New Server Registration dialog.

The Properties Tab

Use fields on the `Properties` Tab (shown in Figure 3.5) to enter the connection details for the new server.

- Enter a descriptive name for the server in the `Name` field; the descriptive name will be displayed in the tree control.
- Enter the host name or IP address in the `Host` field.
- Enter the listener port number of the PostgreSQL or Advanced Server host in the `Port` field.
- Use the `Service` field to specify the name of a service configured in the `pg_service.conf` file.
- Select a maintenance database with the drop-down list box in the `Maintenance DB` field.
- Specify the `Username` that the new server will use when connecting to the database; the specified user must have membership in the `pem-admin` role.
- Enter the password associated with the user in the `Password` field.
- Check the box next to `Store password` to instruct the PEM client to store the password for future connections.
- Check the box next to `Store on PEM Server` to instruct the PEM client to store the definition of the server on the PEM server.

If the connecting role is not a member of `pem_admin` (they are connected using the read-only `pem_user` role), or the user has not logged in to the PEM server, the server definition is stored for the current user only.

- Use the `Colour` selector to specify a background display color for the server.
- Select a server group for the new server using the drop-down list box in the `Group` field. The new server will be displayed in the selected group in the PEM client tree control.
- Use the `Team` field to specify a Team role name. Only PEM users who are members of this role, who created the server initially, or have superuser privileges on the PEM server will see this server when they log on to PEM. If this field is left blank, all PEM users will see the server. For more information about defining a Team role, see Section [4.2.3](#).

Tip: Do not click **OK** (yet) if you wish to register the server as a managed server. To register the server as a managed server, you must complete the **PEM Agent** tab, binding the server to a PEM agent.

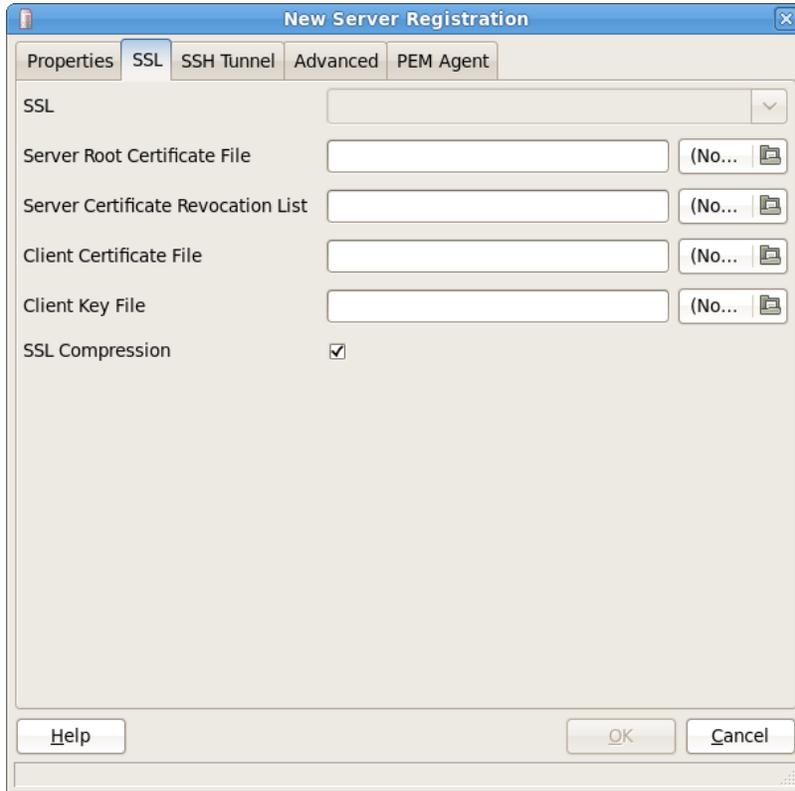


Figure 3.6 - The SSL tab of the New Server Registration dialog.

The SSL Tab

Use fields on the **SSL** tab (shown in Figure 3.6) to specify any SSL options required.

- Use the drop-down list box in the **SSL** field to select an SSL mode:

Specify:	
require	To require SSL encryption for transactions between the server and the agent.
prefer	To use SSL encryption between the server and the agent if SSL encryption is available.
allow	To allow the connection to use SSL if required by the server.
disable	To disable SSL encryption between the agent and the server.
verify-ca	To require SSL encryption, and to require the server to authenticate using a certificate registered by a certificate

	authority.
verify-full	To require SSL encryption, and to require the server to authenticate using a certificate registered by a <i>trusted</i> certificate authority.

For more information about using SSL encryption, see Section 31.17 of the Postgres documentation.

- If applicable, use the `Server Root Certificate File` field to access a file browser, and specify the location of the server root certificate.
- If applicable, use the `Server Certificate Revocation List` field to access a file browser, and specify the location of the certificate revocation list.
- If applicable, use the `Client Certificate File` field to access a file browser, and specify the location of the client certificate.
- If applicable, use the `Client Key File` field to access a file browser, and specify the location of the client key file.
- Check the box next to `SSL Compression` to instruct the server to compress the SSL data stream.

For more information about implementing SSL support, see Chapter 31.17 of the Postgres Core Documentation, available at the EnterpriseDB website at:

<http://www.postgresql.org/docs/9.5/static/libpq-ssl.html>

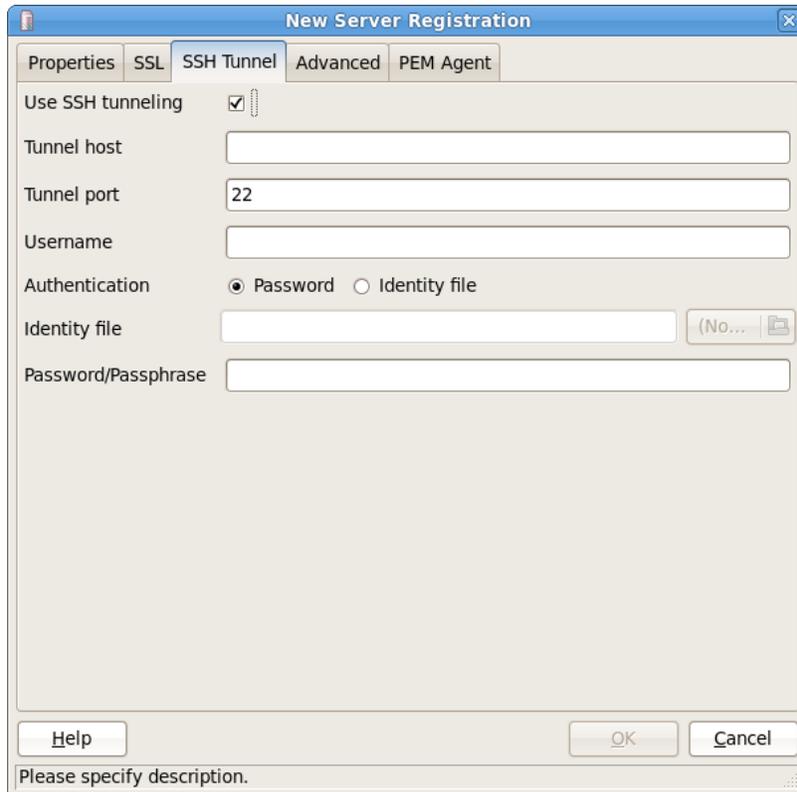


Figure 3.7 - The SSH Tunnel tab of the New Server Registration dialog.

The SSH Tunnel Tab

Use the fields on the SSH Tunnel tab (shown in Figure 3.7) to implement SSH tunneling to the new server; check the box next to Use SSH tunneling to enable the fields:

- Check the box next to Use SSH tunneling to instruct PEM to use SSH tunneling to communicate with the monitored server.
- Specify the name or IP address of the proxy host (through which client connections will be forwarded) in the Tunnel host field.
- Specify the port that should be used for the tunnel in the Tunnel port field.
- Specify the name of a user with connection privileges for the tunnel host in the Username field.
- Specify the type of authentication that will be used when connecting to the tunnel host in the Authentication field.

Select the radio button next to `Password` to specify that the PEM client will provide a password for authentication by the tunnel host. This is the default.

Select the radio button next to `Identity file` to specify that the PEM client will provide a valid private key file when connecting.

- If the tunnel host is expecting a private key file for authentication, use the `Identity file` field to specify the location of the key file.
- If the tunnel host is expecting a password, use the `Password/Passphrase` field to specify a word or phrase that will be provided to the tunnel host for authentication.

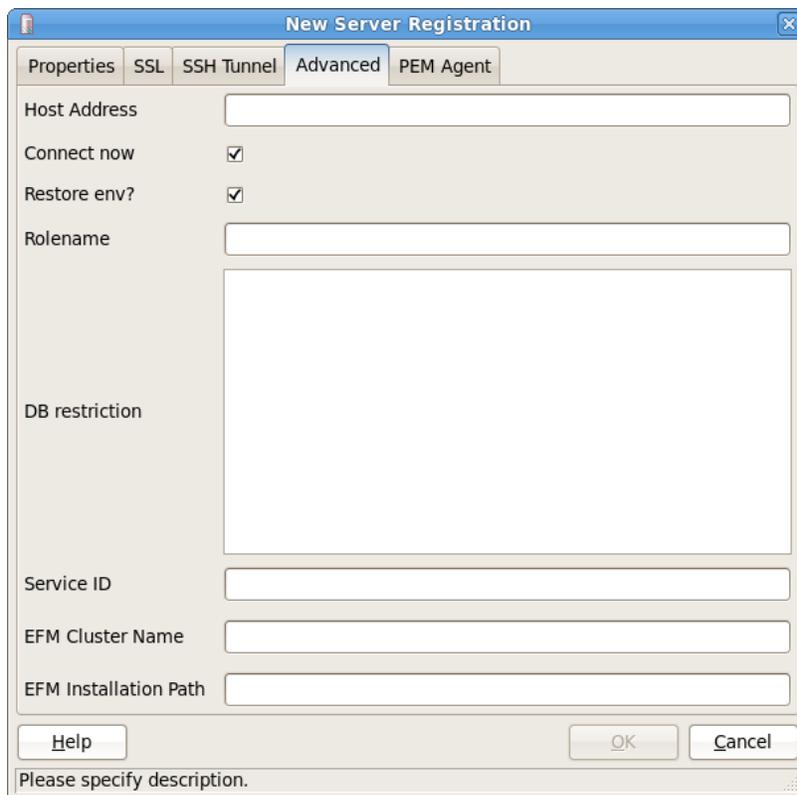


Figure 3.8 - The Advanced tab of the New Server Registration dialog.

The Advanced Tab

Use the fields on the `Advanced` tab (shown in Figure 3.8) to specify connection information for the new server:

- Specify the address of the host in the `Host Address` field.

- Check the box next to `Connect now` to instruct the PEM client to negotiate a connection for the new server after you have completed the `New Server Registration` form and clicked OK.
- Check the box next to `Restore env?` to instruct PEM to restore the browser environment when you reconnect to this server. If you regularly use different databases on the same server you may want to deselect this option.
- If you specify a role name in the `Rolename` field, PEM will connect to the server using the identity specified on the `Properties` tab; after establishing a connection with the server, the client will assume the identity and permissions of the role specified in the `Rolename` field. The role specified on the `Properties` tab must be a member of the role specified on the `Advanced` tab.
- Specify an SQL restriction in the `DB restriction` field to limit the databases displayed in the PEM client. For example, enter: `'live_db', 'test_db'` to display only `live_db` and `test_db` in the PEM browser.
- Specify the name of the PostgreSQL or Advanced Server service in the `Service ID` field on the `Advanced` tab. This allows the PEM server to stop and start the service.

On Unix systems, provide the name of the service script located in `/etc/init.d`

On Windows, provide the ID of the service. You can find the service ID in the `Services Microsoft Management Console` application.

The Advanced Server installer uses a default service ID of `ppas-9.x`, where `x` specifies the version number of the server. For example, the service ID of Advanced Server 9.5 is `ppas-9.5`

The EnterpriseDB one-click installer (for PostgreSQL) uses a default service ID of `postgresql-9.x` where `x` specifies the version number of the server. For example, the service ID of Advanced Server 9.5 is `postgresql-9.5`

- If the server is a member of a Failover Manager cluster, you can use PEM to monitor the health of the cluster and to replace the master node if necessary. To enable PEM to monitor Failover Manager, use the `EFM Cluster Name` field to specify the cluster name. The cluster name is the prefix of the name of the Failover Manager cluster properties file. For example, if the cluster properties file is named `efm.properties`, the cluster name is `efm`.
- If you are using PEM to monitor the status of a Failover Manager cluster, use the `EFM Installation Path` field to specify the location of the Failover Manager

binary file. By default, the Failover Manager binary file is installed in `/usr/efm-2.0/bin`.

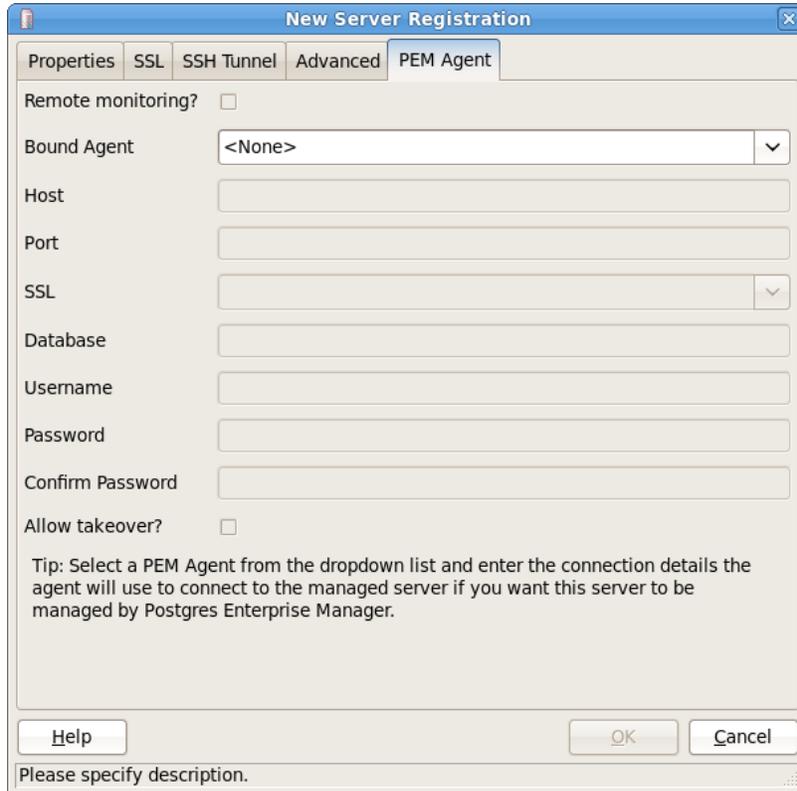


Figure 3.9 - The PEM Agent tab of the New Server Registration dialog.

The PEM Agent Tab

Use fields on the `PEM Agent` tab (shown in Figure 3.9) to create a *binding* between the new server and a PEM agent. A binding provides the link between the PEM server and the PEM agent that allows the PEM server to monitor and manage that agent. The binding also provides the connection parameters that the agent uses when connecting to the PEM server.

It is important to note that the agent may use different connection credentials when connecting to the monitored server than the client uses. For example, the agent on a Unix host may use a Unix Domain Socket to connect to the server, while the client may connect via a connection pooler on an entirely separate host.

- Check the box next to `Remote monitoring ?` to indicate that the PEM agent does not reside on the same host as the monitored server. When remote monitoring is enabled, agent level statistics for the monitored server will not be available for custom charts and dashboards, and the remote server will not be

accessible by some PEM utilities (such as Audit Manager, Capacity Manager, Log Manager, Postgres Expert, and Tuning Wizard).

- Select an Enterprise Manager agent using the drop-down list box to the right of the `Bound Agent` label.
- Enter the IP address that the agent should use when connecting to the Advanced Server host in the `Host` field. You may wish to specify a different `Host` address on the PEM agent tab than the address used on the `Properties` tab. For example, you can configure the PEM agent to connect via a Unix socket to the monitored server, while the client connects using an SSL connection to a forwarded port on an external firewall (that is not appropriate for the agent).
- Enter the port number that the server is monitoring for connections in the `Port` field. By default, the agent will use the port defined on the `Properties` tab.
- Use the `SSL` field to specify an SSL operational mode; specify `require`, `prefer`, `allow`, `disable`, `verify-ca` or `verify-full`.

For more information about using SSL encryption, see Section 31.17 of the Postgres documentation.

- Use the `Database` field to specify the name of the database to which the agent will initially connect.
- Specify the name of the role that agent should use when connecting to the server in the `Username` field.
- Specify the password that the agent should use when connecting to the server in the `Password` field.
- Verify the password by re-entering it in the `Confirm Password` field. If you do not specify a password, you must configure the authentication for the agent manually (for example, by using a `.pgpass` file).
- Use the `Allow takeover` option to specify whether or not the monitoring of the server may be taken over by another agent in a High Availability environment.

The PEM client will connect directly to the managed server, so the `pg_hba.conf` file must contain the appropriate entries to allow those connections. For information about modifying the `pg_hba.conf` file, see

<http://www.postgresql.org/docs/9.5/static/auth-pg-hba-conf.html>

When you've completed the `New Server Registration` dialog, click `OK`. The name of the new server will be included under the `PEM Server Directory` node in the tree control.

3.7 Remotely Starting and Stopping Monitored Servers

PEM allows you to startup and shutdown monitored server instances on remote machines from the PEM client. Each remote server must be properly configured to allow the PEM client to stop or restart the server. To use the PEM client to control the startup or shutdown of a configured server, complete the `New Server Registration` dialog, registering the database server with PEM:

- Specify the `Store on PEM Server` option on the `Properties` tab.
- Specify the name of a service script in the `Service ID` field on the `Advanced` tab:

For `Advanced Server`, the service name is `ppas-9.x`.

For `PostgreSQL`, the service name is `postgresql-9.x`.

To connect to the newly-defined server, right click the server name in the tree control, and select `Connect` from the context menu. Provide a password when prompted.

To start or stop the server, right click the server name in the tree control, and select the `Queue Server Startup` or `Queue Server Shutdown` (shown in Figure 3.5) from the context menu.

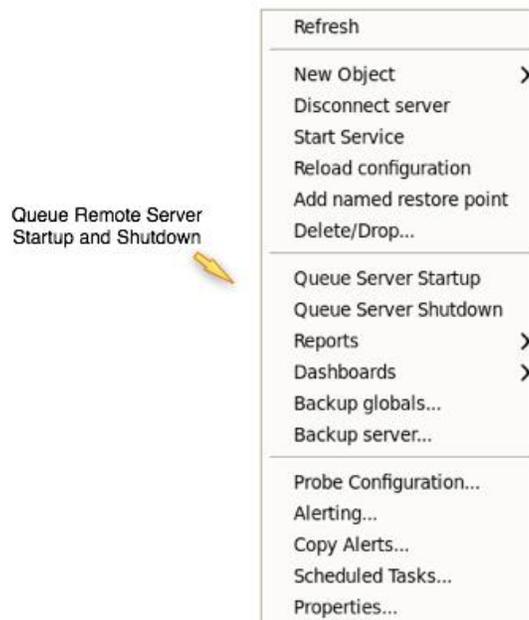


Figure 3.5 - The context menu of a remotely managed server.

Note that currently, PEM only supports the fast shutdown option of the database server.

4 General Database Administration

Postgres Enterprise Manager is based on the most popular GUI tool for PostgreSQL, pgAdmin. If you are already familiar with pgAdmin, you may wish to skip this chapter - it covers some of the basic administrative tasks that can be carried out with PEM, most of which are also features in pgAdmin.

4.1 Editing a Server's Configuration

You can use the PEM client to graphically manage the configuration parameters of a remote Postgres server:

1. Right click on the name of a monitored server in the tree control, and select **Connect** from the context menu. If prompted, provide a password to connect to the server.
2. Navigate through the **Tools** menu to the **Server Configuration** sub-menu, and select the configuration file you wish to edit.

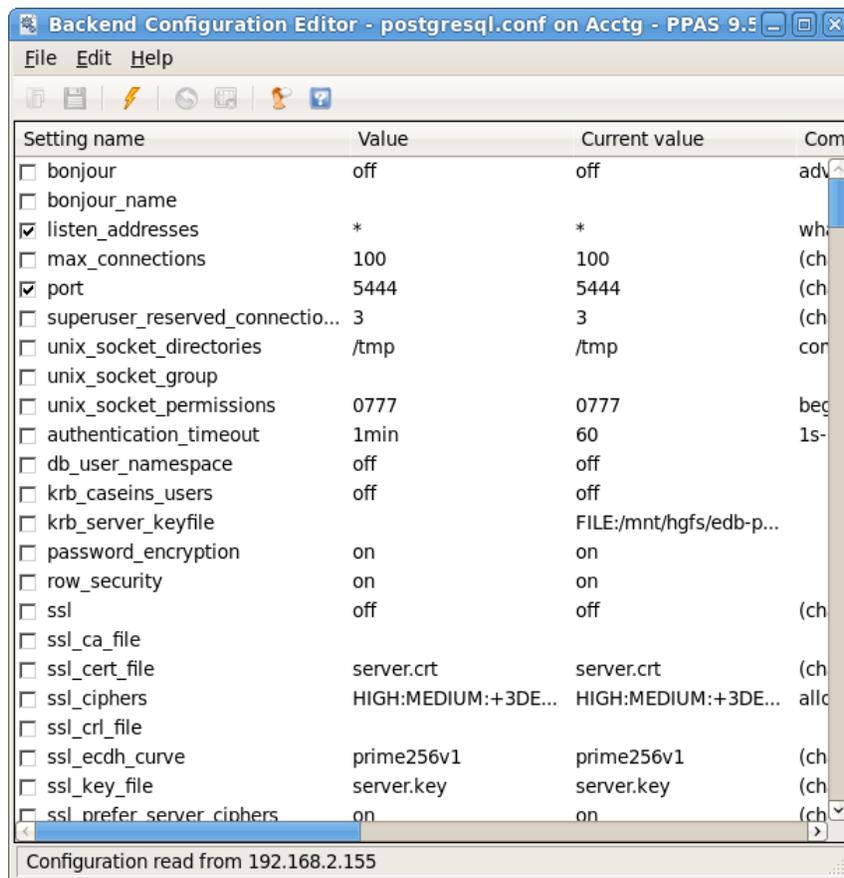


Figure 4.1 - The Configuration Editor dialog.

You can use the `Configuration Editor` (shown in Figure 4.1) to display the contents of the `postgresql.conf` file or `pg_hba.conf` file for the currently selected server. To edit a parameter value, double-click on the parameter name. When you have made any desired changes, you can save the file on the remote server by selecting `Save` from the `File` menu, or by clicking the `Save` toolbar icon.

Warning: *Specifying invalid values for parameters may prevent your Postgres server from starting.*

After saving the configuration file, you must reload the server configuration. To reload the configuration files, navigate through the `Tools` menu and select `Reload configuration`.

Note that some parameter value changes are not dynamic and will not take effect unless the server is stopped and restarted.

4.2 Managing Security

PEM provides a graphical way to manage the security aspects of your Postgres servers. The three most common tasks are:

- The creation and maintenance of login roles
- The creation and maintenance of group roles
- Administering object permissions

4.2.1 Login Roles

A user must have a login account to connect to the Postgres server. Use the `Login Role` dialog (shown in Figure 4.2) to add a new login role or manage the properties of an existing login role on a registered server.

To add a new login role, right click on the `Login Roles` node (located beneath the selected server in the Postgres Enterprise Manager node of the tree control), and select `New Login Role` from the context menu.

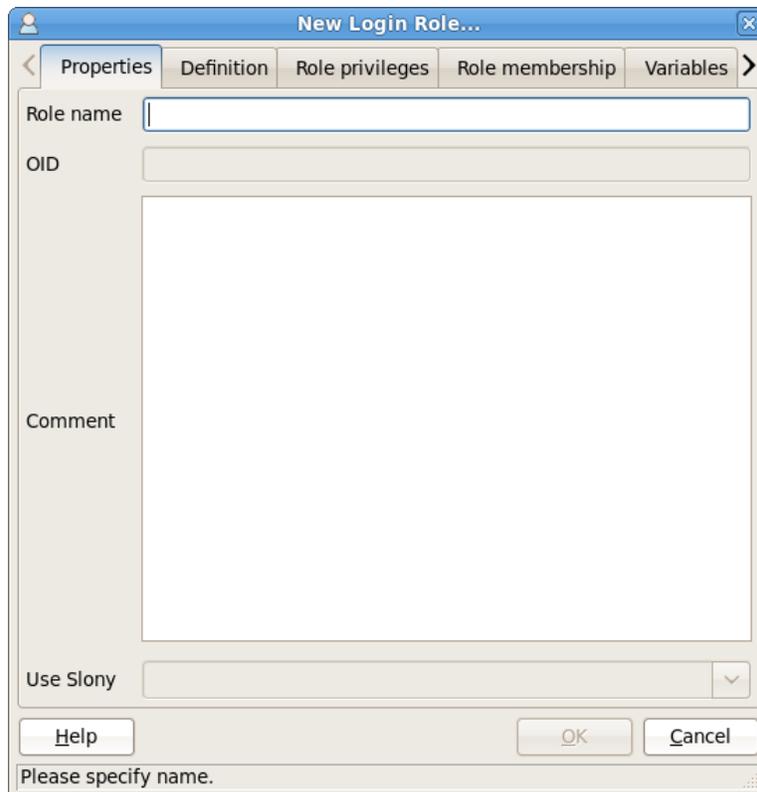


Figure 4.2 - The Login Role dialog

To modify the properties of an existing login role, right click on the name of a login role in the tree control, and select `Properties` from the context menu. To delete a login role, right click on the name of the role, and select `Delete/Drop` from the context menu.

For more complete information on creating and managing a login account, see the PostgreSQL online documentation:

<http://www.postgresql.org/docs/9.5/static/sql-createrole.html>

4.2.2 Group Roles

Group roles can serve as containers, used to dispense system privileges (such as creating databases) and object privileges (e.g. inserting data into a particular table). The primary purpose of a group role is to make the mass management of system and object permissions much easier for a DBA. Rather than assigning or modifying privileges individually across many different login accounts, you can assign or change privileges for a single role and then grant that role to many login roles at once.

Use the `Group Roles` node (located beneath the name of each registered server in the PEM tree control) to create and manage group roles. Options on the context menu provide access to a dialog that allows you to create a new role or modify the properties of an existing role. You can find more information about creating roles at:

<http://www.postgresql.org/docs/9.5/static/sql-createrole.html>

4.2.3 Using a Team Role

When you register a server for monitoring by PEM, you can specify a *Team role* that will be associated with the server. A Team role is a group role that can be used to allow or restrict access to one or more monitored servers to a limited group of role members. The PEM client will only display a server with a specified `Team` to those users who are:

- a member of the Team role
- the role that created the server
- a role with superuser privileges on the PEM server.

To open the `New Group Role` dialog and create a team role, right-click on the `Group Roles` node of the tree control and select `New Group Role...` from the context menu. When the `New Group Role` dialog opens, use the fields provided to specify the properties of the team role. For more information about creating a Team role, see the PEM Installation Guide, available at:

<http://www.enterprisedb.com/download-postgres-enterprise-manager>

4.2.4 Object Permissions

A role must be granted sufficient privileges before accessing, executing, or creating any database object. PEM allows you to assign (`GRANT`) and remove (`REVOKE`) object permissions to group roles or login accounts using the graphical interface of the PEM client.

Object permissions are managed via the graphical object editor for each particular object. For example, to assign privileges to access a database table, right click on the table name in the tree control, and select the `Properties` option from the context menu. Use the options displayed on the `Privileges` tab to assign privileges for the table.

By default, PEM displays only group roles on the `Privileges` tab of the `Properties` dialog. To instruct the PEM client to include login roles in the `User/Group` list on the `Privileges` tab, navigate through the `File` menu, to open the `Options` dialog. Select the `UI Miscellaneous` control node, and then check the box next to `Show users for privileges` to include login roles on the `Privileges` tab.

The PEM client also contains a `Grant Wizard` (accessed through a schema node of the tree control) that allows you to manage many object permissions at once.

4.3 Managing Storage

PostgreSQL uses a named *tablespace* to define an on-disk location (a physical container) in which to store system and user data. Each PostgreSQL host may contain a single tablespace or multiple tablespaces. The PEM client provides a graphical interface that allows you to create and manage PostgreSQL tablespaces.

Use the `Tablespaces` node in the PEM tree control to create and manage tablespaces. The `Properties` editor for tables and indexes allows you to specify the tablespace in which a table or index should reside.

For more information about tablespaces, see the online documentation:

<http://www.postgresql.org/docs/9.5/static/manage-ag-tablespaces.html>

4.4 Creating and Maintaining Databases and Objects

Each instance of the Postgres server manages one or more databases; each user must provide authentication information to connect to the database before accessing the information contained within it. The PEM client provides dialogs that allow you to create and manage databases, and all of the various objects that comprise a database (e.g. tables, indexes, stored procedures, etc.).

Creating a database is easy in PEM: simply right click on any managed server's `Databases` node and select the `New Database...` menu option. You can also access the `New Database` dialog by navigating through the `Edit` menu (on the `Main` menu) to the `New Object` menu, and selecting `New Database`.

Once you have defined a database, you can create objects within the new database. Note that within each database there exist one or more *schemas*. A schema can be thought of as a directory of an operating system disk; it allows the logical separation of database objects (i.e. tables, indexes, SQL functions, and more) inside of a database.

PEM provides graphical dialogs for the creation and maintenance of all supported objects:

- tables
- indexes
- stored procedures
- functions
- triggers
- views
- constraints, etc.

If you are using EnterpriseDB's Advanced Server, you can also use the PEM client to create and manage objects (such as packages) compatible with Oracle databases.

Each managed object is displayed in the tree control. Right click on a named node and use the context menu (or navigate through the top level menu) to perform administrative tasks for the highlighted object.

For example, to create a new table, right click on a `Table` node, select `New Table...` from the context menu. When the `New Table` dialog opens, specify the attributes of the new table (see Figure 4.3).

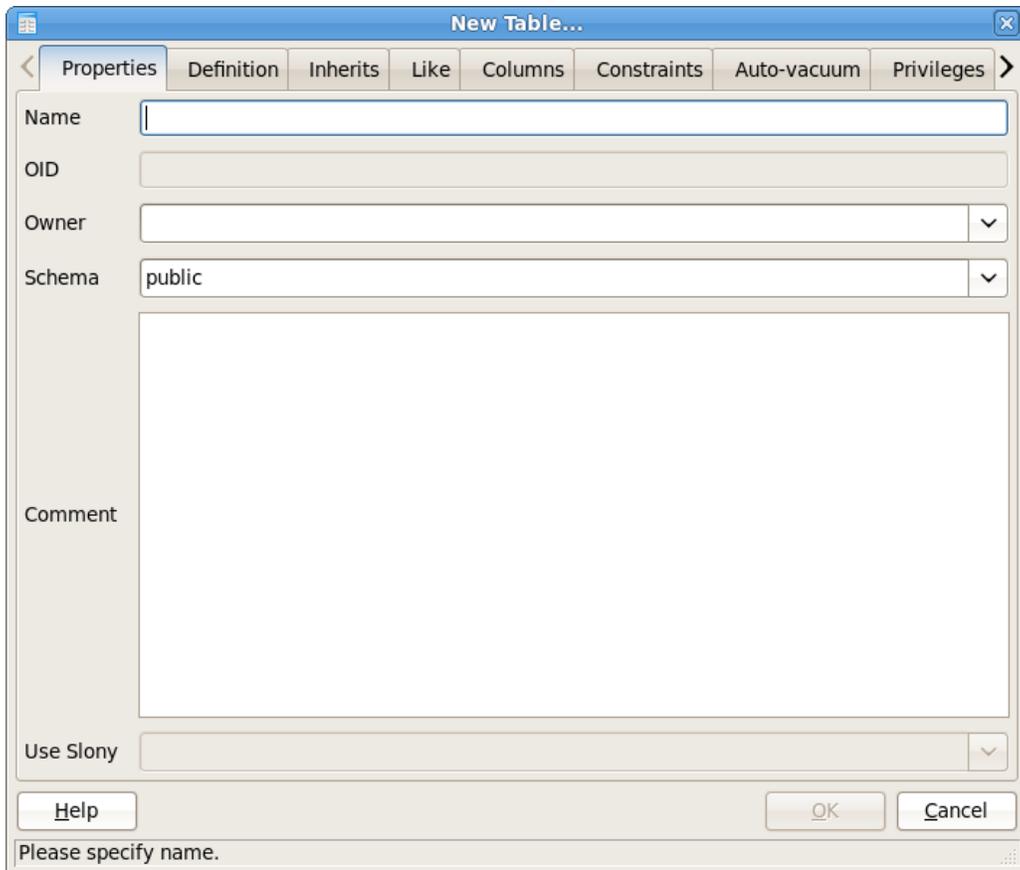


Figure 4.3 - Use PEM's dialogs to create and manage database objects.

PEM provides similar dialogs for the creation and management of other database objects.

4.5 SQL Development

PEM contains a feature-rich Interactive Development Environment (IDE) that allows you to issue ad-hoc SQL queries against Postgres servers.

To invoke the `Query Tool SQL IDE` from within PEM, simply highlight the name of the database you want to query in the tree control, and click the SQL toolbar icon. You can also open the `Query Tool IDE` by selecting `Query tool` from the `Tools` menu.

The `Query Tool` dialog provides an interface that allows you to manually enter in SQL queries, graphically create and execute SQL statements from dragging and dropping objects onto a visual palette, `EXPLAIN` queries and much more.

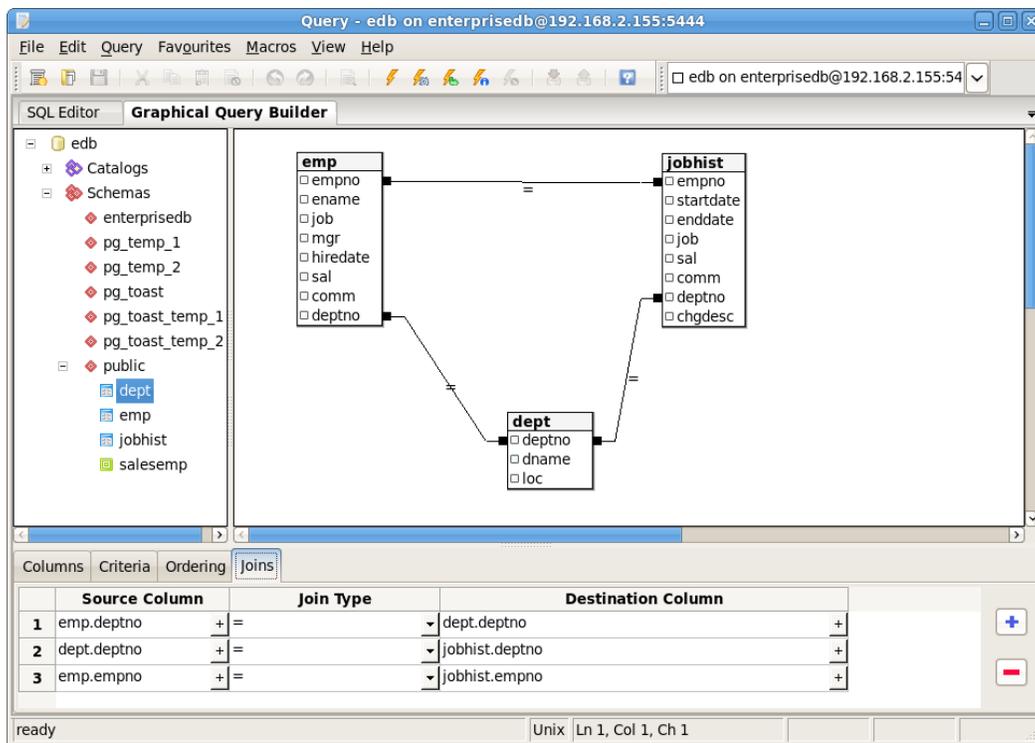


Figure 4.4 - The SQL IDE's Graphical Query Builder.

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 autocompletion code and formatting features that help you write queries.

If you prefer to build your queries graphically, you can use the `Graphical Query Builder` (shown in Figure 4.4) to generate SQL commands. After constructing a graphical query, click the `Execute` button to display the query text in the `SQL Editor` tab, and the results of the query in the `Output` pane (shown in Figure 4.5).

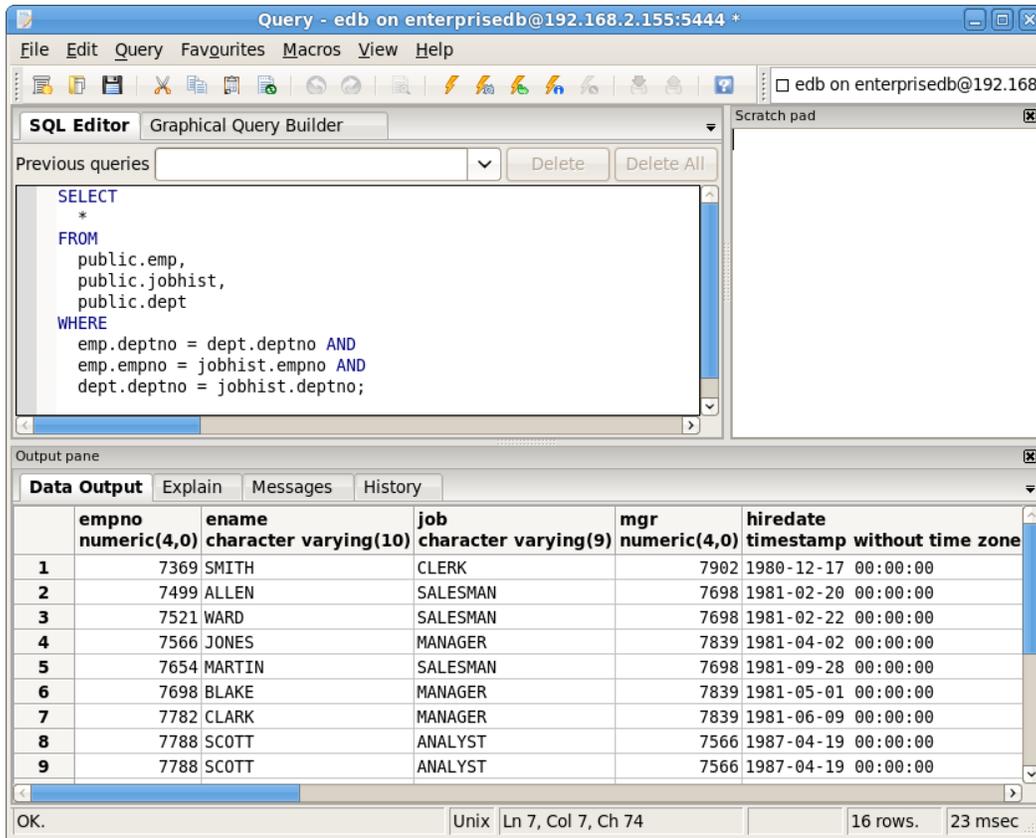


Figure 4.5 - The SQL Editor pane displays the text of the query.

You can manually adjust the query in the SQL Editor pane. As with all PEM features, online Help text is available with the click of a button.

5 Package Deployment

The Package Deployment wizard walks you through the process of scheduling the installation of new packages or upgrades of existing packages. 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 must also enable agent-level probes on the host of the PEM server, and on any system on which a package will be deployed. To access the probe configuration, highlight the name of the PEM agent in the PEM client tree control, and select `Probe Configuration...` from the Management menu to open the Probe Configuration dialog. Enable:

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

To open the Package Deployment wizard, select `Package Deployment...` from the Management menu. The Package Deployment wizard `Welcome...` dialog opens (see Figure 5.1).

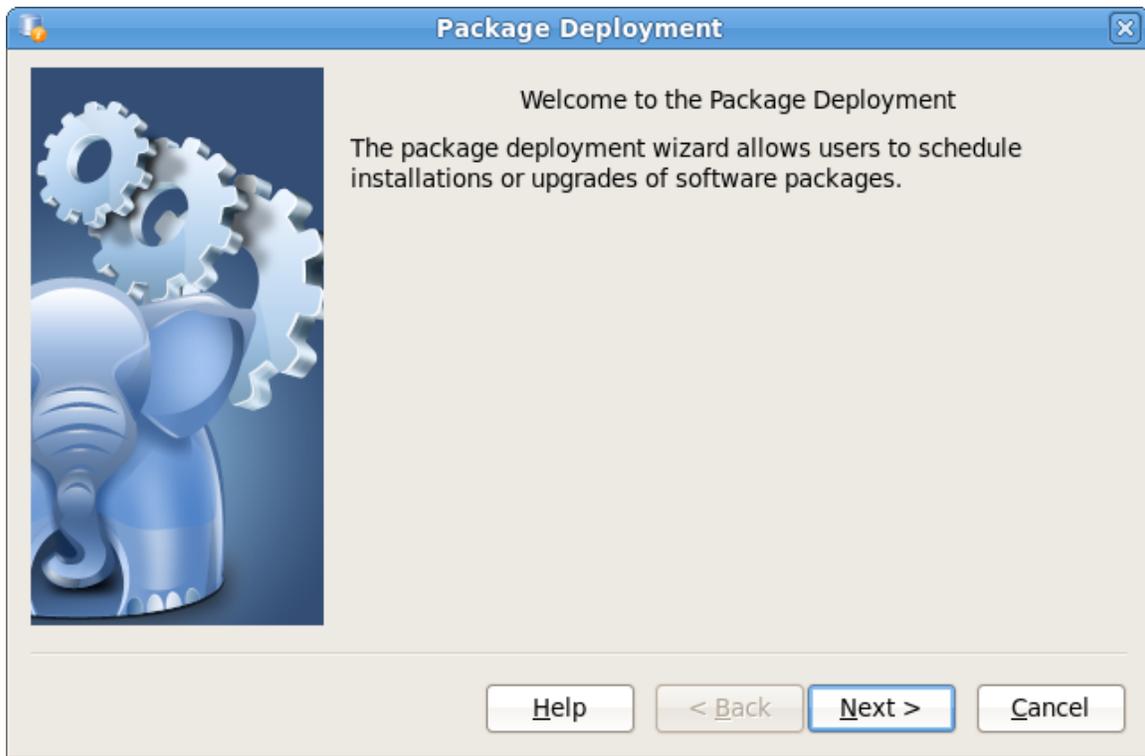


Figure 5.1 - The Package Deployment Welcome dialog.

Click `Next` to continue.

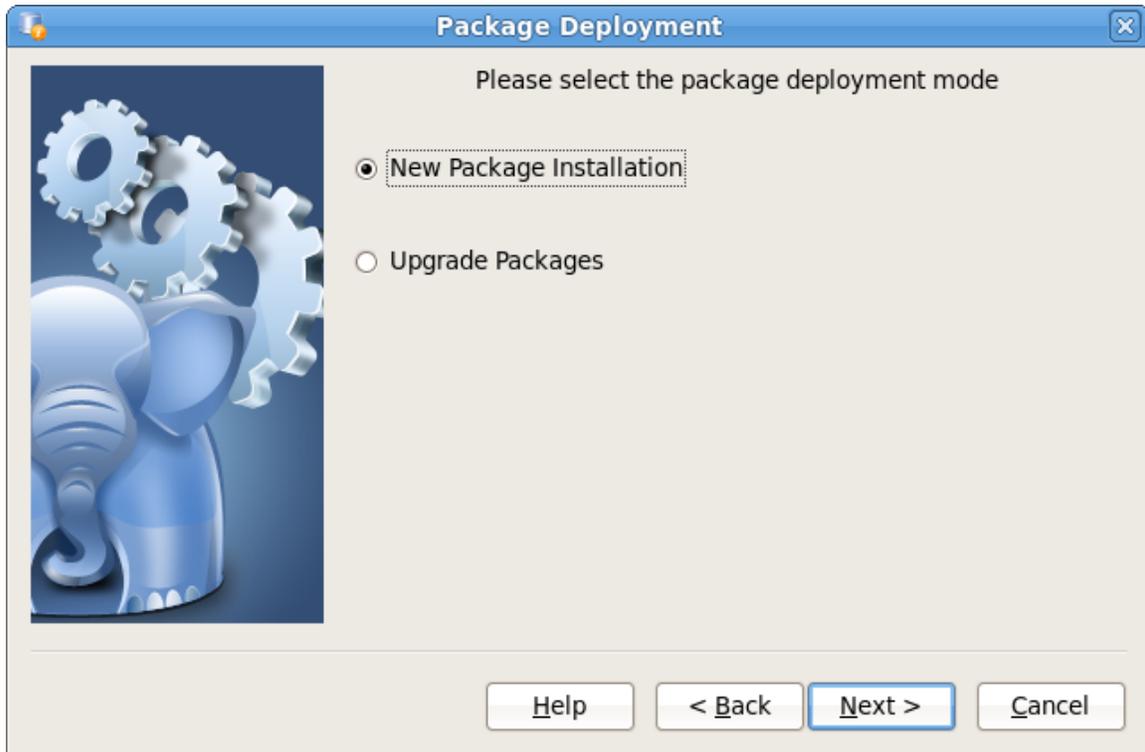


Figure 5.2 - The package deployment mode dialog.

Use the radio buttons on the Package Deployment Mode dialog (see Figure 5.2) 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.

5.1 Installing a New Package

If you select New Package Installation on the Package Deployment Mode dialog, the Package Deployment wizard opens the Agents/Servers dialog (see Figure 5.3), allowing you to specify the agents and servers on which the new applications will be installed.

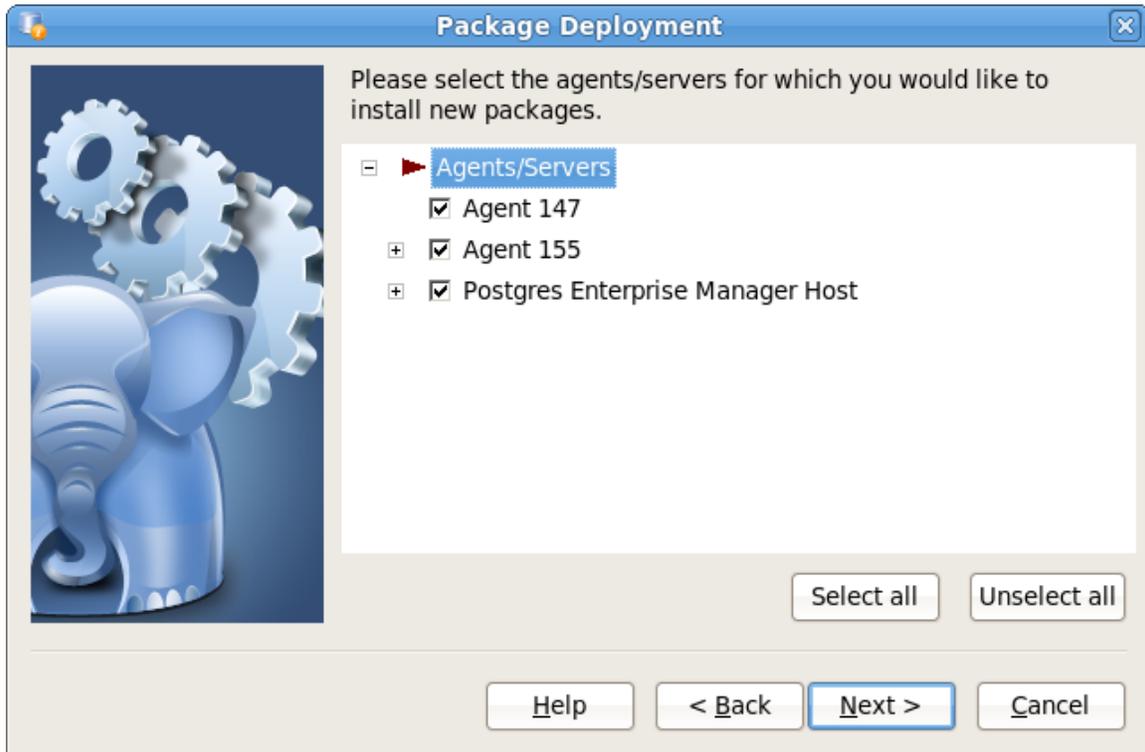


Figure 5.3 - Specify the target Agents and Servers.

Expand the `Agents/Servers` node of the tree control, and check the box next to each server on which you wish to install a new package:

- Use the `Select All` button to mark all of the listed servers for the package installation.
- Use the `Unselect All` button to deselect all of the listed servers.

Check the box next to the name of any agent or server on which you wish to install new software, and click `Next`. The package selection dialog opens (see Figure 5.4).

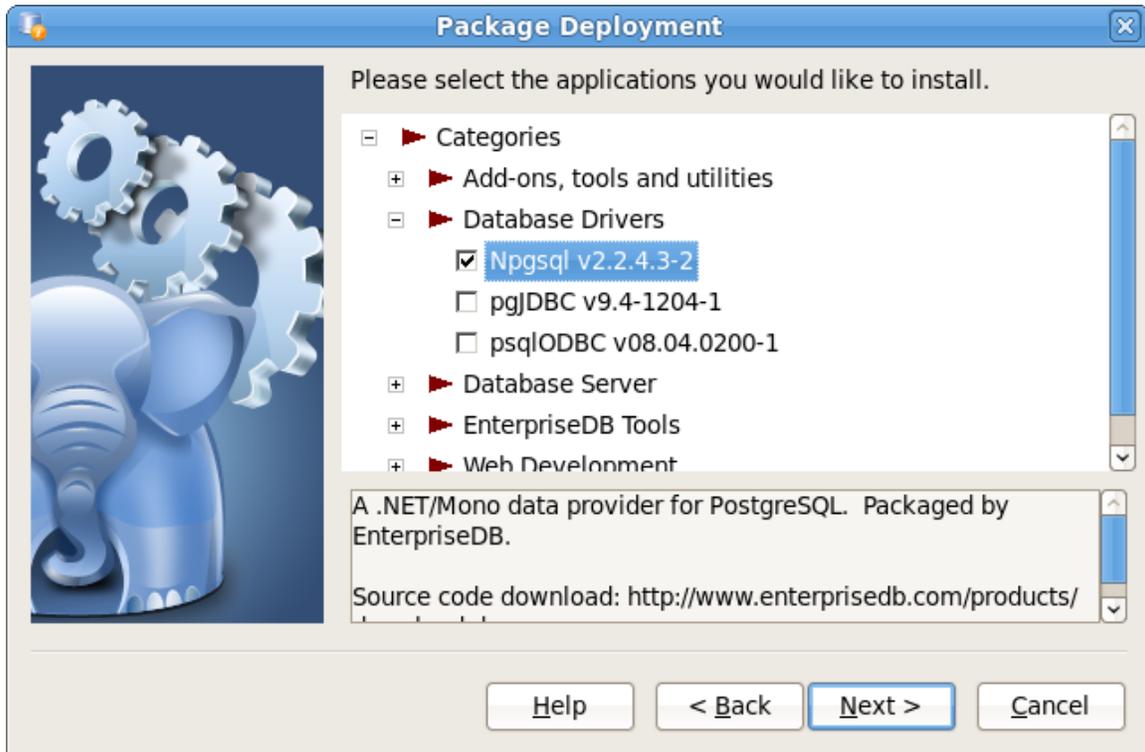


Figure 5.4 - Select which applications are to be installed.

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.

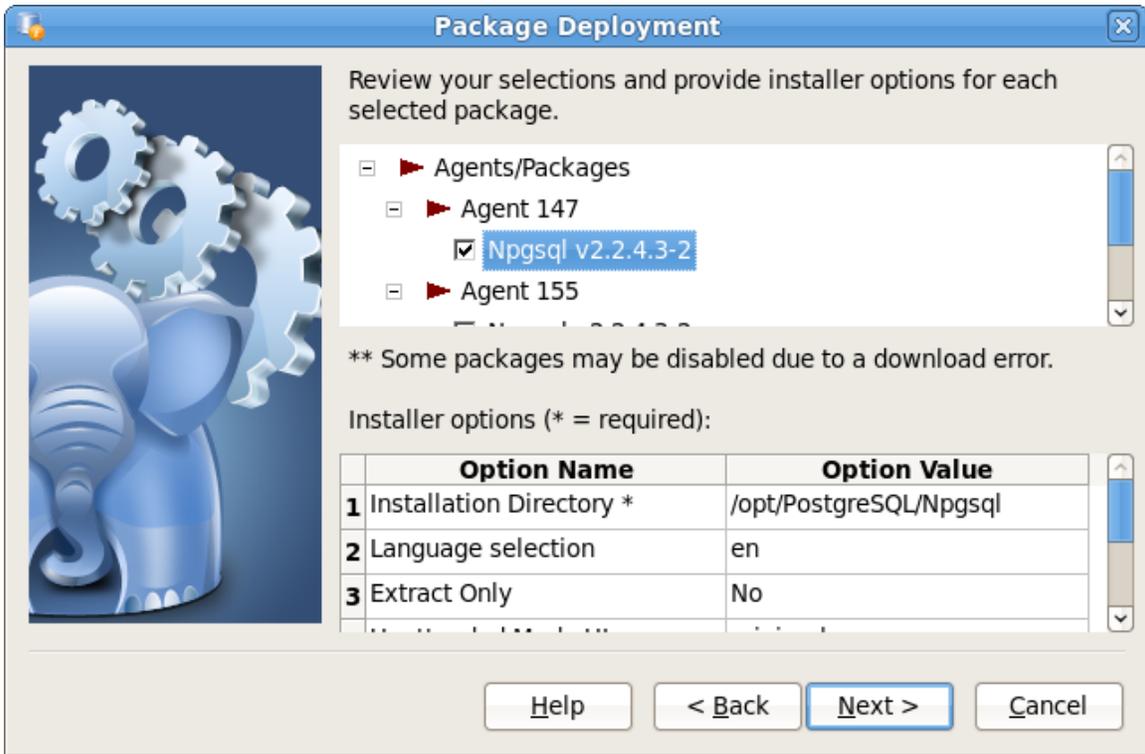


Figure 5.5 - Specify installation options.

Review the list of packages that will be installed, and (if prompted by the Package Deployment wizard) provide any options requested (see Figure 5.5). Click `Next` to continue.

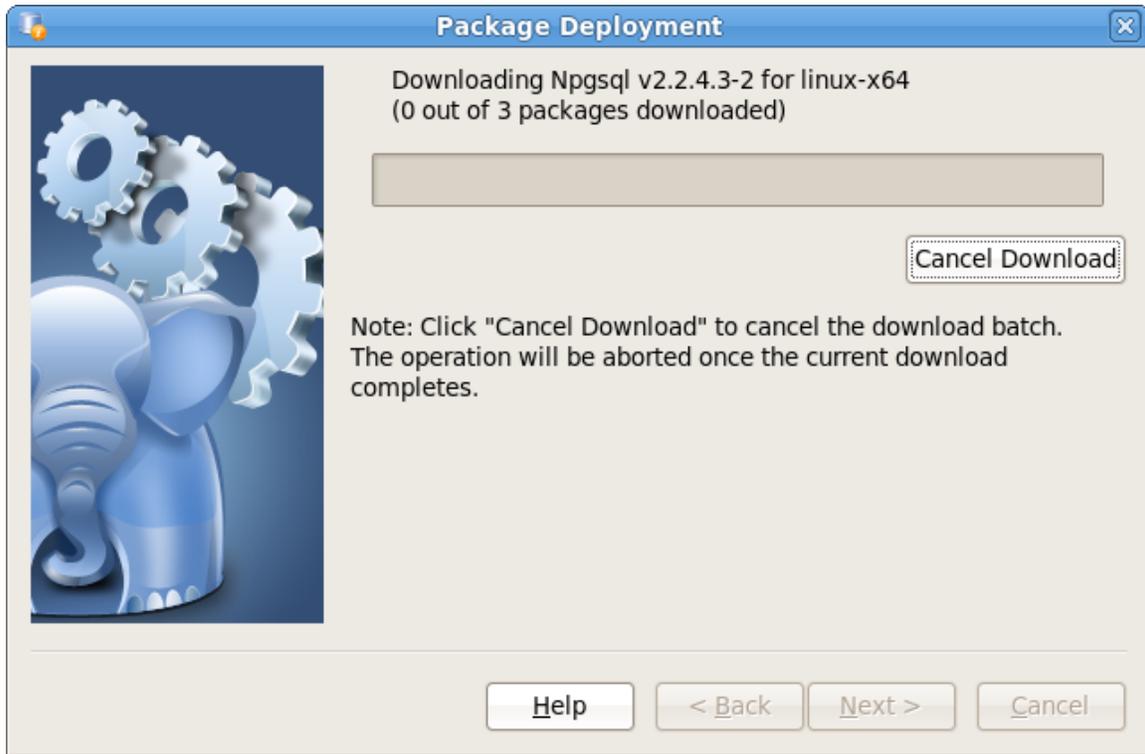


Figure 5.6 - Starting the installer download.

Click the `Start Download` button to instruct the Package Deployment wizard to download application installers (see Figure 5.6). During the download, you can click the `Cancel Download` button to abort the batch download. When the download completes, click `Next`.

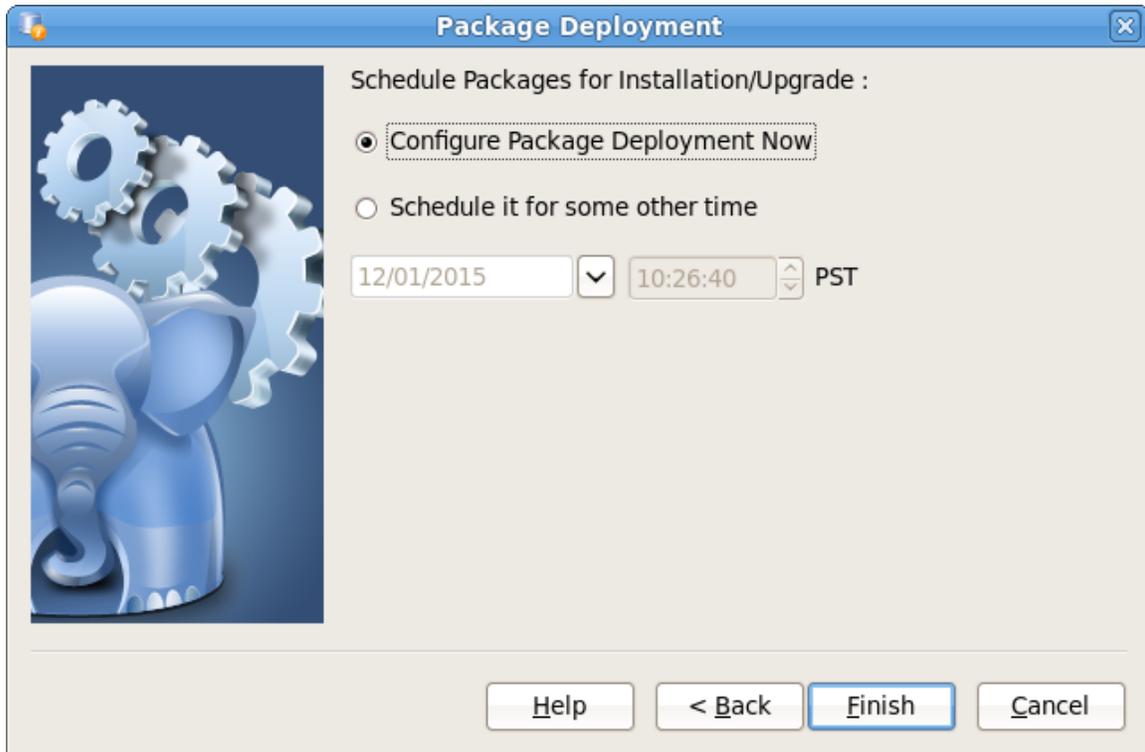


Figure 5.7 - The Package Deployment scheduling dialog.

Use the options on the scheduling dialog (see Figure 5.7) to schedule an installation time for the new packages:

- Select the radio button next to `Configure Package Deployment Now` 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.
- Select the radio button next to `Schedule it for some other time` to enable the date and time fields. Use the date and time fields to specify the 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 included on the `Scheduled Tasks` dialog (accessed through the `Scheduled Tasks...` menu selection on the agent's context menu).

5.2 Upgrading an Installed Package

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

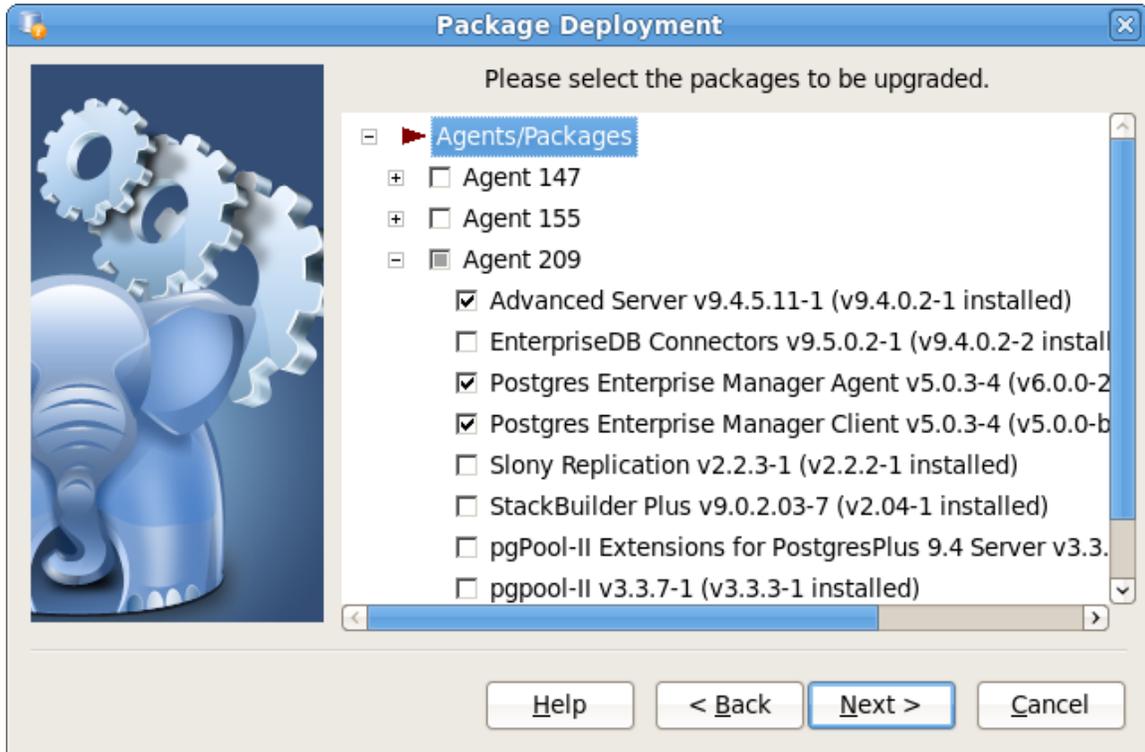


Figure 5.8 - Select the packages that will be upgraded.

Expand the `Agents/Packages` node of the tree control, and check the box next to each agent on which you wish to upgrade packages (see Figure 5.8). Click `Next` to continue.

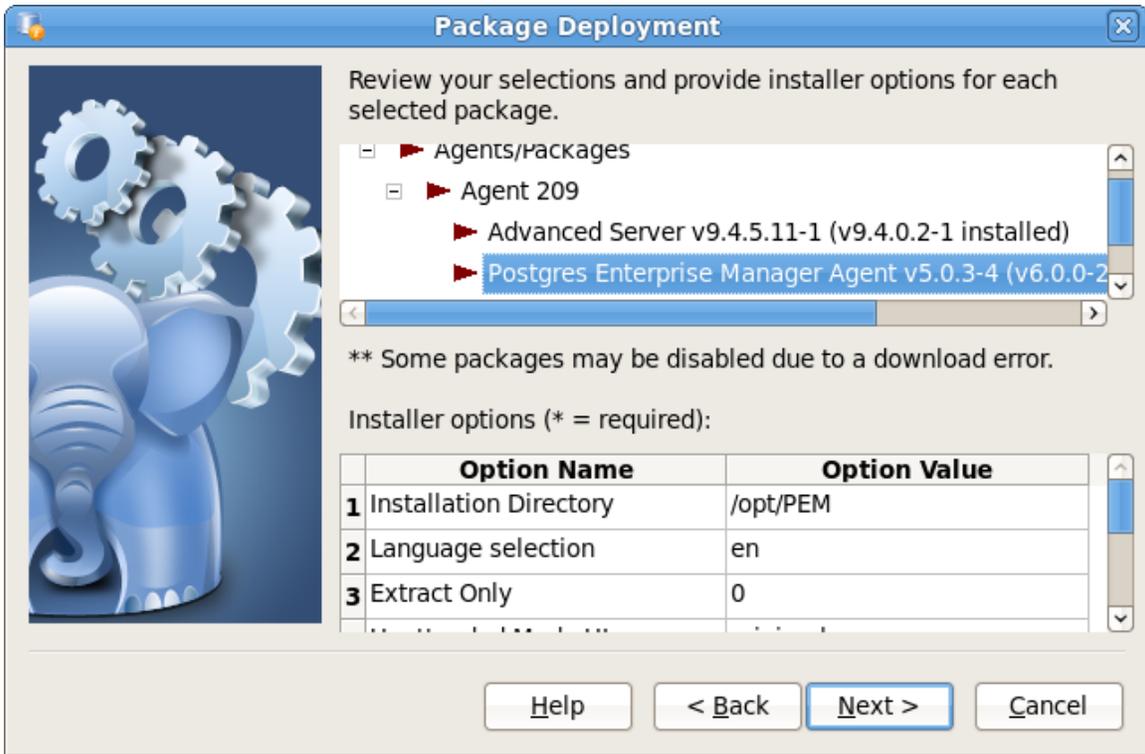


Figure 5.9 - Provide any requested installation options.

Review the list of installed packages, and provide any installation options requested in the Option Name/Option Value fields (see Figure 5.9). When you've reviewed the list, click Next to continue.

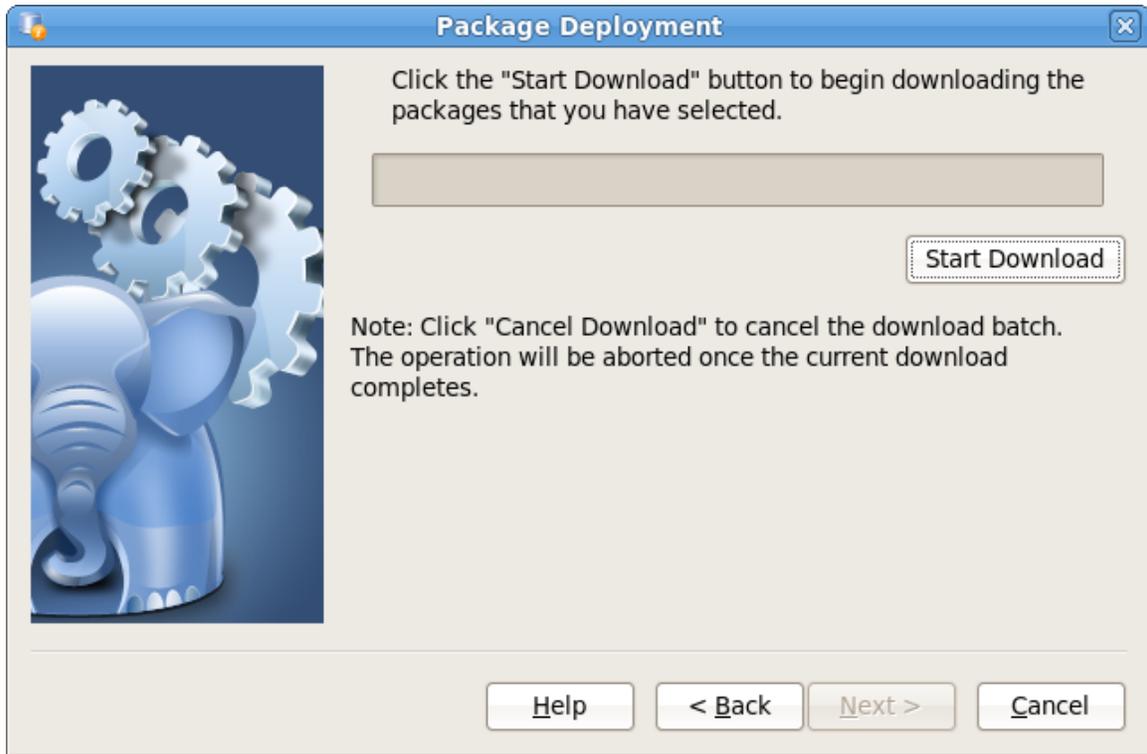


Figure 5.10 - Downloading the application installers.

Click the `Start Download` button to instruct the Package Deployment wizard to download application installers (see Figure 5.10). During the download, you can click the `Cancel Download` button to abort the batch download. When the download completes, click `Next`.

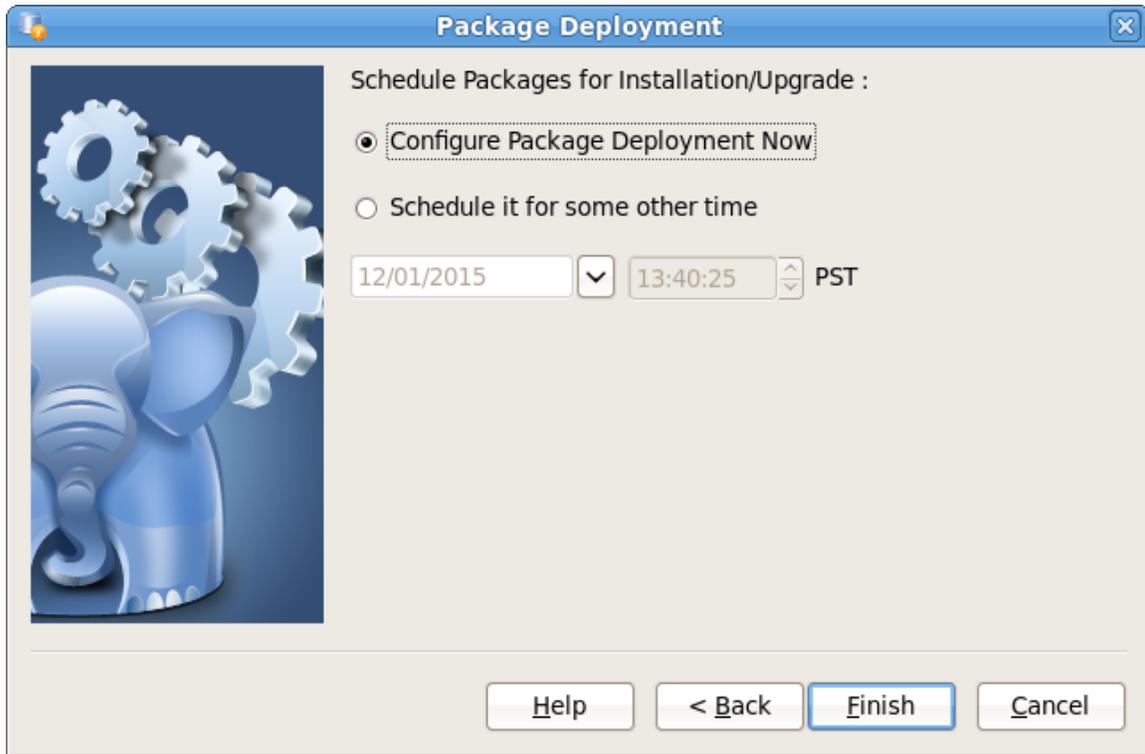


Figure 5.11 - Schedule a time for installation.

Use the options on the scheduling dialog (see Figure 5.11) to schedule an installation time for the new packages:

- Select the radio button next to `Configure Package Deployment Now` 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.
- Select the radio button next to `Schedule it for some other time` to enable the date and time fields. Use the date and time fields to specify the 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 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).

6 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.

6.1 Using Dashboards to View Performance Information

PEM displays performance statistics through a number of *dashboards* that you can navigate in web browser fashion. Each dashboard contains a series of summary views that contain charts, graphs and tables that display the statistics related to the selected object.

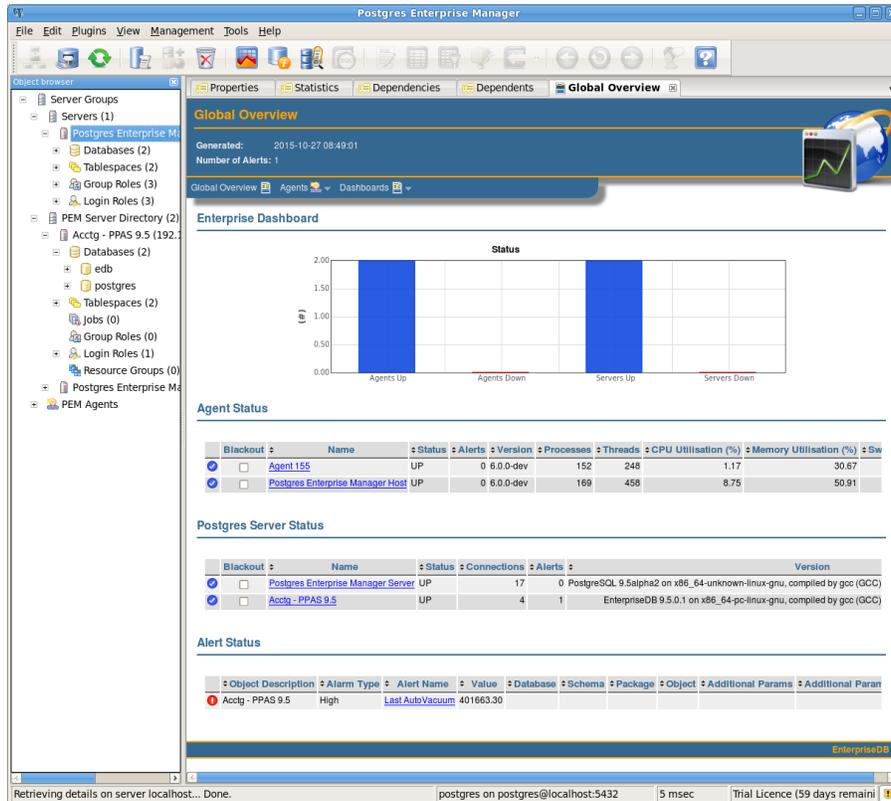


Figure 6.1 - The Global Overview dashboard.

The top-level dashboard is the Global Overview (shown in Figure 6.1). 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 in PEM's tabbed interface (shown in Figure 6.2), opened by default in the right hand side of the client window.

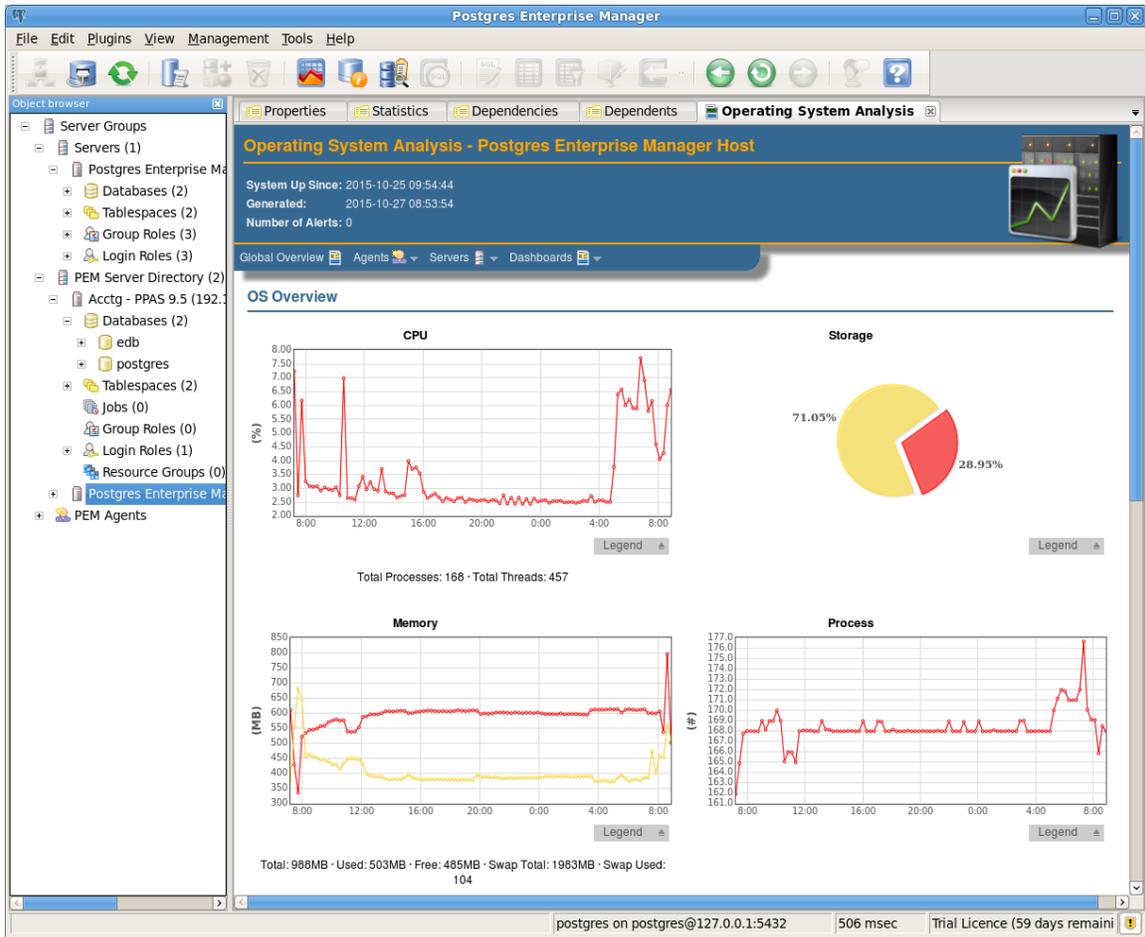


Figure 6.2 - The Operating System Analysis dashboard.

After opening a dashboard, you can navigate to other dashboards within the same tab.

Each dashboard header includes navigation menus that allow you to open dashboards that contain statistical information for the currently selected object, and any object that resides beneath the selected object (in the tree control hierarchy). You can use the `Browser Back` and `Browser Forward` buttons (on the main toolbar) to scroll backward and forward through the previously-viewed dashboards (within a tab). Click the `Refresh` button to update the current dashboard.

Click the `x` on the dashboard tab to close a dashboard.

6.2 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. To open the dashboard editor, select *Manage Dashboards . . .* from the PEM Client's Management menu. When the dashboard editor opens, click the *Manage Dashboards* button to view a drop-down list of previously defined dashboards, or to specify that you would like to create a new dashboard (see Figure 6.3).

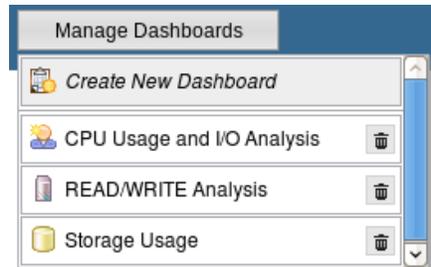


Figure 6.3 – The Manage Dashboards menu.

Select the name of an existing dashboard, or click the *Create New Dashboard* button to begin defining a dashboard. If you select *Create New Dashboard* the *New Dashboard* dialog opens (see Figure 6.4).

Figure 6.4 – The New Dashboard dialog.

Use the fields on the `New Dashboard` dialog to specify general information about the dashboard:

- Specify a title for the dashboard in the `Title` field.
- Provide a description of the dashboard in the `Description` field.
- 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, based on the probes that execute for that level.
- Use the `Teams` selector to specify which roles will have access to the new dashboard.
 - Highlight a name in the list of available roles (the left column), and use the right arrow to move the name into the list of names that have access to the chart.
 - Highlight a name in the list of roles with access to the chart (the right column), and use the left arrow to remove a name from the list of names with access to the chart.
- Check the box next to `Share with all` to instruct the server to allow all `Teams` to access the dashboard.
- Check the box next to `Ops Dashboard?` to instruct the server to create a dashboard that is formatted for display on an Ops monitor. For detailed information about defining a dashboard for an Ops monitor, please see the online help text distributed with the PEM client.

Use the fields within the `Ops Options` portion of the `New Dashboard` dialog to specify title information for the dashboard:

- Check the box next to `Show Title?` to display the dashboard title at the top of the Ops dashboard.
- Use the drop-down list boxes to select a custom font style and font size for the title. The selected font style will be displayed in the `Preview` box.

When you've completed the `New Dashboard` dialog, click `OK` to continue.

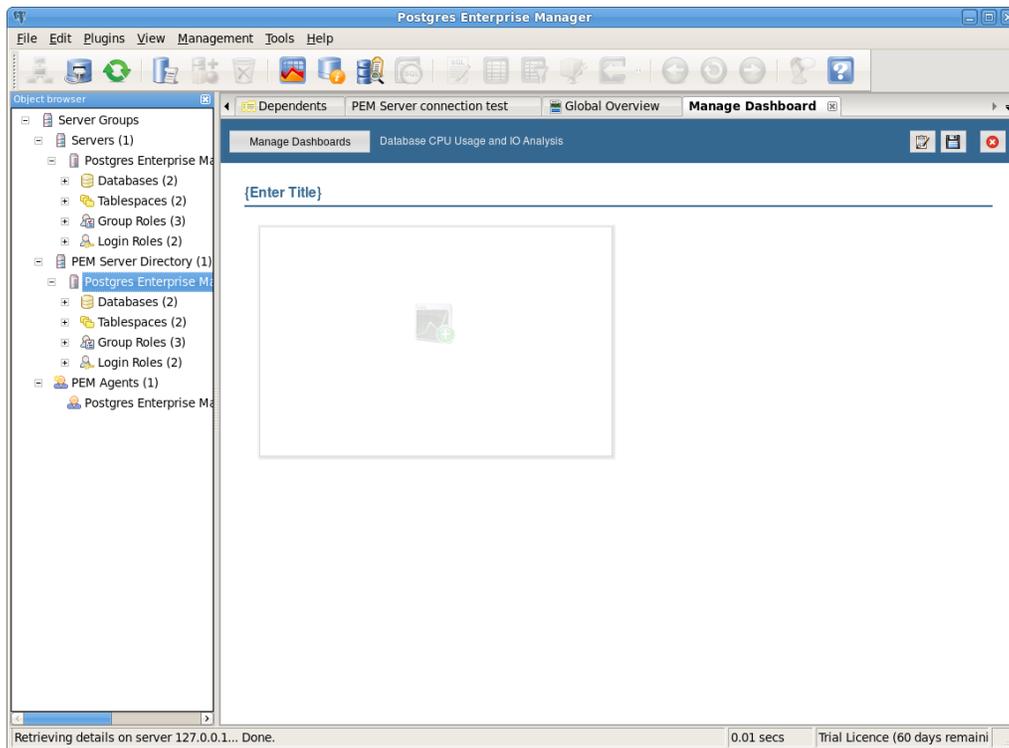


Figure 6.5 – Creating a custom dashboard.

Click the {Enter Title} label to add a section header to the dashboard. Note that before adding a chart to a new section of the dashboard, you must replace the {Enter Title} section header placeholder with a section header (see Figure 6.5).

After adding the section header, click the Add Chart icon (located in the chart placeholder) to open a selection dialog that allows you to choose the chart that will appear in that section. Expand a category to view the charts that are available in the category; highlight the name of a chart to view the chart's configuration information and a list of the metrics that are displayed on the chart.

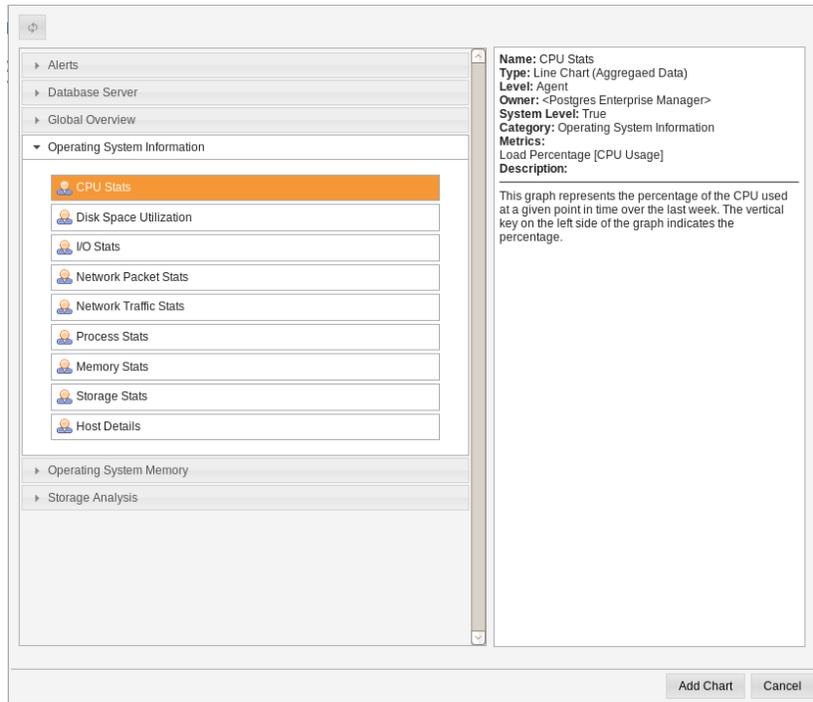


Figure 6.6 – Adding a chart to a custom dashboard.

When you've selected a chart that you wish to display in the specified location on the dashboard, click the **Add Chart** button (see Figure 6.6).

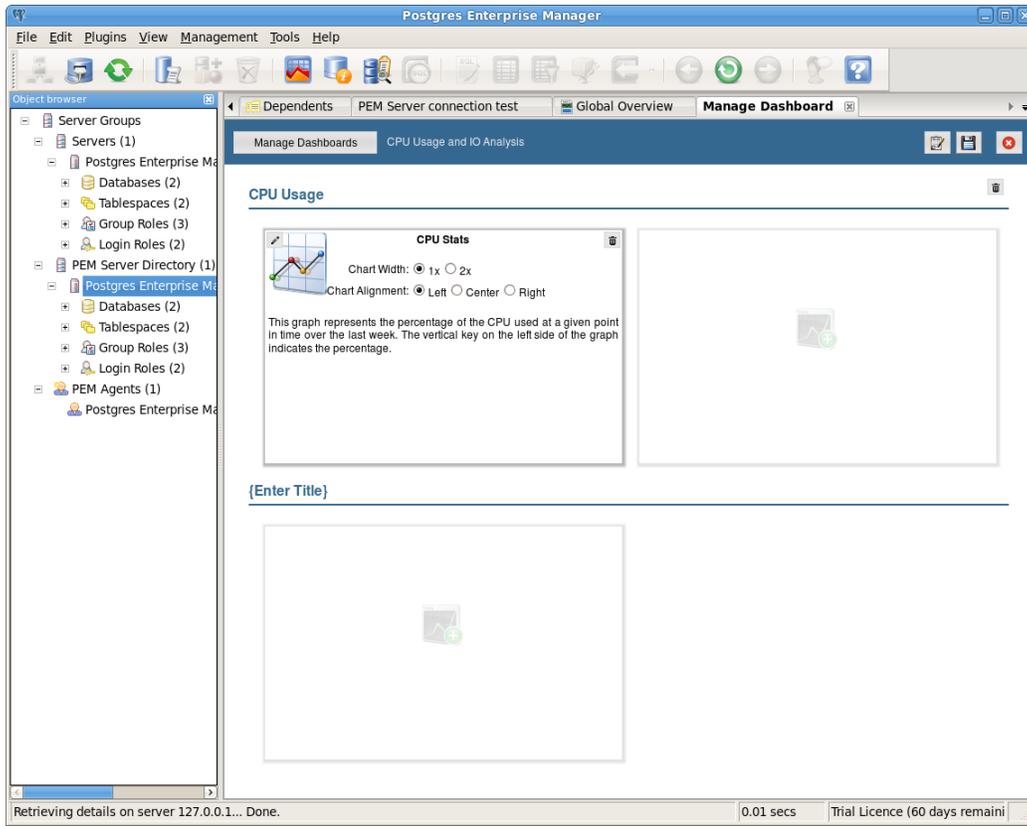


Figure 6.7 – Specifying chart positioning details.

The dashboard editor displays the title and description of the selected chart in the chart placeholder on the dashboard (see Figure 6.7). Beneath the chart title, radio buttons allow you to select the display details:

- Use the radio buttons next to `Chart Width` to indicate the number of columns the chart should occupy on the dashboard; select the `1x` radio button to indicate one column, or `2x` to indicate two columns. Please note that a dashboard is two columns wide.
- Use the radio buttons next to `Chart Alignment` to indicate the position of the chart within the section:
 - Select the `Left` radio button to indicate that the chart should occupy the left side.
 - Select the `Center` radio button to indicate that the chart should be centered.

- Select the `Right` radio button to indicate that the chart should occupy the right side.
- If you are adding the chart to a dashboard designed for display on an Ops monitor, use the radio buttons next to `Show Chart Title` to specify if the dashboard should include a chart title. Select `Yes` to display a title, or `No` to omit the title.

To replace a chart with a different chart, select the edit icon (located in the upper-left hand corner of the chart placeholder). The chart selection dialog will open, allowing you to select a replacement chart; when you've made your selection, click the `Update Chart` button to replace the current chart with the selected chart.

To delete a chart from the dashboard, select the delete icon (located in the upper-right hand corner of the chart placeholder). PEM will open a popup, prompting you to confirm that you would like to delete the selected chart.

To add another chart to the dashboard, select the `Add Chart` icon on another chart placeholder, and repeat the selection process.

To make changes to the dashboard title or definition, click the `Edit` icon, located in the upper-right hand corner of the dashboard editor.

To save your new dashboard or modifications to an existing dashboard and exit, click the `Save` icon, located in the upper-right corner of the dashboard editor.

To exit without saving your changes, click the `Close` icon, located in the upper-right corner of the dashboard editor.

6.2.1 Creating an Ops Dashboard

You can use the PEM client's Custom Dashboard dialog 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 (see Figure 6.8).



Figure 6.8 – An Ops Dashboard.

To create an Ops dashboard, check the box next to Ops Dashboard?, and use the Ops Options portion of the PEM New Dashboard dialog to specify dashboard details. After defining an Ops dashboard, use the dashboard editor to add charts and graphs and optional titles to the new dashboard. After adding charts and tables to the Ops dashboard, use the disk icon in the upper-right corner of the dashboard editor to save your work.

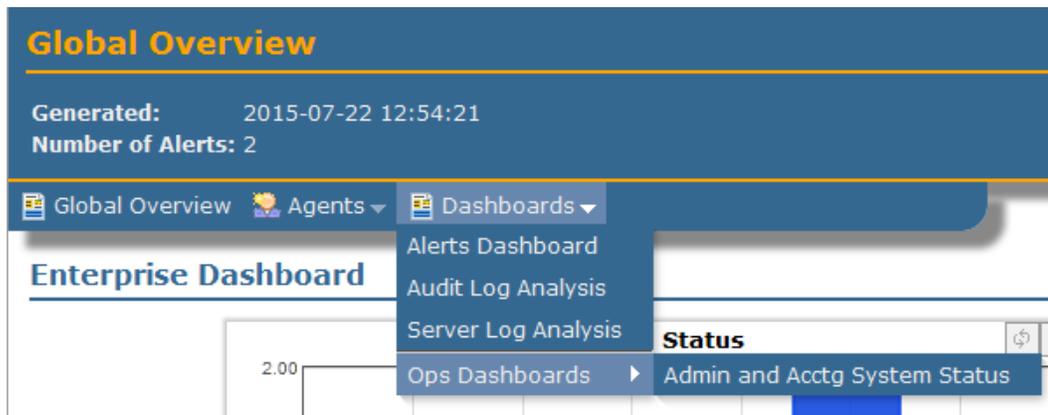


Figure 6.9 – Accessing an Ops dashboard.

Navigate through the `Dashboards` menu of the hierarchy level specified in the `Level` field on the `New Dashboard` dialog to open an `Ops` dashboard. For example, if you specified a value of `Global`, the dashboard will be listed under the `Dashboards` menu of the `Global Overview` (see Figure 6.9).

6.3 Creating a Custom Chart

You can use the PEM chart editor (see Figure 6.10) to create or modify a custom line chart or table; after defining a chart, you can display the chart on a custom dashboard. To open the chart editor, select `Manage Charts...` from the PEM Management menu.

Use the `Manage Charts` button to view a list of options that allow you to create a new chart, or modify an existing chart:

- Select the name of a previously defined chart to modify an existing chart.
- Click the `Create New Chart` button to begin defining a chart.
- Select the `Import Capacity Manager Template` button to select a capacity manager template on which to base the chart.

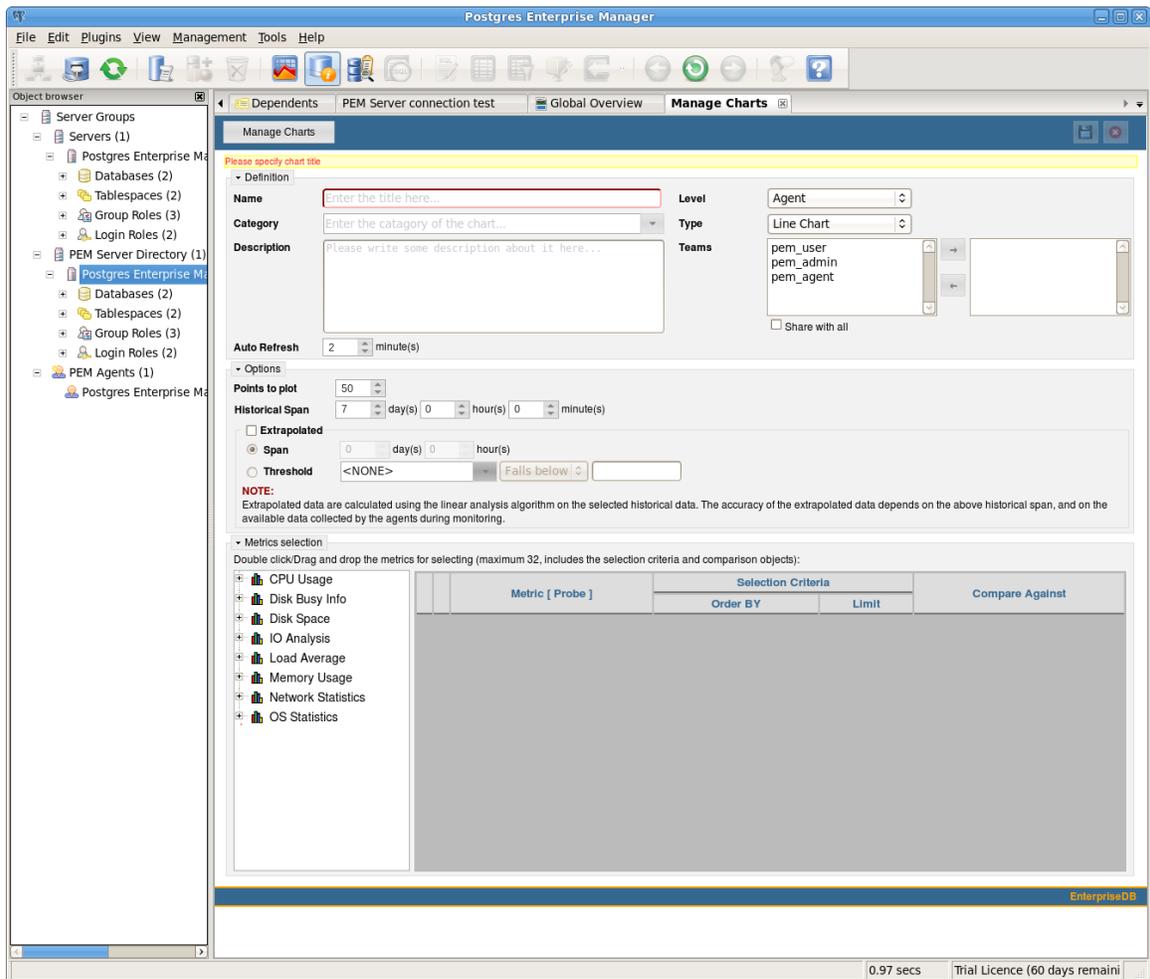


Figure 6.10 – Defining a new chart.

Use the `Definition` button to open or close the `Definition` panel. Use the fields in the `Definition` panel to specify general information about the chart:

- Specify the name of the chart in the `Name` field.
- Specify the category in which this chart will be displayed in the `Category` field. When adding a custom chart to a custom dashboard, the chart will be displayed for selection in the `Category` specified.
- 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.
- Use the `Auto Refresh` field to specify the number of minutes between chart updates - choose a value from 1 to 999. The default auto refresh rate is 2 minutes.
- Use the `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 has access to a unique set of probes that return the information that you can include in your chart.
- Use the `Type` drop-down listbox to specify the type of chart that you would like to create. Select either a `Line Chart` or a `Table`.
- Use the `Teams` selector to specify which roles will have access to the new chart:
 - Highlight a name in the list of available roles (the left column), and use the right arrow to move the name into the list of names that have access to the chart.
 - Highlight a name in the list of roles with access to the chart (the right column), and use the left arrow to remove a name from the list of names with access to the chart.
 - Check the `Share with all` checkbox to indicate that all users may access the chart.

Use the `Options` button to open or close the `Options` panel. Use the fields in the `Options` panel to specify configuration details for the chart:

- Use the `Points to plot` field to specify the number of points that PEM should plot on the chart.
- Use the `Historical Span` field to specify the number of minutes, hours, or days of historical data that should be included on the chart.

- Check the box next to `Extrapolated` to instruct PEM to generate extrapolated data based on the specified historical data, and enable the `Span` and `Threshold` radio buttons.
- Select the radio button next to `Span`, and use the day and hour selectors to specify the period of time spanned by the metrics on the chart. Specify the number of `Days` and `Hours` for which data will be displayed.
- Select the radio button next to `Threshold`, and use the threshold selectors to specify a maximum or minimum value for the chart.

Use the `Metrics selection` button to open or close the `Metrics selection` panel. Use the fields in the `Metrics selection` panel to select or modify the metrics that are displayed on a line chart.

If you are creating a line chart, the probes that correspond to the selected level (`Agent`, `Server`, or `Database`) will be displayed in a tree control in the `Metrics selection` panel. To include a given metric on your chart, expand the tree control and double-click the name of the metric, or drag the metric to the selection panel.

Columns on the selection panel allow you to design the content of your line chart:

- Use the double-ended arrow control as an anchor for re-ordering the items within the selection panel. Click an arrow, and drag the associated item to a new location within the list.
- Click the garbage can icon to delete an item from the selection panel.
- The selection panel displays the name of the metric in the (non-modifiable) `Metric [Probe]` column.
- Use the selection boxes in the `Selection Criteria` column to specify the order of the data displayed and limits for the metric.
- PEM supports comparisons of cross-hierarchy metrics. You can use the `Compare Against` column to select a metric from another agent, server, or database to compare to the specified metric.
 - Click the pencil icon to open a tree control that allows you to select objects for comparison. Expand the tree control to access objects that are eligible for comparison; double-click an object name to move it into the `Selected objects for comparison` box.

When you've selected the objects that will be compared to the metric, click `Ok` to preserve your selections and return to the `Metrics selection`

panel, or `Cancel` to return to the `Metrics` selection panel without saving your selections.

- Click the `copy` icon to copy the comparable objects for this metric to all of the metrics for the same probe.

If you are creating a table, the bottom panel will display a list of probes that correspond to the selected level (`Agent`, `Server`, or `Database`). Select a probe name to view the available metrics gathered by that probe. Check the box next to the name of a metric to include the metric on your table, and to access fields that allow you to specify the order of the data displayed, and limits for the metric.

When you've finished defining the chart, click the disk icon in the upper-right hand corner of the chart editor to save your edits.

6.3.1 Controlling and Customizing Charts, Graphs and Tables

Use the icons in the upper-right corner of each graphic on a dashboard to control and customize the charts, graphs and tables displayed in the PEM client for your current user session (see Figure 6.11).

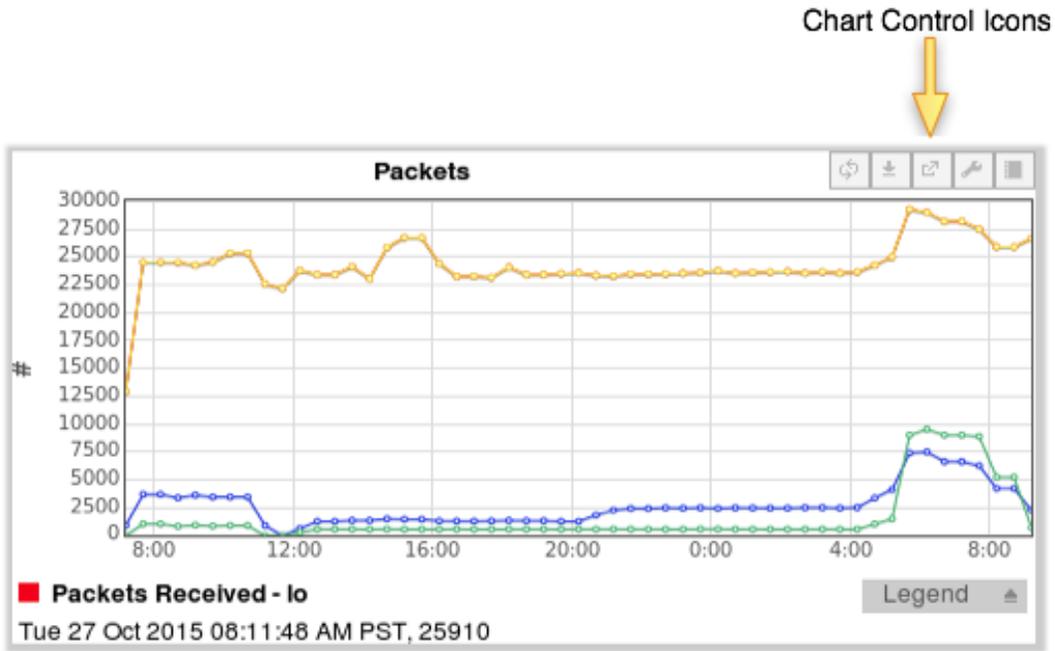


Figure 6.11 – The PEM Client chart control icons.

Select an icon to:

-  Refresh the content of a chart, graph or table.
-  Download an image of the chart or graph.
-  Expand the chart or graph to full-screen.
-  Personalize the chart, graph, or table settings for the current user.
-  View information about the chart, graph, or table.

For more information about customizing the graphics displayed on the PEM dashboards, please see the PEM client online help.

6.4 Customizing Probes

PEM uses *probes* to retrieve statistics from a monitored server, database, operating system or agent. A probe is a scheduled event that returns a set of performance metrics about a specific monitored object. You can use the `Probe Configuration` dialog (shown in Figure 6.12) to specify when each probe is executed.

To review or modify the currently defined probes for each server and its underlying objects (e.g. databases, tables, etc.), right click on the object name in the tree control and select `Probe Configuration` from the context menu. You can also access the `Probe Configuration` dialog by highlighting an object name, and selecting `Probe Configuration` from the `Management` menu.

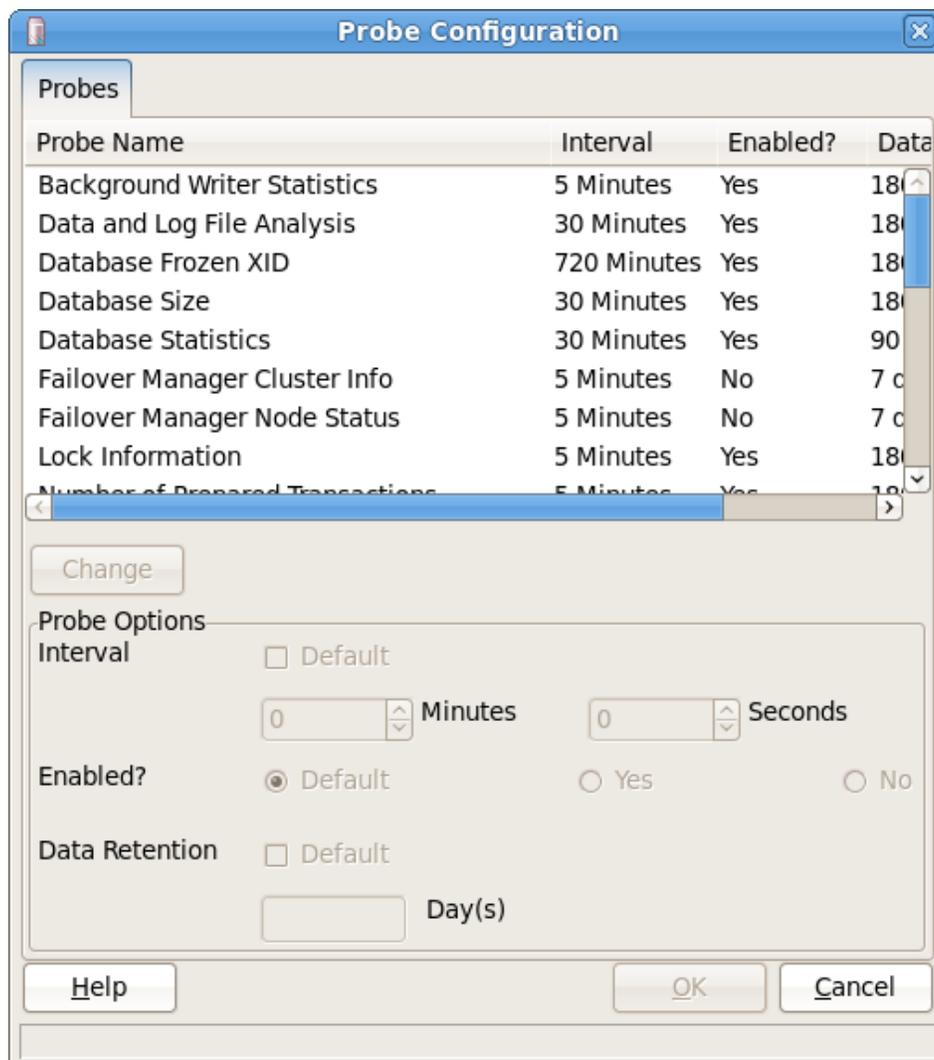


Figure 6.12 - The Probe Configuration dialog.

Probes monitor a unique set of metrics for each specific object type (server, database, database object, or agent). You can modify the properties associated with a probe, specifying:

- How often the probe executes.
- How long its information is retained for historical reporting purposes.
- If the probe is enabled or disabled.

The `Probe Configuration` dialog displays a list of the metrics that are collected for the selected node. The `Probe Configuration` dialog may also display information about probes that cannot be modified from the current node. If a probe cannot be modified from the current dialog, the background (the area behind the node name) is greyed-out; when highlighted, the status bar will display the node level from which the statistic may be configured.

6.5 Customizing Alerts

The PEM server comes with a number of pre-defined alerts that are actively monitoring your servers. You can also create custom alerts that will notify you when resource utilization exceeds user specified thresholds. 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 [6.5.3](#).

6.5.1 Creating a Custom Alert

Use the Alerting dialog (shown in Figure 6.13) to define or modify an alert. To open the Alerting dialog, right click on the name of a monitored object in the tree control, and select Alerting... from the context menu, or highlight the object name and select Alerting... from the Management menu.

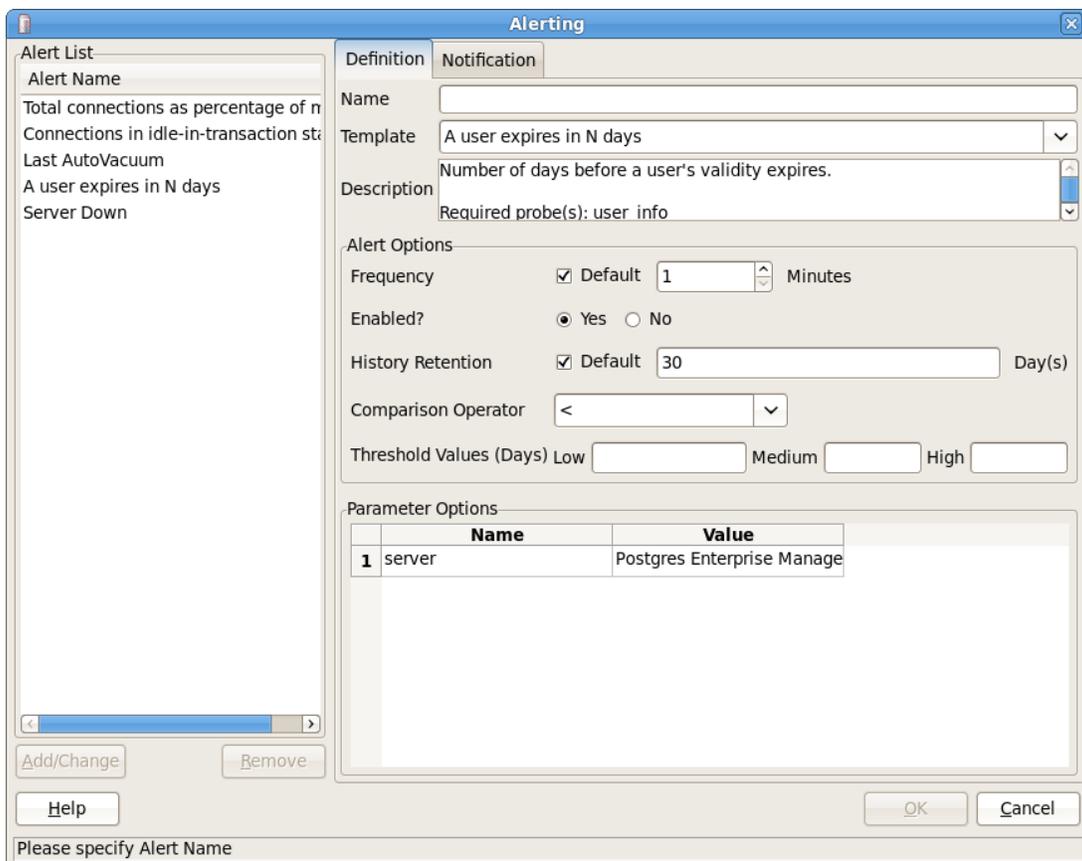


Figure 6.13 - The PEM Alerting dialog.

The `Alert List` displays a list of the currently defined alerts for the tree control node from which the dialog was opened.

The fields on the `Definition` tab (located in the right pane of the dialog) describe the properties of an alert. To define a new alert:

- Enter the name of the alert in the `Name` field.
- Use the drop-down listbox in the `Template` field to select an alert template that will specify the resource or metric evaluated by the alert.
- Highlight a template name to view a description of the template in the `Description` field.
- Use the fields in the `Alert Options` box to define the properties of the alert:
- Use the spin control in the `Frequency` field to specify how often the alert should check if the alert conditions are satisfied. When the spinner is set to the default (recommended) value for the selected template, the box next to `Default` is checked.
- Use the radio buttons next to `Enabled?` to enable or disable the alert; select `Yes` to enable the alert, and `No` to disable the alert.
- Use the `History Retention` field to specify the number of days that PEM will store data collected by the alert. By default, PEM will recommend storing historical data for 30 days.

The `Comparison Operator` and `Threshold Values` fields work together 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 determined by the `Comparison Operator`), PEM will raise a `Low`, `Medium` or `High` level alert:

- Use the `Comparison Operator` drop-down listbox 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 Values` fields to specify the values that PEM will compare to the system values to determine if an alert should be raised:
 - 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.

Please note that you must specify values for all three thresholds (Low, Medium, and High).

The `Parameter Options` box contains a table of the parameters that are required by the template; the table displays both parameters for which the system will provide a value and parameters for which the user will provide a value.

- Specify a value for any parameter that displays a prompt in the `Value` column.

Use the fields on the `Notification` tab (shown in Figure 6.14) to specify how PEM will behave if an alert is raised.

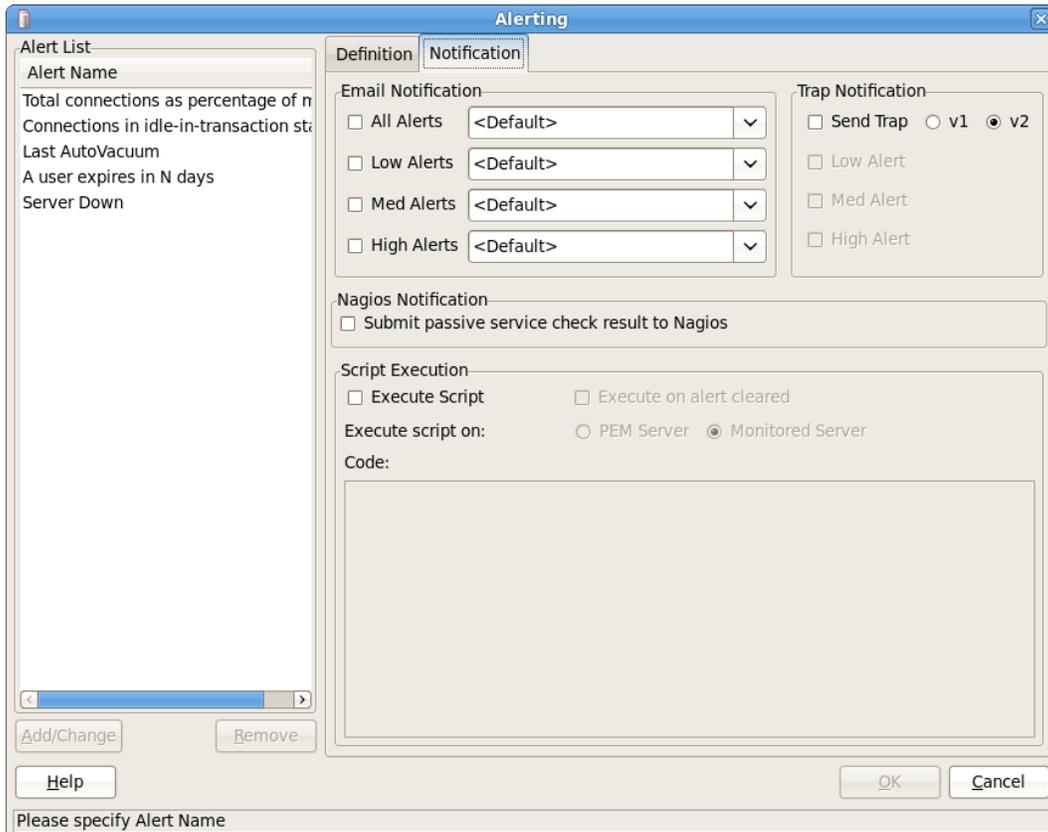


Figure 6.14 - The Notification tab.

PEM can send an email notification or execute a script if an alert is triggered or cleared. You can use the `SMTP Email Group` dialog to define the address of the user or users that will be notified; to access the `SMTP Email Group` dialog, navigate through the `PEM client Management` menu to the `Server` context menu, and select `Email Groups....`

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.

- Check the box next to an alert level, and 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:

- Check the `Send Trap` checkbox to send SNMP traps or SNMP notifications when the state of this alert changes. Check the box next to `Send Trap` and select the `v1` radio button to send SNMP v1 traps, or the `v2` radio button to send SNMP v2 notifications to the receiver.
- Use the `Low Alert`, `Med Alert` and `High Alert` checkboxes to specify the level of alert that will trigger the trap. For example, if you check the box next to `High Alert`, a notification will be sent when an alert with a high severity level is triggered.

Check the box next to `Submit passive service check result to Nagios` 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 the online help text available through the PEM client.

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.

- Check the box next to `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 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:

Placeholder	Description
<code>%AlertName%</code>	This placeholder will be replaced with the name of the triggered alert.
<code>%ObjectName%</code>	This placeholder will be replaced with the name of the server or agent on which the alert was triggered.
<code>%ThresholdValue%</code>	This placeholder will be replaced with the threshold value reached by the metric when the alert triggered.
<code>%CurrentValue%</code>	This placeholder will be replaced with the current value of the metric that triggered the alert.
<code>%CurrentState%</code>	This placeholder will be replaced with the current state of the alert.
<code>%OldState%</code>	This placeholder will be replaced with the previous state of the alert.
<code>%AlertRaisedTime%</code>	This placeholder will be replaced with the time that the alert was raised, or the most recent time that the alert state was changed.

Please Note - the status bar will assist you in defining an alert by prompting you for information still missing from the `Alerting` dialog.

When you have defined the alert attributes in the fields on the right side of the dialog, click the `Add/Change` button to add the new alert to the `Alert List`. To exit the `Alerting` dialog and save any changes to the `Alert List`, click `Ok`; to exit without saving additions or modifications to the list, select `Cancel`.

For detailed information about using the `Alerting` dialog, please see the PEM client's online help text.

6.5.1.1 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.

6.5.2 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.

Detailed information about each configuration step is available in the PEM client online help.

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

6.5.3 Creating a Custom Alert Template

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 (of the type specified in the template that is associated with the alert) to determine if a situation exists that requires administrative attention.

The `Alert Templates` dialog (shown in Figure 6.15) provides an interface that allows you to define a custom alert template or view and modify the definitions of existing alert templates. To open the `Alert Template` dialog, select the `Alert Templates...` menu option from the `PEM` client `Management` menu.

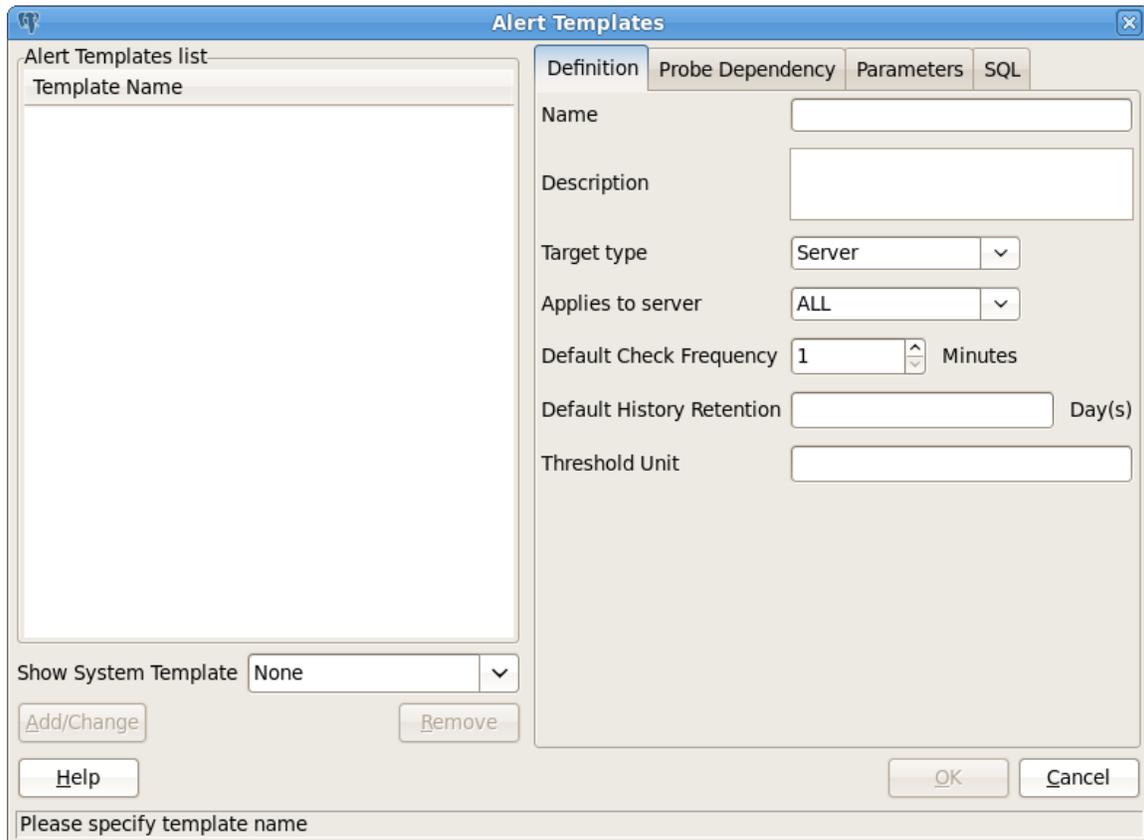


Figure 6.15 - The Definition tab of the Alert Templates dialog.

To view or modify an existing template, highlight the template name in the `Alert Templates` list in the left pane. Fields on the `Definition` tab (located in the right pane) specify general information about the template:

- Use the `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 to which the alert will be applied; you can specify a single server type, or `ALL`.
- Use the `Default 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.
- Use the `Default 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.

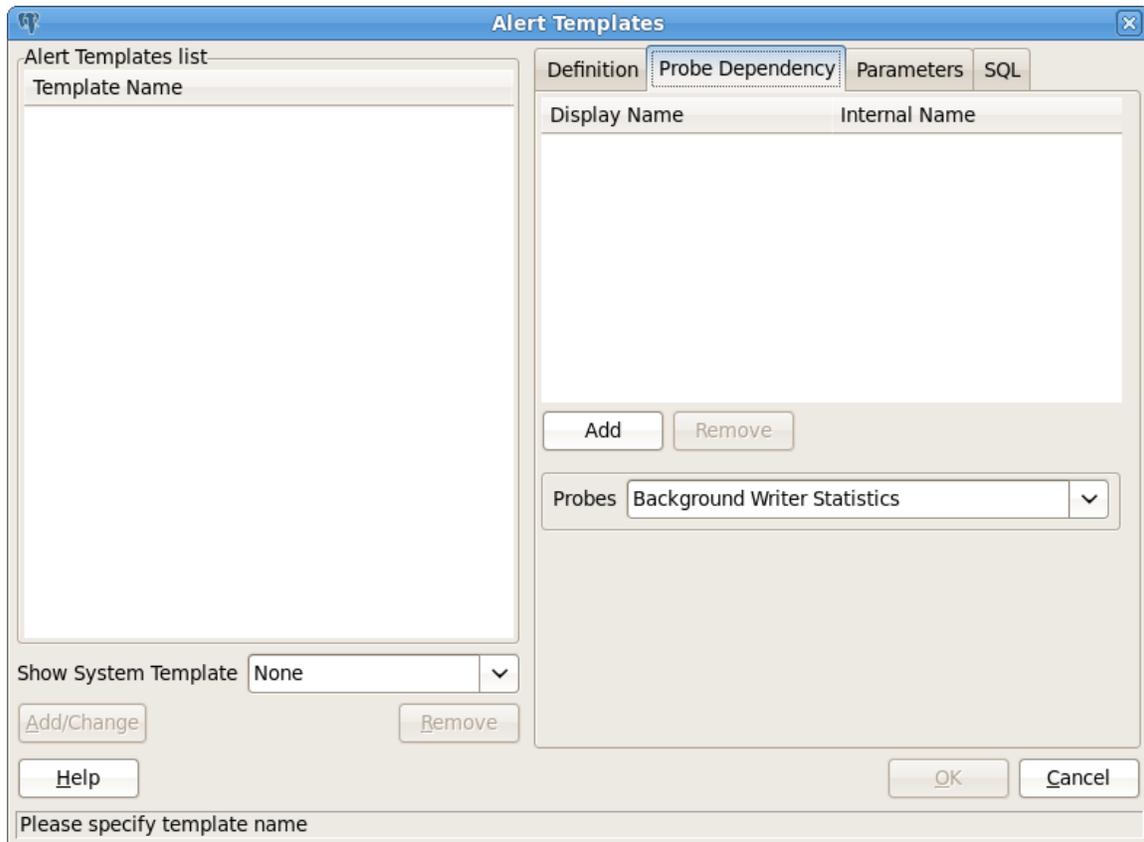


Figure 6.16 - The Probe Dependency tab of the Alert Templates dialog.

Use the fields on the `Probe Dependency` tab (see Figure 6.16) 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 one or more probes from the selected probes list, highlight the probe name(s), and click the `Remove` button.

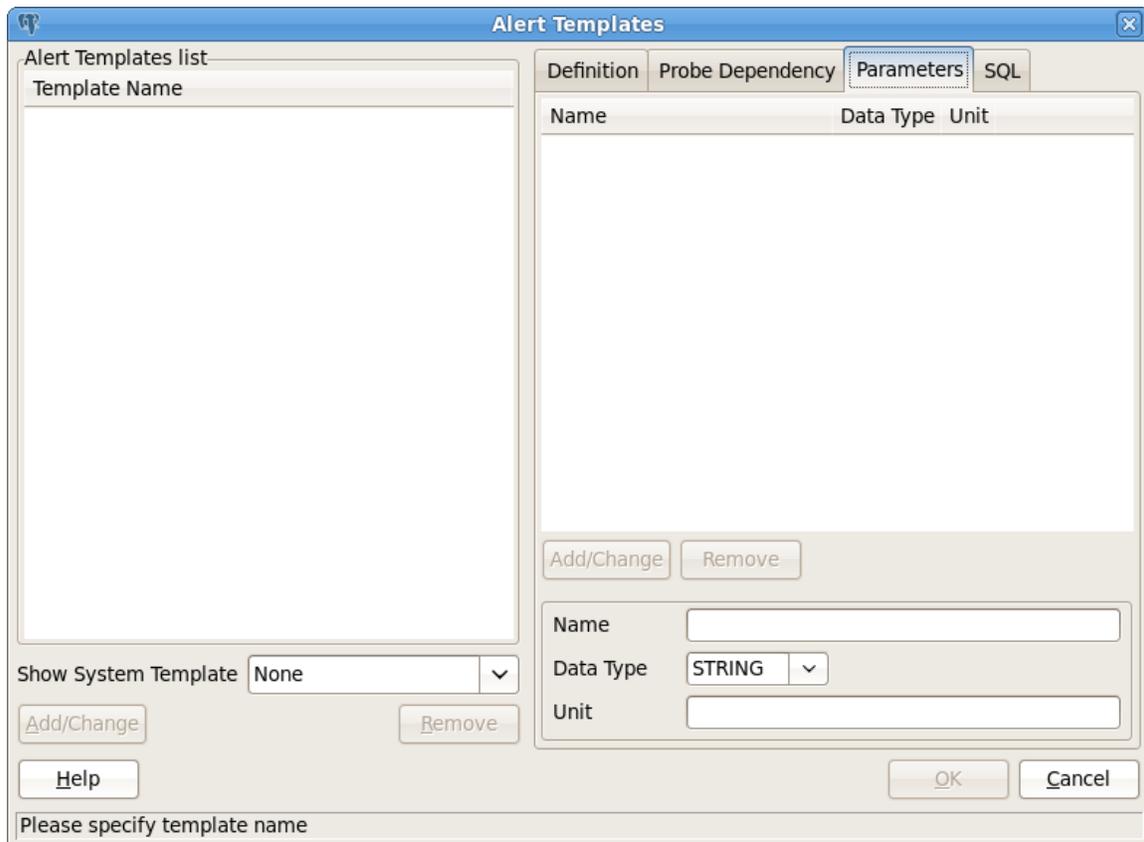


Figure 6.17 - The Parameters tab of the Alert Templates dialog.

Use fields on the `Parameters` tab (see Figure 6.17) to define the parameters that will be used in the SQL code specified on the `SQL` tab:

- 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.

When you've defined a new parameter, click the `Add/Change` button to save the definition and add the parameter to the parameter list.

To modify an existing parameter definition, highlight a parameter name in the list, modify the parameter values in the fields at the bottom of the tab, and click `Add/Change` to preserve the changes. To remove one or more parameter definitions, highlight the parameter name(s) and click the `Remove` button.

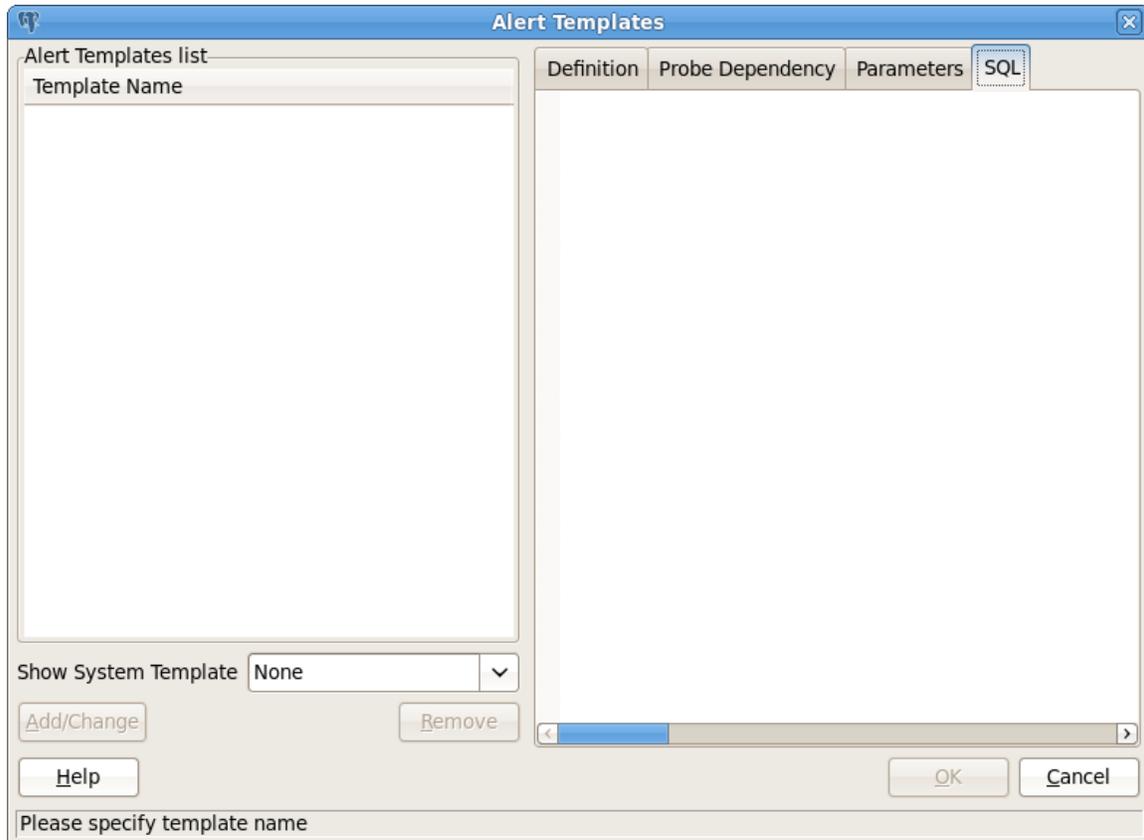


Figure 6.18 - The SQL tab of the Alert Templates dialog.

Use the SQL tab (shown in Figure 6.18) 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.

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}'

Click the **Add/Change** button to save the alert template definition and add the template name to the **Alert Templates** list; then click **OK** to save your work and exit.

After saving a custom alert template, you can use the **Alerting** dialog to define an alert based on the template. For more information about creating a custom alert, see [Section 6.5.1](#).

6.6 Viewing and Responding to Alerts

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*.

Alerts call your attention to conditions on registered servers that require your attention.

You can view alert information in a number of places inside PEM:

- The `Global Overview` dashboard contains a count of all alerts that have occurred, along with a breakdown of total alerts by monitored server and a listing of the most recent alerts that have occurred.
- The `Alerts` dashboard displays summarized statistics for all alerts that have occurred across your servers as well as a detailed listing of each alert that has been identified.
- When an alert is triggered, a flashing  icon is displayed in the lower right-hand corner of the main window. Click the icon to open the `Alerts` dashboard.

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

7 Capacity Manager

PEM contains built-in capabilities for performing database capacity planning. Capacity planning helps DBAs by providing answers to questions like:

- How much storage will my database need six months from now?
- How fast is my database growing?
- What objects are responsible for the growth in my database?
- Will my server be able to support another database instance?
- Is the performance of my database getting better, staying the same, or getting worse?

Capacity planning for databases typically involves two things:

- Historical trend analysis involves viewing selected database statistics over various time periods so that trends can easily be spotted.
- Forecasting entails using historical statistical information and projecting the values of various statistics (e.g. a database's size) will be in the future.

7.1 Performing Trend Analysis

PEM automatically collects a wide range of performance metrics about storage usage, memory usage, I/O traffic and more. The performance metrics are stored in a metadata repository that is created when PEM is installed.

The hard part of capacity planning operations (automatic data collection) is transparently handled for you by PEM.

All that remains is to use PEM's Capacity Manager interface to build desired trend analysis and forecasting reports.

Of course, you have full control over what and how often data is collected, but you can also take advantage of the product's defaults and have the recommended statistics gathered for you automatically.

To open the Capacity Manager dialog, select Capacity Manager... from the Management menu. The Capacity Manager dialog (shown in Figure 7.1) provides

quick access to a list of the available metrics, and the options for producing capacity planning reports.

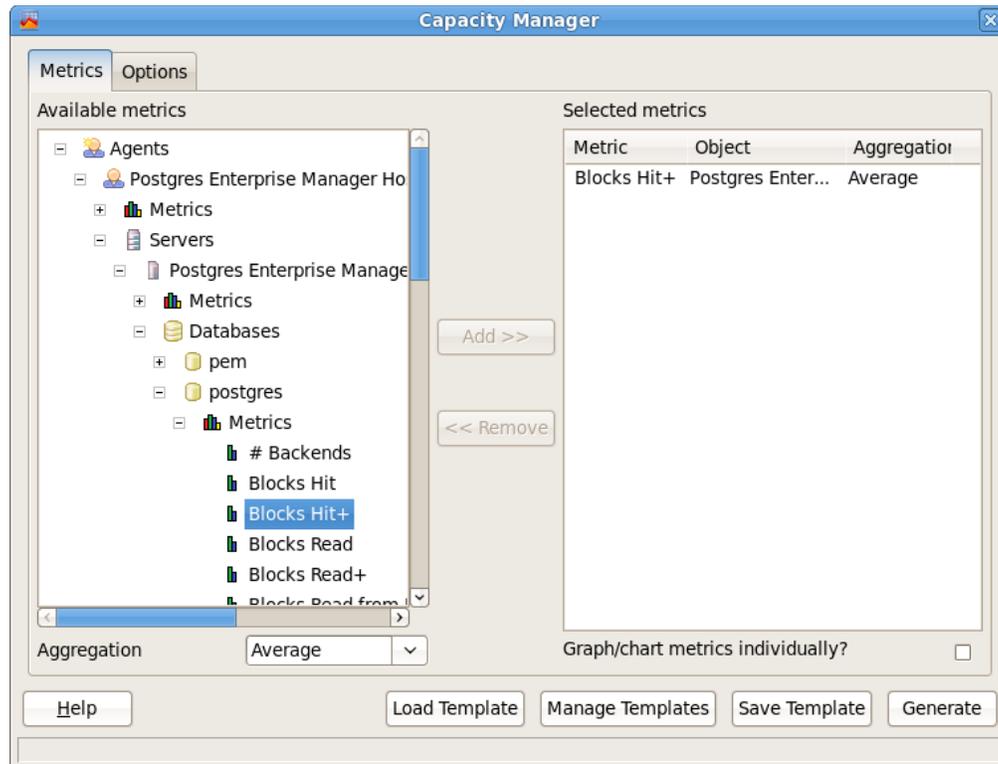


Figure 7.1 - The Capacity Manager dialog

The Capacity Manager dialog displays two tabs – Metrics and Options – that you can use to define capacity planning reports.

The Metrics tab displays a tree control that allows you to easily navigate all of your managed servers and select statistics that you wish to analyze. For example, to follow the growth of a particular database, you would expand that node in the tree control, highlight Database Size in the listed Metrics, and click the Add >> button to add it to the Selected metrics pane.

Capacity Manager can plot multiple statistics on one graph or produce a separate graph for each distinct metric. A checkbox located in the lower right corner of the Metrics tab of the Capacity Manager dialog enforces this option.

When you have specified the metrics that will be included in the analysis, you can then specify the timeframe over which the analysis will be performed. Use the fields on the Options tab of the Capacity Manager dialog, to specify the Time Period covered by the report (see Figure 7.2).

Capacity Manager allows you to create both graphical and tabular reports for historical trend analysis and future forecasts analysis reports. The `Options` tab provides fields that allow you to specify the form that the resulting report will take:

- A graph
- A data table
- Both a graph and data table

Finally, you can specify where Capacity Manager reports are displayed or written (see Figure 7.2). PEM will display the report in either a new or existing tab within the PEM client or write the report to a file on the host of your client workstation.

Figure 7.2 – Specify the time period, type, and destination of the Capacity Manager report.

Specify a Report destination, and click the `Generate` button to generate the report, as shown in Figure 7.3.

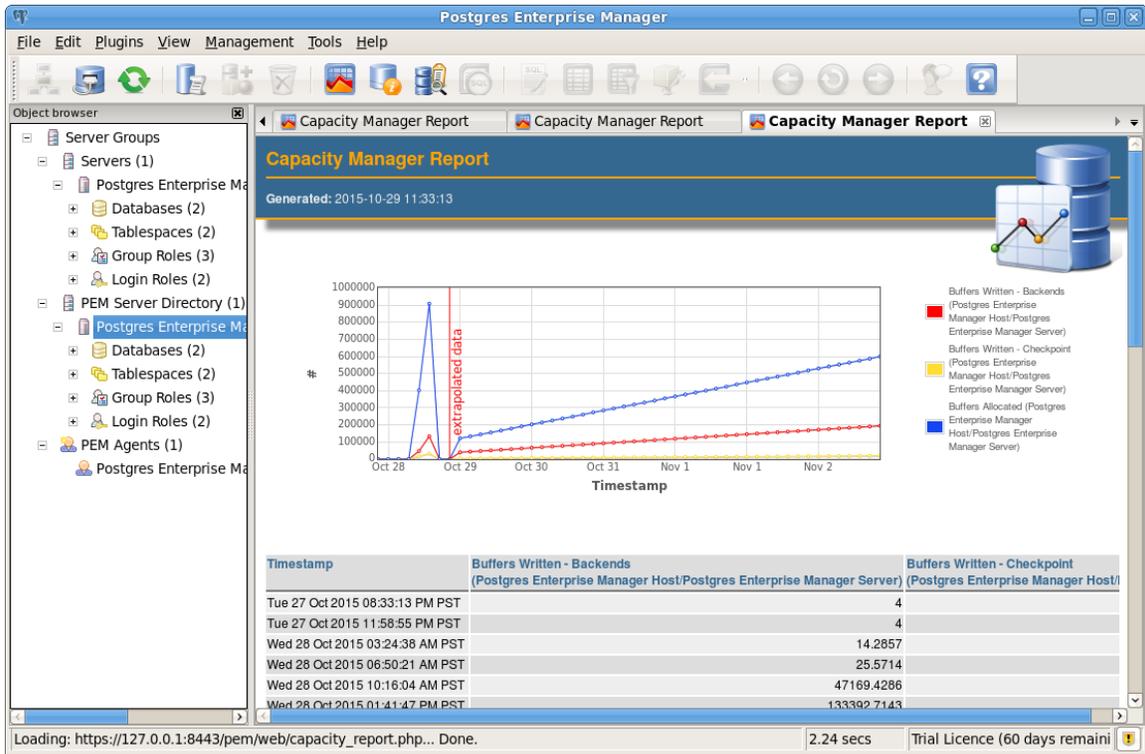


Figure 7.3 - The Capacity Manager report

7.2 Forecasting Future Trends

Capacity Manager uses historical metrics to forecast future trends. To create forecasting reports with Capacity Manager, simply select your desired metrics, and use the Capacity Manager Options tab to specify an End Time for the report that is in the future (see Figure 7.4).

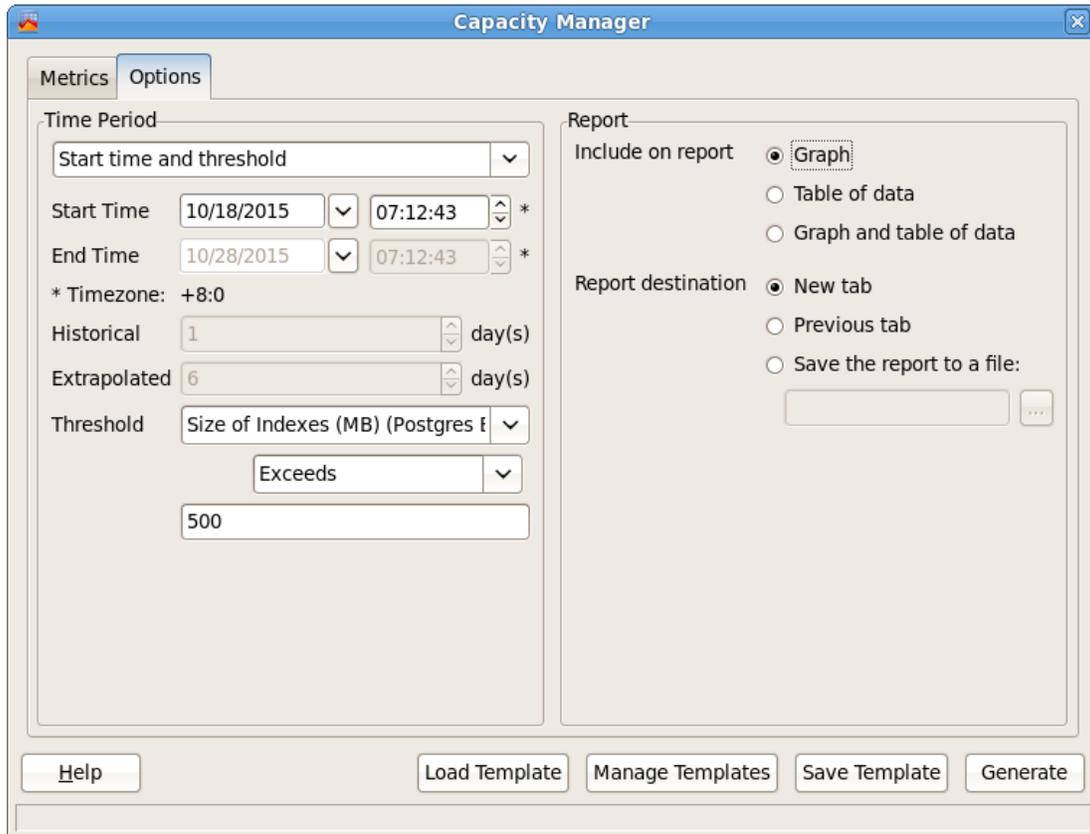


Figure 7.4 - Using the Options dialog to forecast future trends.

For example, you might use Capacity Manager to predict when you will need to increase the database storage available on your system. Use the drop-down listbox to select the projection criteria that PEM will use to extrapolate data. Select from:

- Start time and end time
- Start time and threshold
- Historical days and extrapolated days
- Historical days and threshold

After specifying the projection criteria, and specifying dates and thresholds for the report, click the `Generate` button. Capacity Manager will use historical usage data to predict your future resource requirements.

8 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.

8.1 Setting the Advanced Server Instance Service ID

To configure logging for an Advanced Server instance, the server must be registered as a PEM-managed server, and the registration information must include the name of a service script. When registering a new service, 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 8.1).

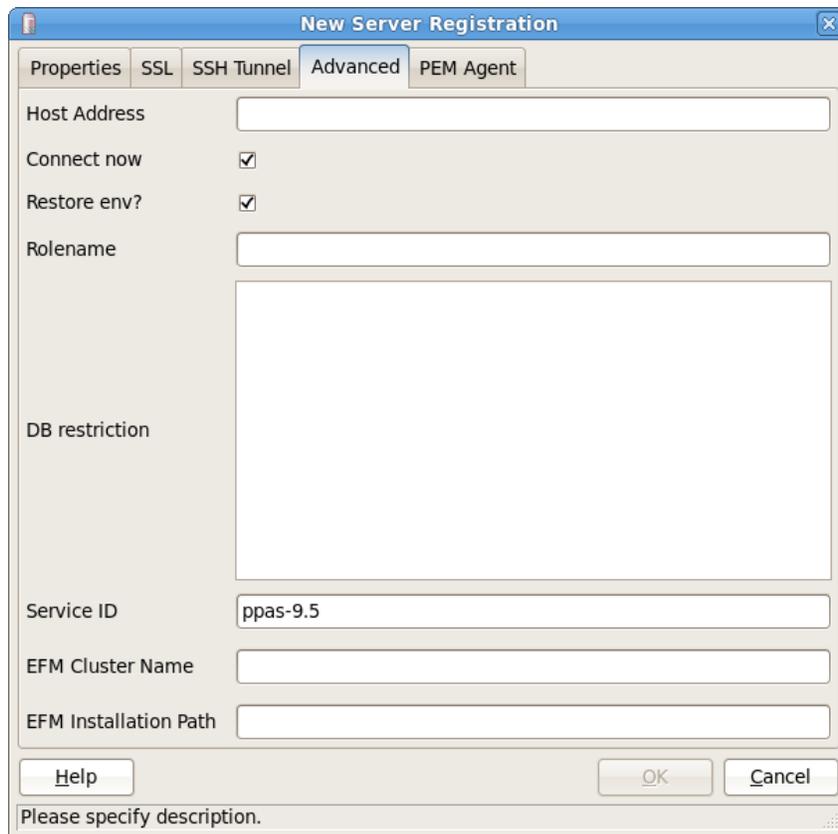


Figure 8.1 - The Service ID of the Advanced Server instance.

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

- The name of the Advanced Server version 9.5 service script is `ppas-9.5`.
- The name of the PostgreSQL version 9.5 service script is `postgresql-9.5`.

8.2 Setting the EDB Audit Configuration Probe

Before configuring audit logging of Advanced Server servers, you should enable the EDB Audit Configuration probe. To open the Probe Configuration dialog, right click on the name of a registered Advanced Server server in the tree control, and select Probe Configuration from the context menu. You can also access the Probe Configuration dialog by highlighting the Advanced Server name, and selecting Probe Configuration 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 8.2).

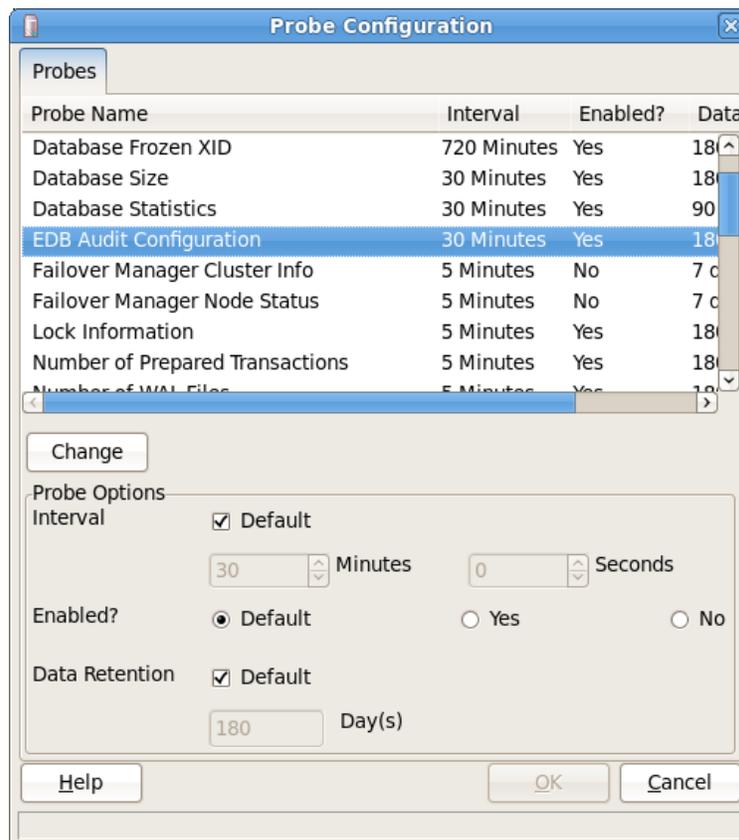


Figure 8.2 - The EDB Audit Configuration probe.

If EDB Audit Configuration is not enabled, use the Enabled? radio buttons on the Probe Options dialog to enable it.

8.3 Configuring Audit Logging with the Audit Manager

To open the Audit Manager, select `Audit Manager...` from the Management menu. The Audit Manager Welcome dialog opens as shown in Figure 8.3.

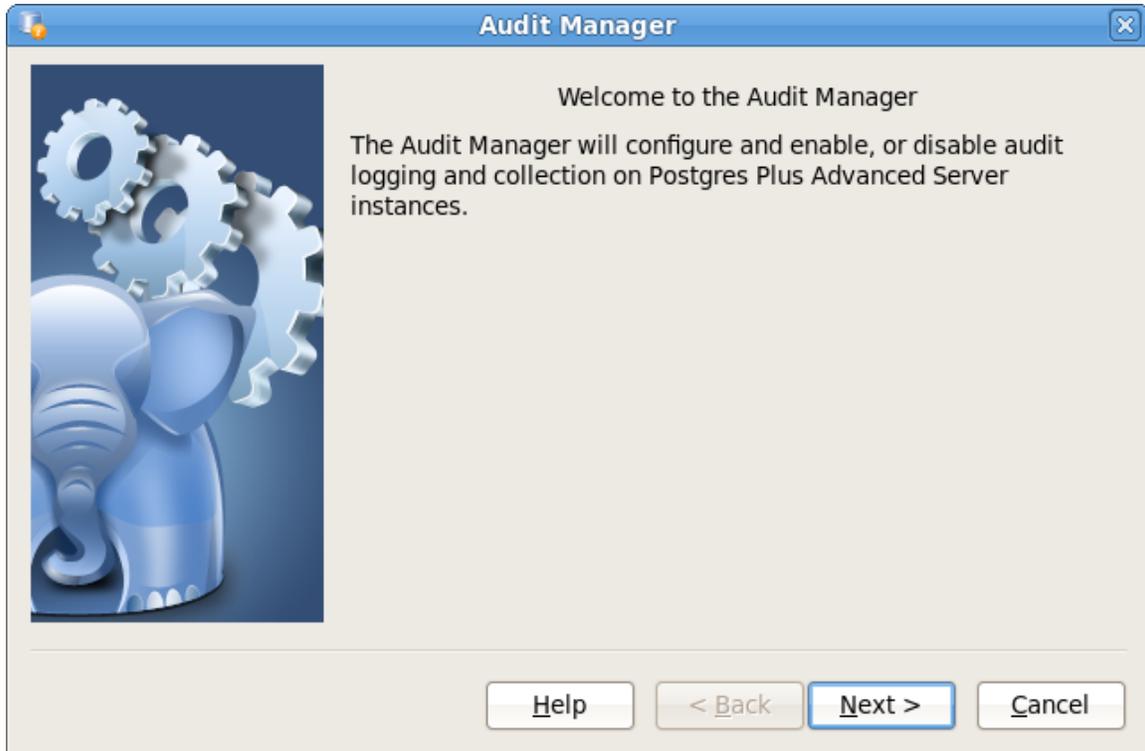


Figure 8.3 - The Audit Manager Welcome dialog.

Click `Next` to continue.

Select the servers you wish to configure for auditing (shown in Figure 8.4). Click `Next` to continue.

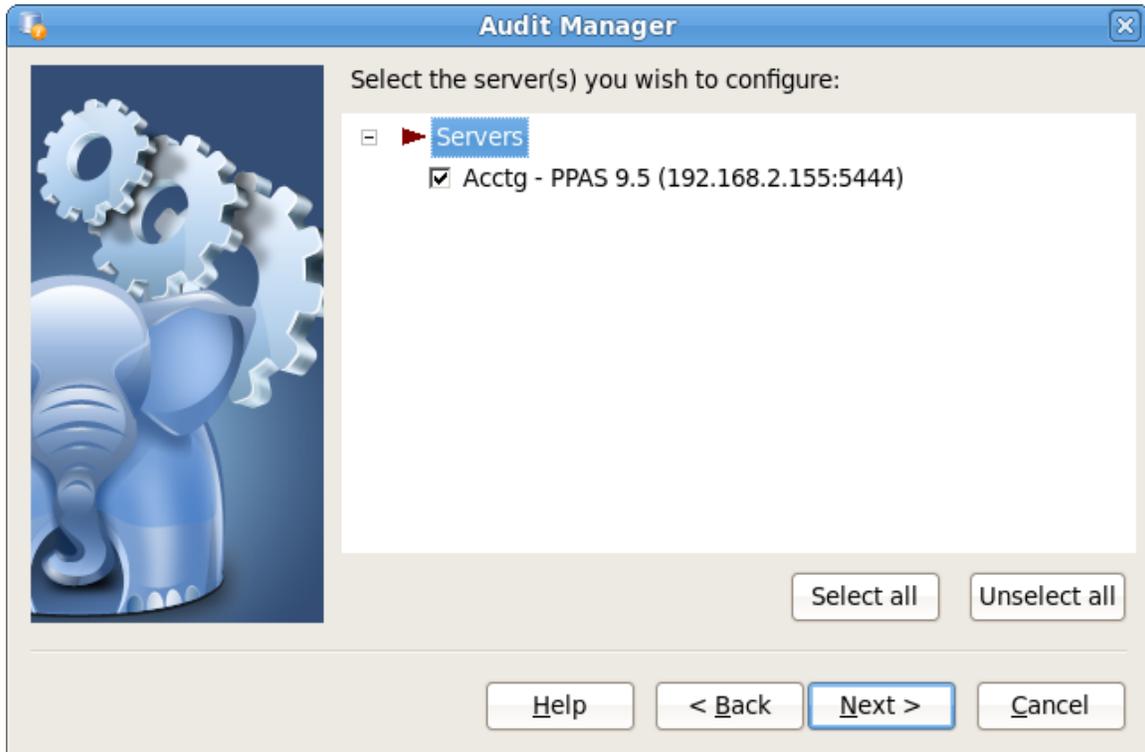


Figure 8.4 – Select the servers you wish to configure for auditing.

The Auditing Parameters Configuration dialog lets you enable or disable auditing and choose how often log records are collected into PEM (see Figure 8.5).



Figure 8.5 – The Auditing Parameters Configuration dialog.

Use the fields on the Auditing Parameters Configuration dialog to specify auditing preferences:

- Use the Auditing Status radio buttons to Enable or Disable auditing.
- Check the Enable Log Collection checkbox to instruct PEM to periodically gather the log records so you can later view them in the Audit Log dashboard. When enabled, the PEM agent will parse the audit logs, and store the result in the `pemdata.audit_logs` table on the PEM server.
- Use the Collection Frequency drop-down list to specify how often PEM should collect the log records.
- Use the Log Format radio buttons to specify the raw log format that will be written on each server. When Enable Log Collection is checked, PEM will use CSV format.
- 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 the `Audit File Name` to specify a format for the log file name. By default, the format is `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

Click `Next` to continue to the `Audit Log Configuration` dialog (see Figure 8.6).



Figure 8.6 – The Audit Log Configuration dialog.

Use the Audit Log Configuration dialog to determine the types of activities to be logged during auditing:

- Specify `All` to log all connection attempts, `Failed` to log only failed connection attempts, or `None` for no connection logging of Log Connection Attempts.
- Specify `All` to log all disconnection attempts or `None` for no disconnection logging of Log Disconnection Attempts.
- Check the Log Select Statements checkbox to log `SELECT` statements.
- Check the Log Error Statements checkbox to log SQL statements that result in an error.
- Check the Log DML Statements checkbox to log data manipulation language SQL statements such as `INSERT`, `UPDATE`, and `DELETE`.
- Check the Log DDL Statements checkbox to log data definition language SQL statements such as `CREATE`, `DROP`, and `ALTER`.

Click `Next` to continue to the Auditing Parameters Log Rotation dialog (see Figure 8.7).

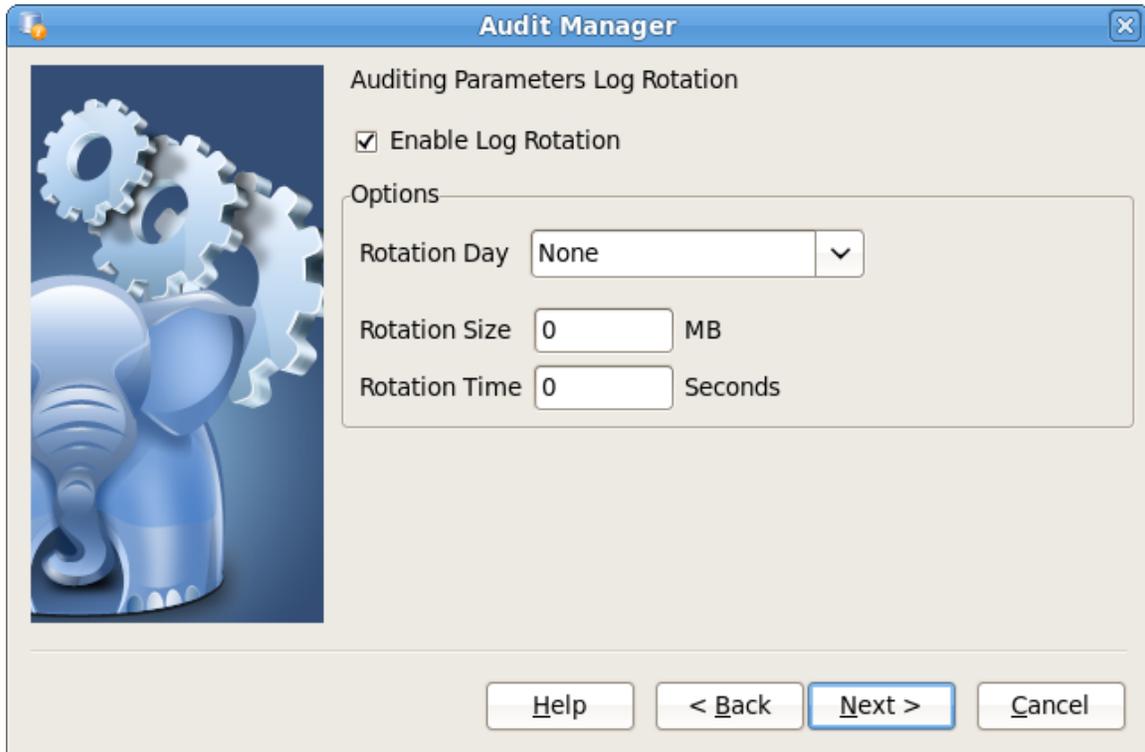


Figure 8.7 – The Auditing Parameters Log Rotation dialog.

Use the Auditing Parameters Log Rotation dialog to set factors controlling audit log file rotation.

- Check the `Enable Log Rotation` checkbox to periodically rotate the log file. If the log file is not rotated, all records will be saved in a single file that may grow to an unmanageably large size over time.
- Use the `Rotation Day` drop-down list to specify a rotation schedule for the log file. You can specify:
 - `Everyday` to instruct the server to rotate the log file each day
 - The name of a day on which the file will be rotated
 - `None` to indicate that log rotation should occur based on file size and/or length of time between rotations
- Use the `Rotation Size` field to specify the maximum size of the log file; the log file will be rotated upon reaching the given file size.

- Use the `Rotation Time` field to specify the length of time between rotations; the log file will be rotated after the specified the number of seconds have passed.

Click `Next` to continue to the `Schedule Auditing Changes` dialog (see Figure 8.8).



Figure 8.8 – The Schedule Auditing Changes dialog.

Use the `Schedule Auditing Changes` dialog to determine when auditing configuration changes are to take effect.

- Select `Configure Auditing 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.
- Select `Schedule it for some other time` if you want 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 Task Viewer to review a list of Scheduled jobs. To open the Task Viewer, right click on the name of a server or agent and select Scheduled Tasks from the context menu.

8.4 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 8.9).

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).

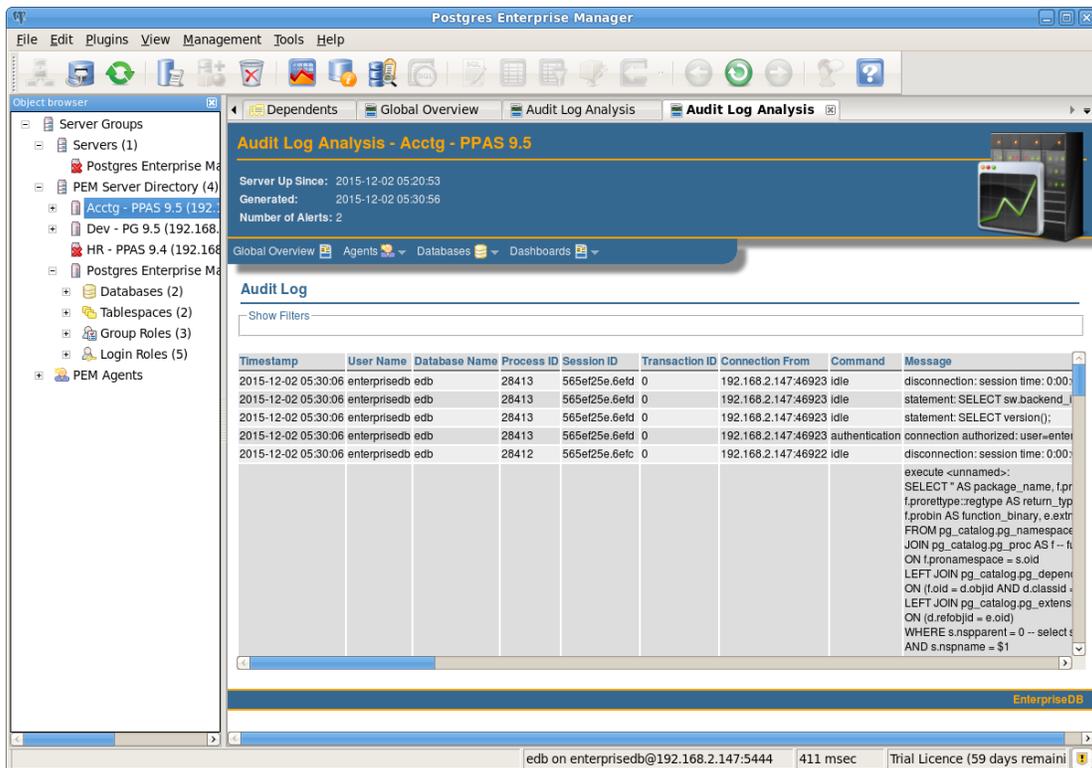


Figure 8.9 – The Audit Log dashboard.

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 8.10).

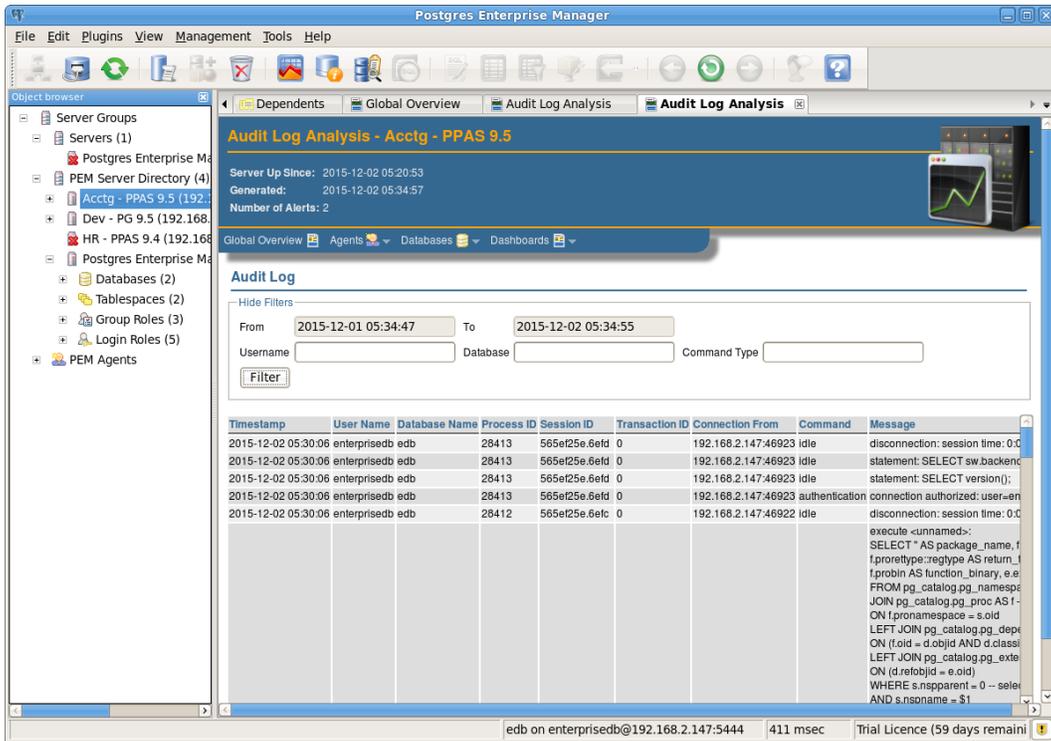


Figure 8.10 – The Audit Log dashboard filters panel.

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.

9 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. For more information, see [Section 8.1, *Setting the Advanced Server Instance Service ID*](#).

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 9.1).



Figure 9.1 - The Log Manager welcome dialog.

Click **Next** to continue to the **Servers** dialog (see Figure 9.2).

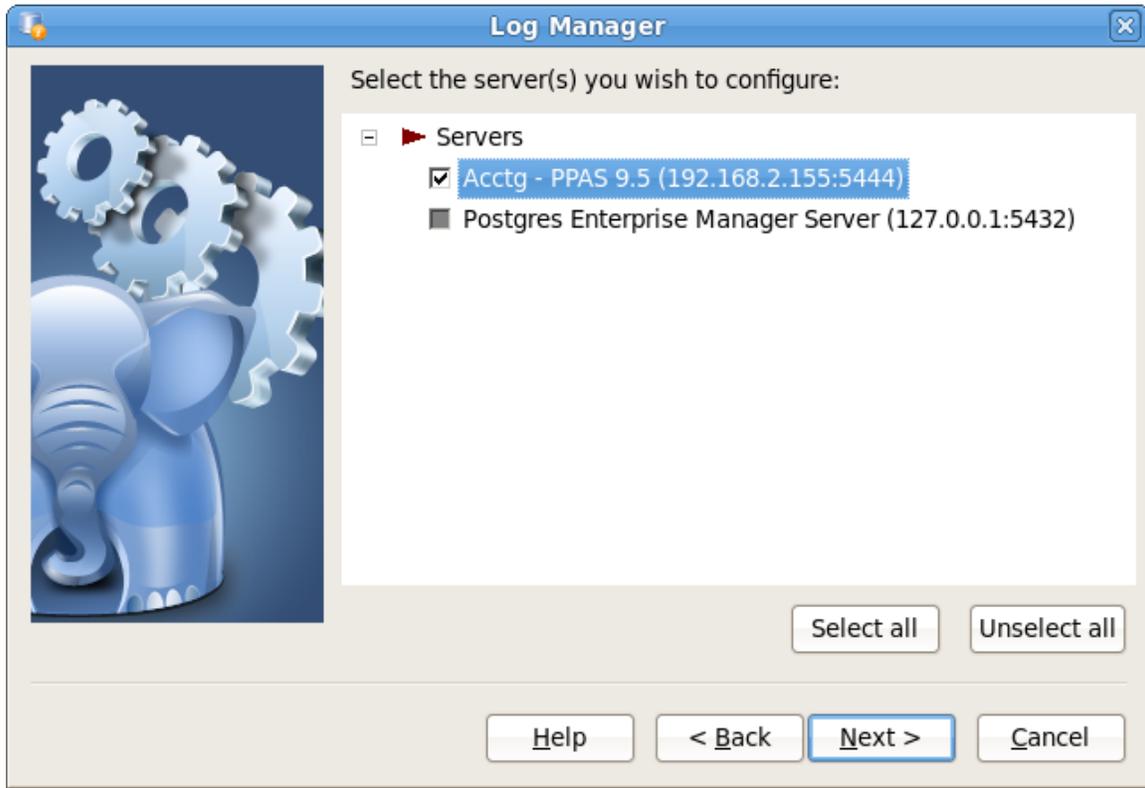


Figure 9.2 - The Log Manager server selection dialog.

The **Servers** dialog displays a list of the server connections monitored by PEM. Only those servers that specify a **Service ID** on the **Advanced** tab of the **Properties** dialog are active. Check the box next to the name of a server (or servers) to which the Log Manager wizard will apply the specified configuration.

Click **Next** to continue.

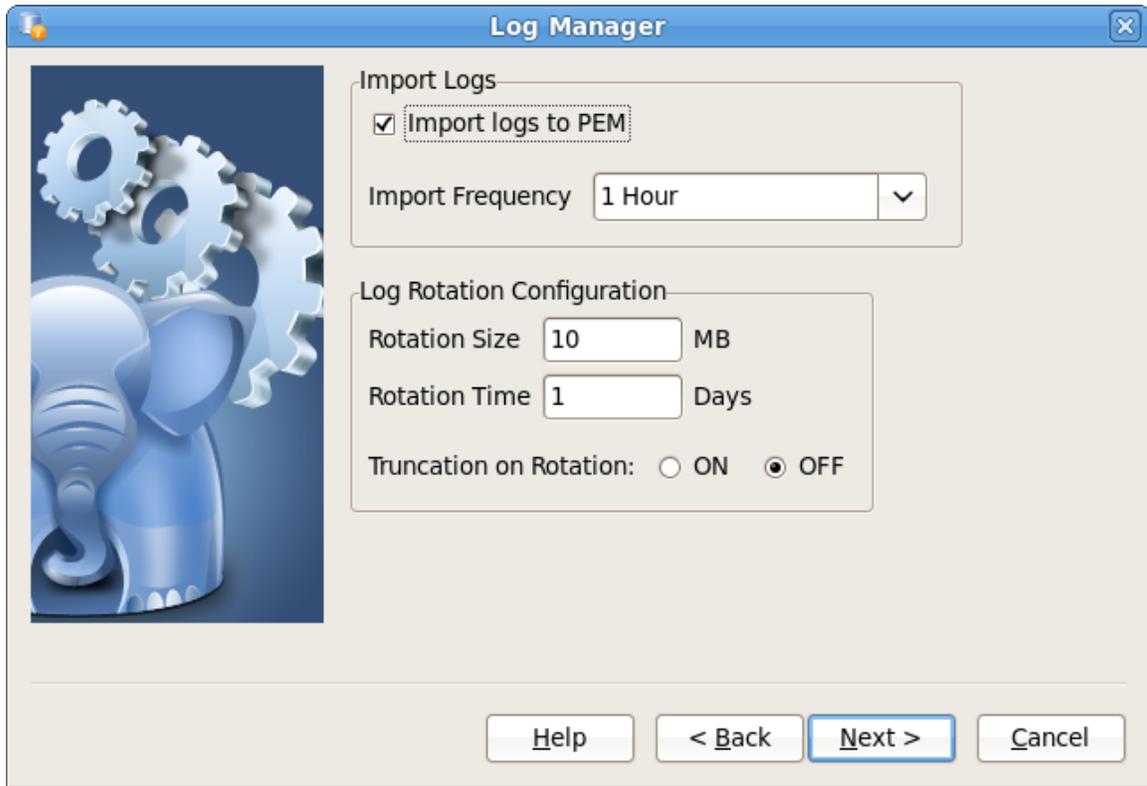


Figure 9.3 - The Log Manager scheduling dialog.

Use the options on the Scheduling dialog (as shown in Figure 9.3) to schedule log file import frequency, and to select a time that logging configuration changes will be applied.

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

- Check the box next to the `Import logs to PEM` label to specify that 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` panel to specify the maximum length (lifespan or size) of a log file:

- Use the `Rotation Size` 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 `Rotation Time` 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` radio buttons 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.



Figure 9.4 – The Where to Log dialog.

Use the fields on the `Where to Log` dialog (shown in Figure 9.4) 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:

- Check the box next to `stderr` to specify that log files should be written to `stderr`.
- Check the box next to `csvlog` 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.
- Check the box next to `syslog` to specify that log files should be written to the system log files.
- On Windows, check the box next to `eventlog` to specify that log files should be written to the event log.
- Use the `Log Collector` radio buttons to instruct the server to re-direct captured log messages (directed to `STDERR`) into log files.
- Use the `Log Silent Mode` radio buttons to instruct the server to run silently in the background, disassociated from the controlling terminal.
- Use options in the `Log Directory` box 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. By 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:

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

Click `Next` to continue.

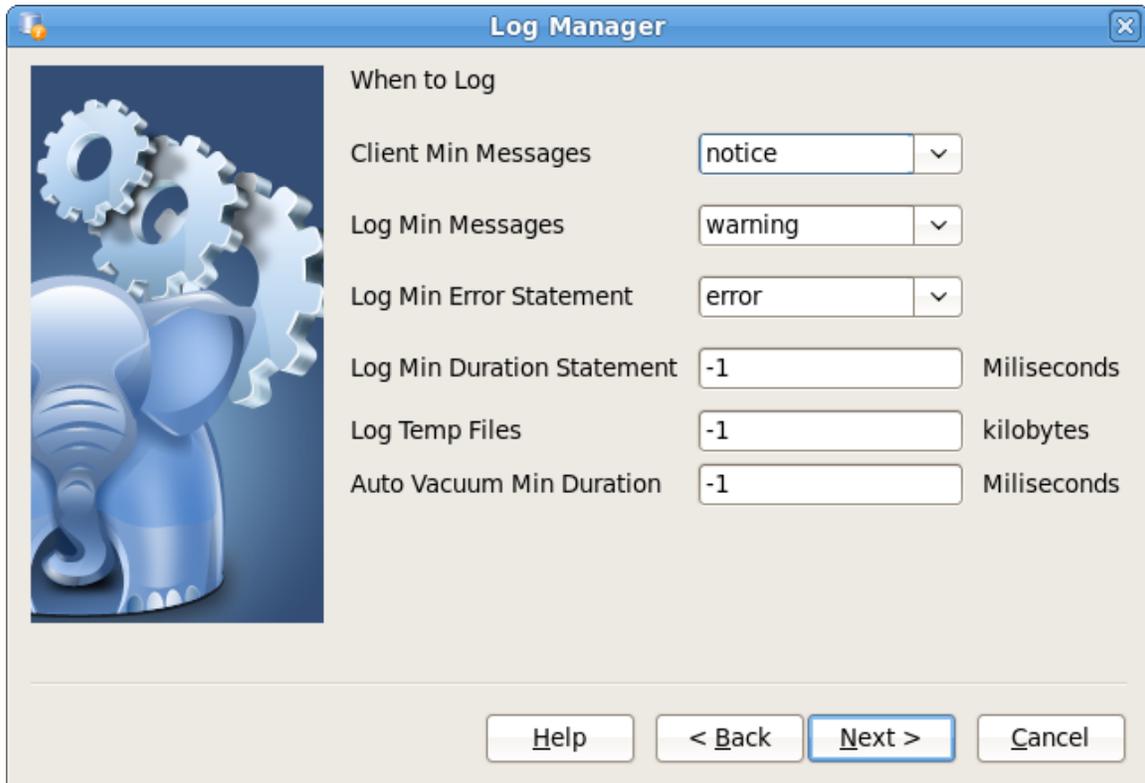


Figure 9.5 - The Log Manager When to Log dialog.

Use the fields on the `When to Log` dialog (shown in Figure 9.5) 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 which severity levels are sent to the client application.
- Use the `Log Min Messages` drop-down list box to specify which severity levels are 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 `Auto Vacuum 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 9.6 - The Log Manager What to Log dialog.

Use the fields on the `What to Log` dialog (shown in Figure 9.6) to specify log entry options that are useful for debugging and auditing.

The checkboxes 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:

- Check the box next to `Parse Tree` to instruct the server to include the parse tree in the log file.
- Check the box next to `Rewriter Output` to instruct the server to include query rewriter output in the log file.
- Check the box next to `Execution Plan` 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` option is enabled, 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 options in the `General Options` box to instruct the server to include auditing information in the log file:

- Check the box next to `Checkpoints` to include checkpoints and restartpoints in the server log.
- Check the box next to `Connections` to include each attempted connection to the server (as well as successfully authenticated connections) in the server log.
- Check the box next to `Disconnections` to include a server log entry for each terminated session that provides the session information and session duration.
- Check the box next to `Duration` to include the amount of time required to execute each logged statement in the server log.
- Check the box next to `Hostname` 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.
- Check the box next to `Lock Waits` 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.

Click `Next` to continue.



Figure 9.7 - The Log Manager What to Log dialog.

Use the fields on the second `What to Log` dialog (shown in Figure 9.7) to specify the amount of information written to the log files.

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/9.5/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 9.8 - The Schedule Logging Changes dialog.

Use options on the `Schedule Config Changes` dialog (see Figure 9.8) to specify when logging configuration changes will be applied:

- Select the `Configure Logging Now` radio button to specify that the server will restart when you have completed the Log Manager wizard.
- Select the `Schedule it for some other time` radio button to enable date and time list boxes; use the date and time selectors to specify a convenient time for the server to restart.

Note that when you apply the configuration changes specified by the Log Manager wizard, the server will be restarted, 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. You can use the `Scheduled Tasks` dialog to confirm that the configuration file update and server restart have been scheduled or performed as expected. To open the `Scheduled Tasks` dialog, right-click on the name of the selected server in the PEM client tree control, and select `Scheduled Tasks...`

9.1 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.

The screenshot displays the 'Server Log Analysis' dashboard in the Postgres Enterprise Manager. The dashboard header includes the title 'Server Log Analysis - Acctg - PPAS 9.5' and key metrics: 'Server Up Since: 2015-12-02 05:51:49', 'Generated: 2015-12-02 06:15:32', and 'Number of Alerts: 1'. Below the header is a 'Server Log' section with a 'Show Filters' button and a table of log entries. The table has the following columns: Timestamp, User Name, Database Name, Process ID, Session ID, Transaction ID, Connection From, Command, and Message. The log entries show various database activities such as 'SELECT', 'authentication', and 'disconnection'.

Timestamp	User Name	Database Name	Process ID	Session ID	Transaction ID	Connection From	Command	Message
2015-12-02 06:10:08	enterprisedb	edb	30711	565efbbf.7717	0	192.168.2.147:48088	idle	disconnection: session time: 0:00:00.096 user=er...
2015-12-02 06:10:08	enterprisedb	edb	30711	565efbbf.7717	0	192.168.2.147:48088	SELECT	duration: 21.220 ms
2015-12-02 06:10:07	enterprisedb	edb	30711	565efbbf.7717	0	192.168.2.147:48088	SELECT	duration: 0.387 ms
2015-12-02 06:10:07	enterprisedb	edb	30709	565efbbf.7715	0	192.168.2.147:48086	idle	disconnection: session time: 0:00:00.208 user=er...
2015-12-02 06:10:07	enterprisedb	edb	30709	565efbbf.7715	0	192.168.2.147:48086	SELECT	duration: 4.737 ms
2015-12-02 06:10:07	enterprisedb	edb	30711	565efbbf.7717	0	192.168.2.147:48088	authentication	connection authorized: user=enterprisedb databa...
2015-12-02 06:10:07	enterprisedb	edb	30710	565efbbf.7716	0	192.168.2.147:48087	idle	disconnection: session time: 0:00:00.198 user=er...
2015-12-02 06:10:07	enterprisedb	edb	30710	565efbbf.7716	0	192.168.2.147:48087	SELECT	duration: 0.850 ms
2015-12-02 06:10:07	enterprisedb	edb	30711	565efbbf.7717	0	192.168.2.147:48087	SELECT	connection received: host=192.168.2.147 port=48...
2015-12-02 06:10:07	enterprisedb	edb	30710	565efbbf.7716	0	192.168.2.147:48087	SELECT	duration: 0.307 ms
2015-12-02 06:10:07	enterprisedb	edb	30710	565efbbf.7716	0	192.168.2.147:48087	authentication	connection authorized: user=enterprisedb databa...
2015-12-02 06:10:07	enterprisedb	edb	30710	565efbbf.7716	0	192.168.2.147:48087	authentication	connection received: host=192.168.2.147 port=48...
2015-12-02 06:10:07	enterprisedb	edb	30709	565efbbf.7715	0	192.168.2.147:48086	SELECT	duration: 0.372 ms
2015-12-02 06:10:07	enterprisedb	edb	30709	565efbbf.7715	0	192.168.2.147:48086	authentication	connection authorized: user=enterprisedb databa...
2015-12-02 06:10:07	enterprisedb	edb	30709	565efbbf.7715	0	192.168.2.147:48086	authentication	connection received: host=192.168.2.147 port=48...
2015-12-02 06:10:07	enterprisedb	edb	30708	565efbbf.7714	0	192.168.2.147:48085	idle	disconnection: session time: 0:00:00.138 user=er...

Figure 9.9 - The Server Log Analysis dashboard.

The header information on the Server Log Analysis dashboard (shown in Figure 9.9) 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 9.10).

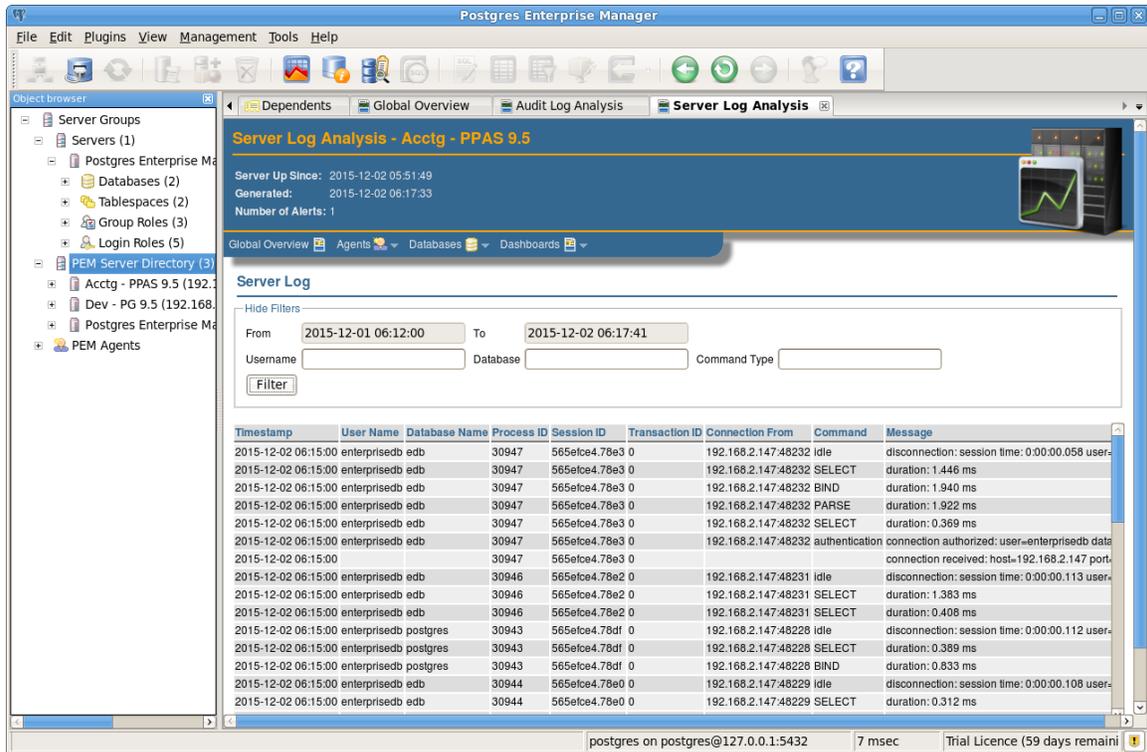


Figure 9.10 - Defining a Server Log filter.

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.

10 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. For more information about using the `Log Manager` wizard, see [Section 9](#).

To open the `Log Analysis Expert` wizard select the `Postgres Log Analysis Expert` option from the `Management` menu in the PEM client. When the wizard's `Welcome` dialog (see [Figure 10.1](#)) opens, click `Next` to continue.

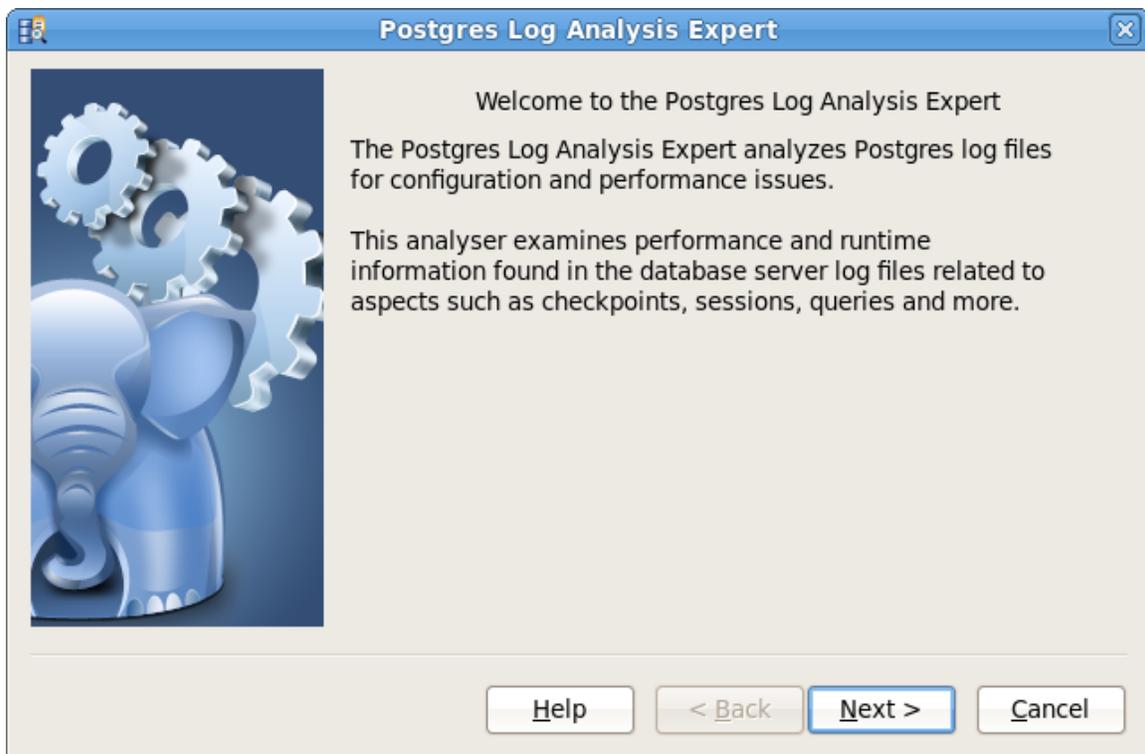


Figure 10.1 - The Log Analysis Expert Welcome dialog.

The wizard displays a list of `Analyzers` from which you can select (see [Figure 10.2](#)). Each `Analyzer` generates a corresponding table, chart, or graph that contains information gleaned from the log reports.

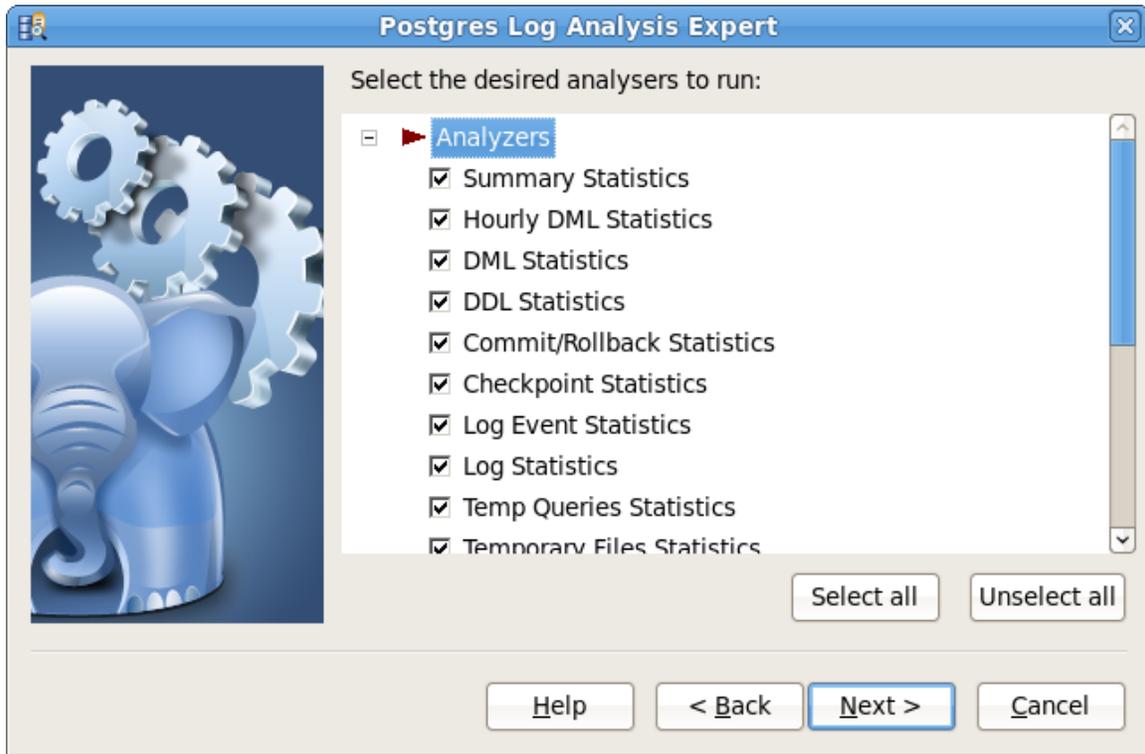


Figure 10.2 - The Log Analysis Expert Analyzers list.

Check the box to the left of an Analyzer to indicate that the Log Analysis Expert should prepare the corresponding table, chart or graph. You can also:

- Click the `Unselect all` button to un-check all of the boxes before navigating through the list and selecting only the tables, charts or graphs that you wish Log Analysis Expert to generate.
- Click the `Select all` button to instruct Log Analysis Expert to review the server logs and generate a report for each analyzer selected.

After making your selections, click `Next` to continue to the `Servers` tree control (see Figure 10.3).

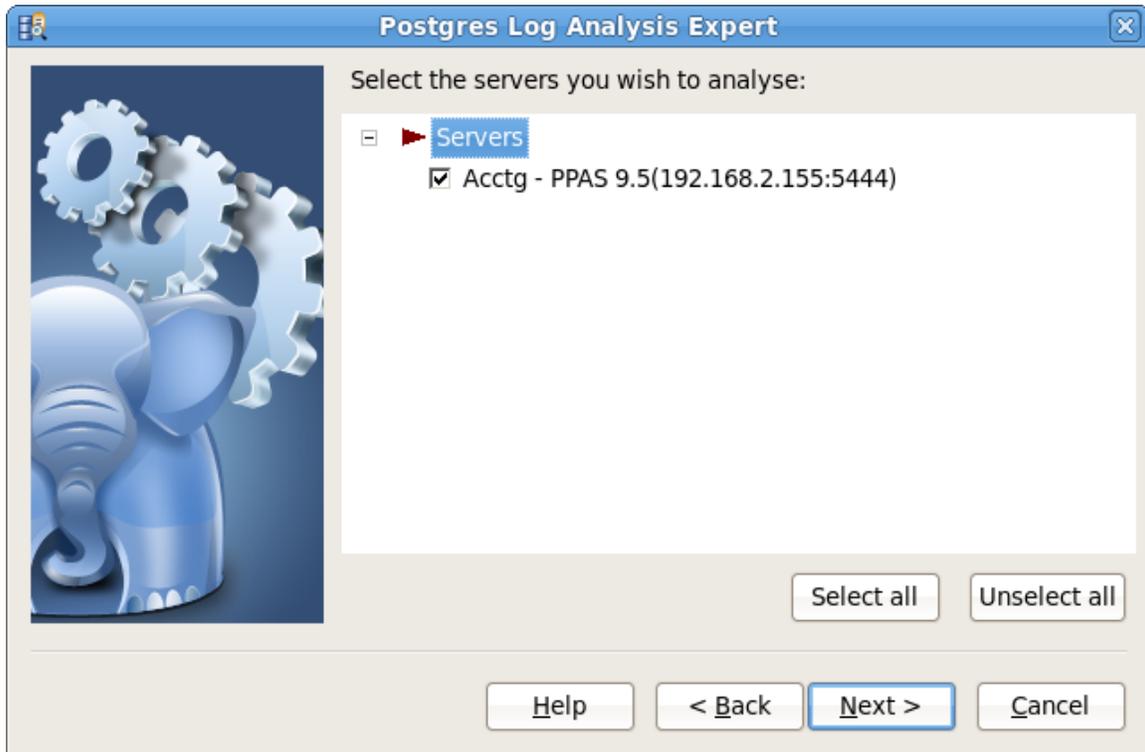


Figure 10.3 - The Servers selection dialog.

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.

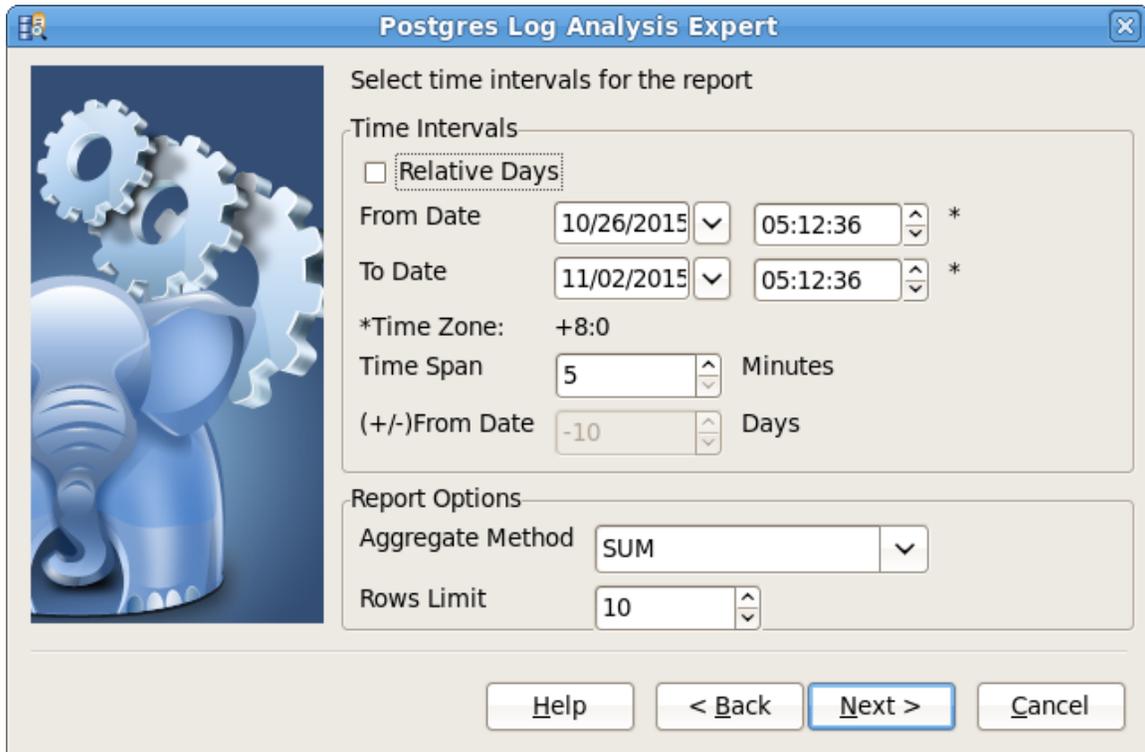


Figure 10.4 - Specify a time range for analysis.

Use the fields in the `Time Intervals` section of the dialog (see Figure 10.4) to specify the time range that the Log Analysis Expert is supposed to analyze:

- Check the box next to `Relative Days` to enable the `(+/-)From date` field and specify the number of days before or after the date and time selected in the `From Date` field.
- Use the `From Date` field to specify the starting date and time for the analysis.
- Use the `Time Span` selector 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 `(+/-) From Date` selector to indicate the number of days before or after the `From Date` that should be included in the analysis.

- Use the options in the `Report Options` section of the dialog 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 - 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 `Rows Limit` selector to indicate the maximum number of rows to include in a table.

Click `Next` to continue.

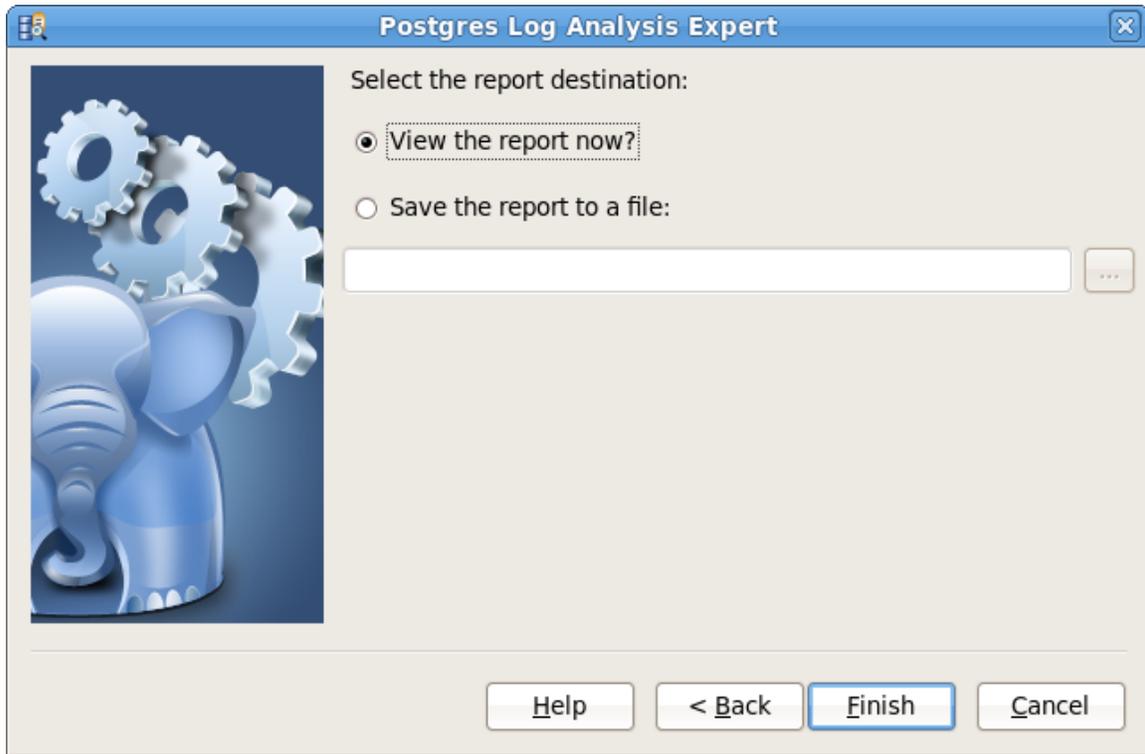
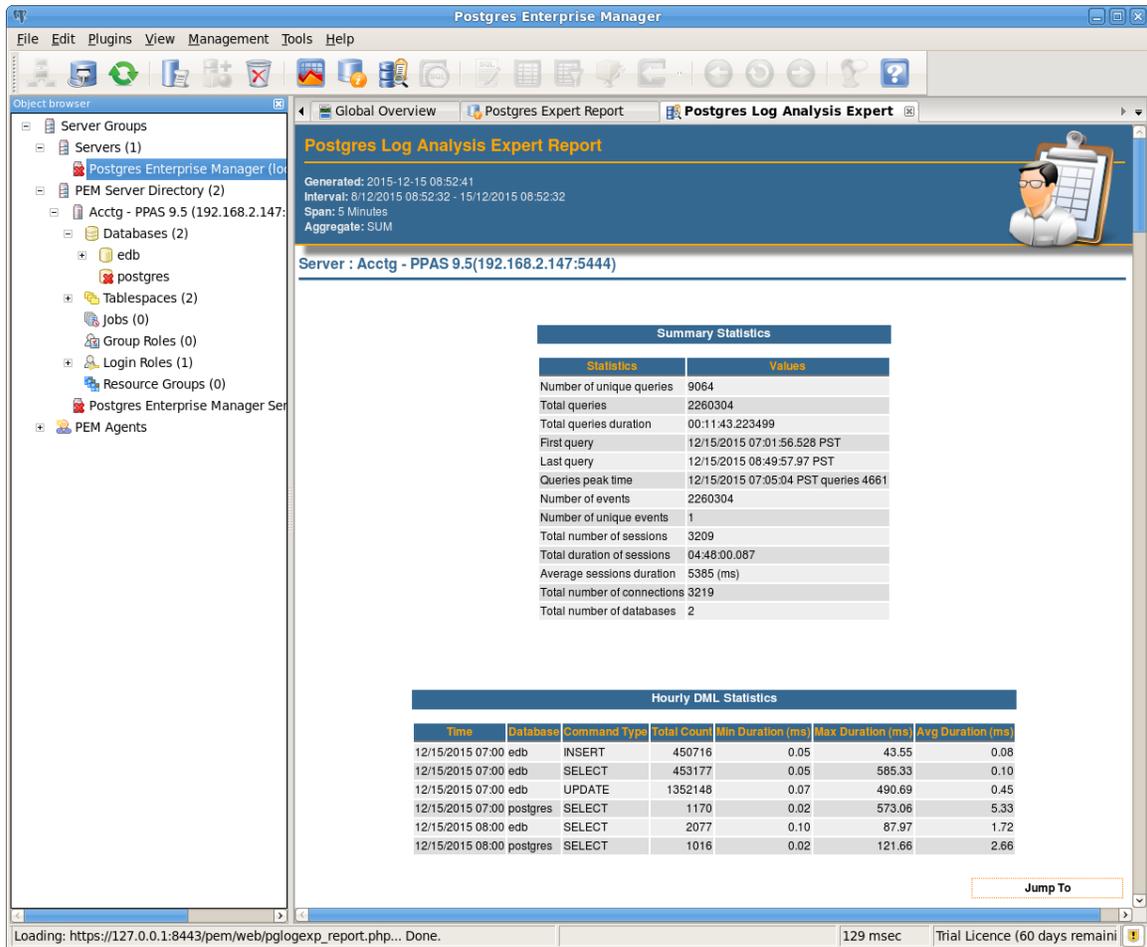


Figure 10.5 - Specify a report destination.

You can select the default option and click `Finish` to view the Log Analysis Expert report, or check the box next to `Save the report to a file` to save a copy of the report to an HTML file for later use (see Figure 10.5). If you wish to save the report to a file, specify a filename in the field provided; alternatively, use the button to the right of the field to open a browser dialog and use a browser to specify a log file location.

10.1 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. If the report contains an analysis of more than one monitored server, the graphs will be displayed in sets; first the graphs, tables and charts that display statistics for one server, then the graphs for the next server in the report.



The screenshot displays the Postgres Enterprise Manager interface with the 'Postgres Log Analysis Expert Report' open. The report is generated for the server 'Acctg - PPAS 9.5(192.168.2.147:5444)'. The report includes the following data:

Generated: 2015-12-15 08:52:41
Interval: 8/12/2015 08:52:32 - 15/12/2015 08:52:32
Span: 5 Minutes
Aggregate: SUM

Server : Acctg - PPAS 9.5(192.168.2.147:5444)

Summary Statistics	
Statistics	Values
Number of unique queries	9064
Total queries	2260304
Total queries duration	00:11:43.223499
First query	12/15/2015 07:01:56.528 PST
Last query	12/15/2015 08:49:57.97 PST
Queries peak time	12/15/2015 07:05:04 PST queries 4661
Number of events	2260304
Number of unique events	1
Total number of sessions	3209
Total duration of sessions	04:48:00.087
Average sessions duration	5385 (ms)
Total number of connections	3219
Total number of databases	2

Hourly DML Statistics						
Time	Database	Command Type	Total Count	Min Duration (ms)	Max Duration (ms)	Avg Duration (ms)
12/15/2015 07:00	edb	INSERT	450716	0.05	43.55	0.08
12/15/2015 07:00	edb	SELECT	453177	0.05	585.33	0.10
12/15/2015 07:00	edb	UPDATE	1352148	0.07	490.69	0.45
12/15/2015 07:00	postgres	SELECT	1170	0.02	573.06	5.33
12/15/2015 08:00	edb	SELECT	2077	0.10	87.97	1.72
12/15/2015 08:00	postgres	SELECT	1016	0.02	121.66	2.66

Jump To

Loading: https://127.0.0.1:8443/pem/web/pglogexp_report.php... Done. 129 msec Trial Licence (60 days remaini)

Figure 10.6 - The Log Analysis Expert Report.

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

11 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.

PEM provides the SQL Profiler component to assist in both locating and optimizing poorly-running SQL code. Users of Microsoft SQL Server's Profiler will find PEM's SQL Profiler very similar in operation and capabilities.

11.1 Setup and Configuration

Before using SQL Profiler, you must:

1. Download and install the SQL Profiler product into each managed database instance you wish to profile.
2. Edit the `postgresql.conf` parameter file and 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
```

3. 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 in each database that you will use the SQL Profiler against.
4. Stop and re-start the server for the changes to take effect.

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>

11.2 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.

Capturing a new SQL trace is very simple. To open SQL Profiler, select the *SQL Profiler* menu option from the *Management* menu. Upon startup, SQL Profiler will present you with a dialog asking if you would like to start a new trace, or open an existing trace.

If you choose to start a new trace (capturing SQL statements executed against the currently selected server) SQL Profiler will open the *Create New SQL Trace* dialog, as shown in Figure 11.1.

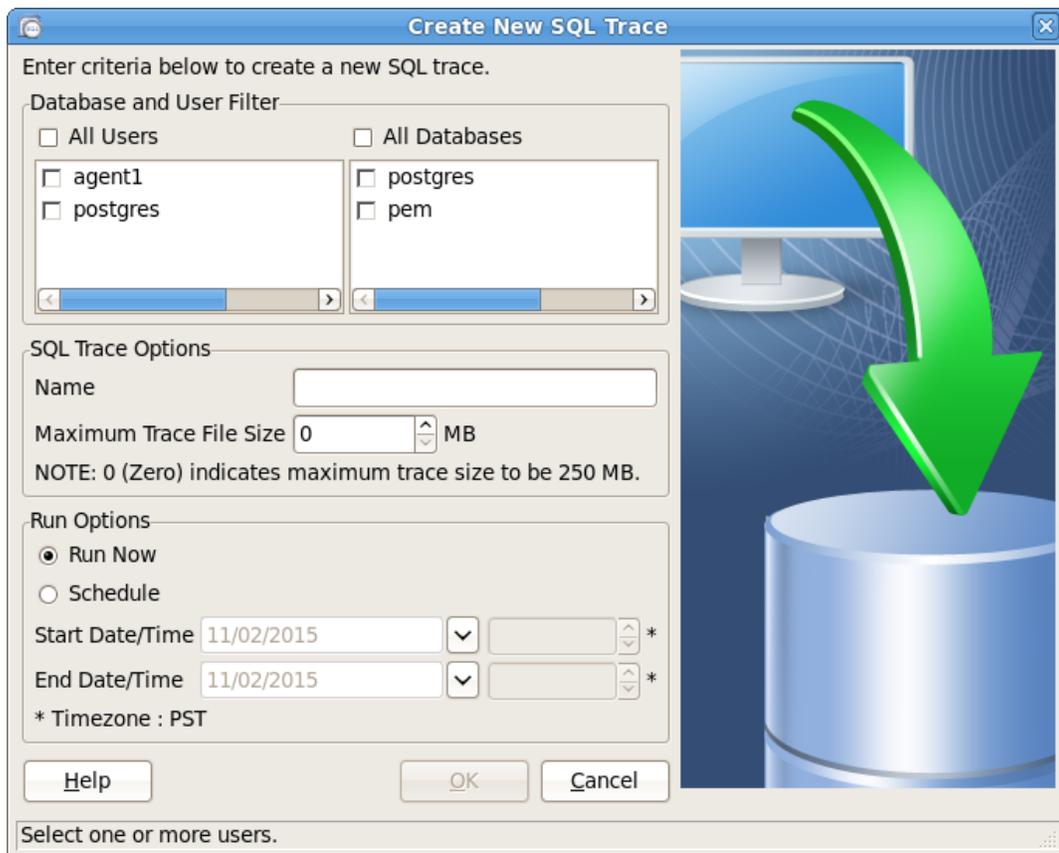


Figure 11.1 - The Create New SQL Trace dialog

Use fields on the dialog to selectively filter the captured SQL statements by user account and by database. You may choose to limit your trace, or to capture every SQL statement sent by all users against all databases.

You can also name your trace for future reference and set a maximum file size for the resulting trace file. Lastly, you can choose to begin capturing SQL immediately in an ad-hoc fashion (if you choose this option, you must manually stop the trace), or you can schedule the trace to run at a later time.

Scheduling a trace is particularly beneficial for capturing workloads during off hours. For example, you may want to capture the SQL statements that are executed against a server from 2 - 4 am. You can configure your trace to run during that timeframe and then examine the output at your leisure when you arrive at work in the morning.

When you are satisfied with all your selections, click the OK button to start the trace.

11.3 Analyzing a SQL Trace Output

When you start a new trace, SQL Profiler displays the collected trace data in the top pane of the SQL Profiler dialog. Initially the output will be blank; to see SQL statements that have been captured thus far, click the Refresh toolbar icon, or select Refresh from the Trace menu. If the toolbar is not visible, select Tool Bar from the View menu to add the toolbar to the display.

The screenshot shows the SQL Profiler - Resource Trace (20151102161034000000) window. The top pane displays a table of trace data with columns for Start Time, Duration, and Query. The Properties pane on the left shows metrics for the selected query, and the Explain pane on the right shows a graphical plan for the selected query.

	Start Time	Duration	Query
104	2015-11-02 08:14:29.216108-08	0.15	SELECT c.relname AS index name, r.relname AS table name, i.indk
105	2015-11-02 08:14:29.228927-08	0.012	SELECT version();
106	2015-11-02 08:14:29.230493-08	0.131	SELECT c.relname AS table name, c.relhaskey AS has primary key
107	2015-11-02 08:14:39.275292-08	0.013	SELECT version();
108	2015-11-02 08:14:39.278957-08	5.609	SELECT
109	2015-11-02 08:14:49.352683-08	0.013	SELECT version();
110	2015-11-02 08:14:49.355136-08	0.142	SELECT c.relname AS table name, c.relhaskey AS has primary key

The Properties pane shows the following metrics:

- Executed (#): 15
- Execution (%): 0
- Duration (%): 0
- Rows affected(%): 0
- Page faults (%): 0
- Page reclaims (%): 0
- Swaps (%): 0
- File system in (%): 0

The Explain pane shows a graphical plan with the following components:

- pg_namespace
- Hash
- pg_class
- Hash Join

Figure 11.2 - The SQL Profiler trace dialog

The SQL Profiler trace viewer includes three panes, shown in Figure 11.2:

- The top pane displays SQL statements collected in the trace.
- The lower-left pane displays the full SQL text and statistics for the statement highlighted in the top pane.
- The lower-right pane displays either a graphical or text-based EXPLAIN plan for the statement highlighted in the top pane.

SQL Profiler provides a number of features that will help you find poorly-running SQL code on your servers. Some of the most useful options are:

- To display key statistical information collected by SQL Profiler, open the `Select Columns` dialog by selecting the `Columns` option from the `View` menu. The `Select Columns` dialog allows you to customize the columns SQL Profiler displays in the trace output.
- Sorting data in the `Trace Data` pane is easy; just double-click the column header and SQL Profiler will sort the data within the table by that column. Double-click the column header again to reverse the sort order. This allows you to find, for example, the longest running SQL statement very quickly.
- You can filter the contents of a trace to further restrict which SQL statements are displayed. To open the `Trace Filter` dialog, click the `Filter` toolbar button, or select the `Filter` menu option from the `Edit` menu. You can use the `Trace Filter` dialog to construct a filter to display only the SQL statements that match your specific criteria.
- The `Properties` pane displays the complete SQL statement for the currently highlighted row in the `Trace Data` pane; you can use cut and paste to copy the SQL statement to the `Query Tool SQL IDE` for re-working.
- Select the `Metrics` tab in the `Properties` pane to display a variety of statistics for the currently highlighted statement. The metrics include information about the number of times a statement has been executed, the overall percentage of execution time vs. all other collected statements, and more.
- The graphical `Explain` pane, provides a graphical interpretation of the execution plan for the selected query. Click a segment of the graph to display statistics for that portion of the query plan.
- To export the contents of a trace to file, choose the `Export` option from the `File` menu. The `Export data to file` dialog allows you to re-format the trace data as it is saved to file.

To stop an active trace, click the `Stop Trace` toolbar button, or select `Stop Trace` from the `Trace` menu.

To view the contents of a previous trace, click on the `Open Trace` toolbar button, or select `Open Trace` from the `File` menu, and select a saved trace. SQL Profiler will close the current trace, and display the newly selected trace.

11.4 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.

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

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).

Note again, that the Index Advisor is only included in EnterpriseDB's Advanced Server, versions 9.0 and above.

12 Tuning Wizard

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

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

To run the Tuning Wizard, select the `Tuning Wizard` option from the `Management` menu of the PEM client. The Tuning Wizard opens as shown in Figure 12.1.

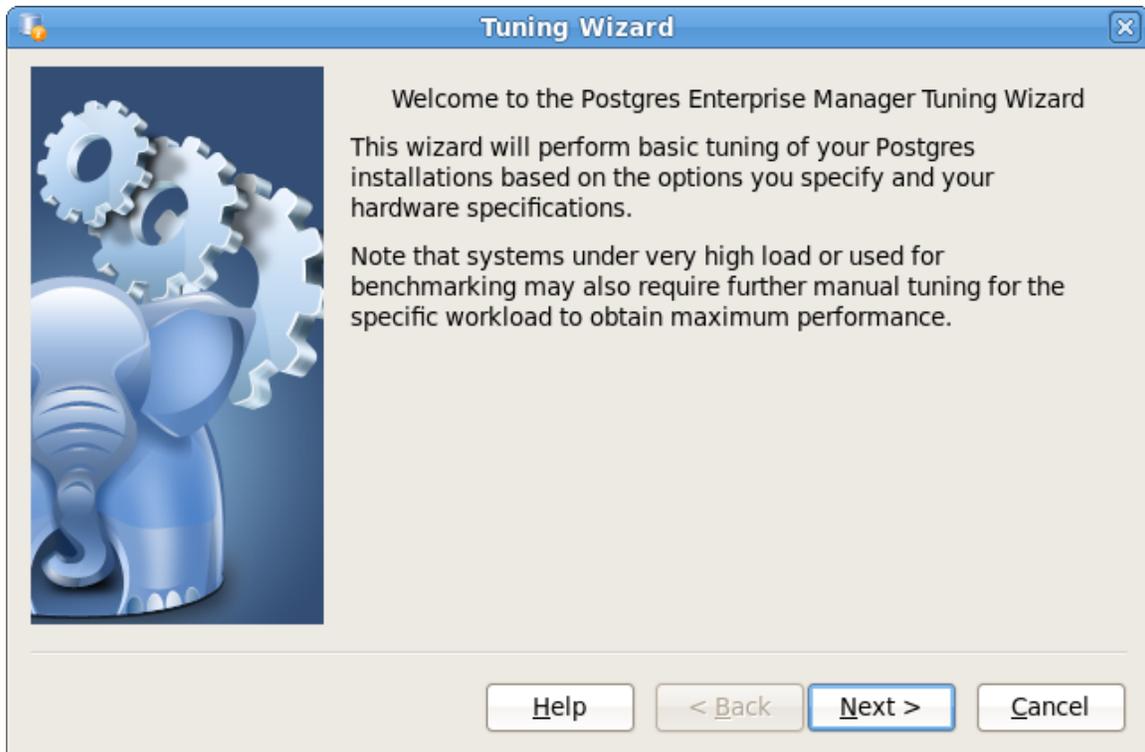


Figure 12.1 - The Tuning Wizard Welcome dialog

Click `Next` to continue to the server selection dialog (shown in Figure 12.2).

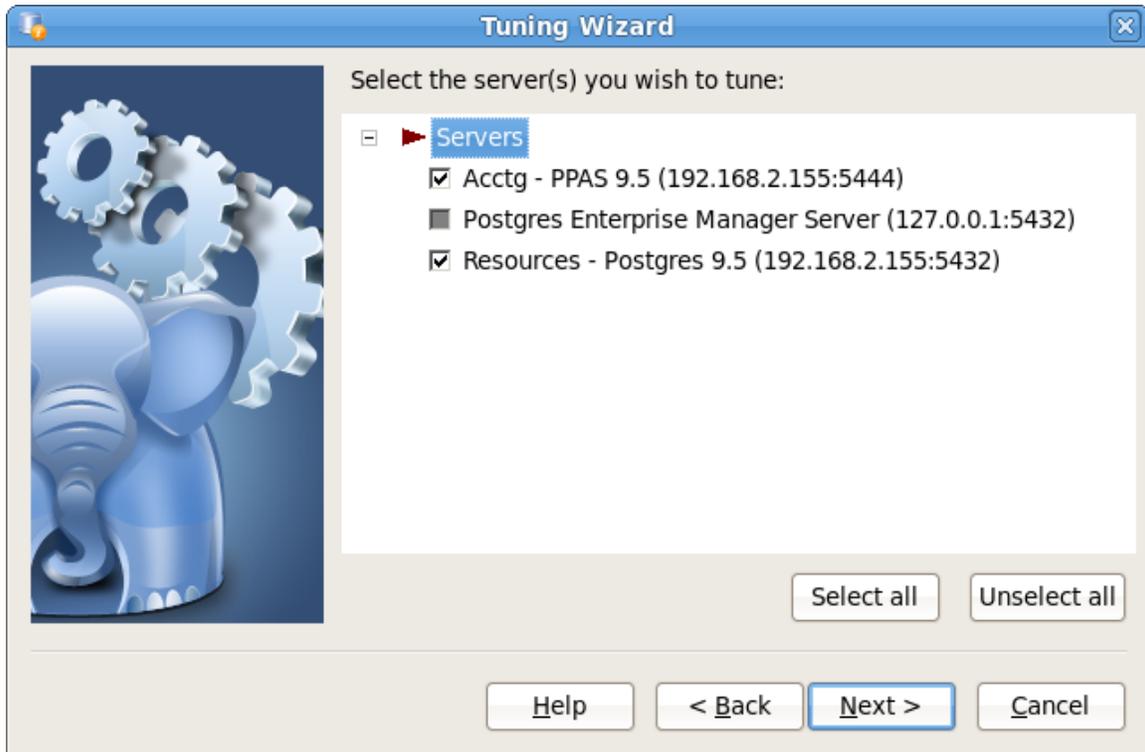


Figure 12.2 - The server selection dialog.

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. Use the `Select all` button to mark all servers for tuning, or the `Unselect all` button to clear all of the checkboxes.

Click `Next` to continue to the `Machine Utilization` dialog (shown in Figure 12.3).

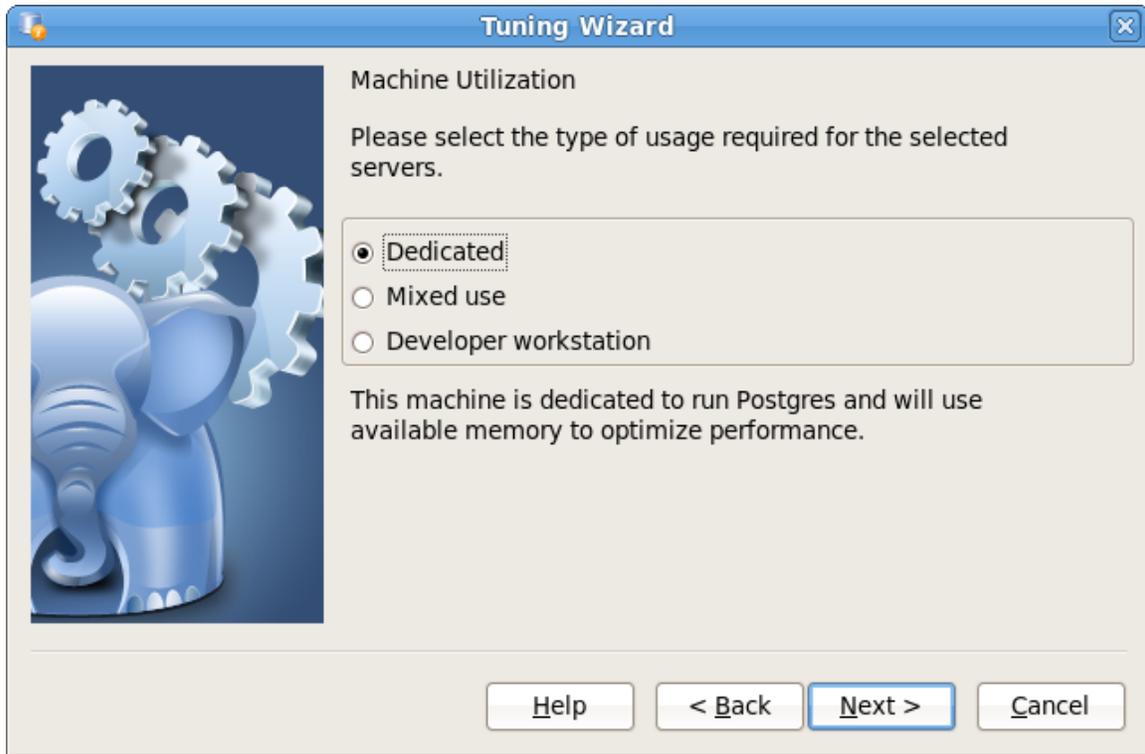


Figure 12.3 - The Machine Utilization dialog.

Use the radio buttons on the `Machine Utilization` dialog 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 the `Dedicated` radio button to dedicate the majority of the system resources to the database server.
- Select the `Mixed use` radio button to dedicate a moderate amount of system resources to the database server.
- Select the `Developer workstation` radio button to dedicate a relatively small amount of system resources to the database server.

Click `Next` to continue to the `Workload Selection` dialog (shown in Figure 12.4).

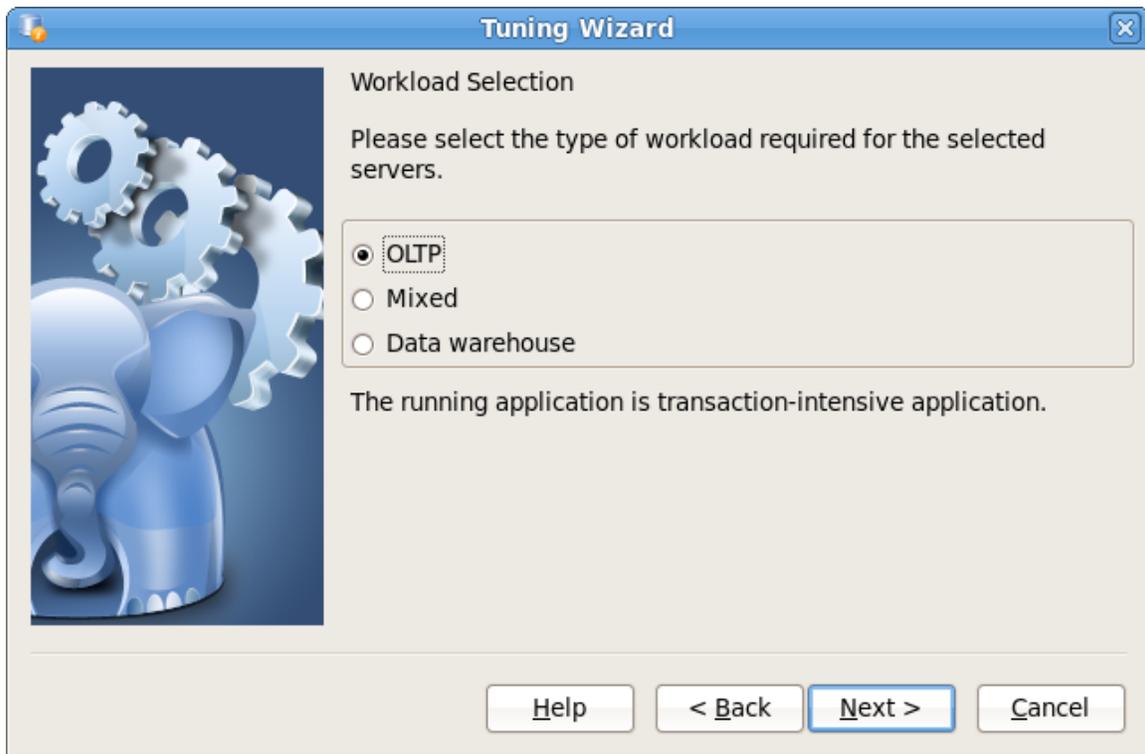


Figure 12.4 - The Tuning Wizard Workload Selection dialog.

Use the radio buttons on the `Workload Selection` dialog to specify the type of workload typically performed on the selected server:

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

Click `Next` to continue to the `Tuning Changes Summary` dialog (see Figure 12.5)

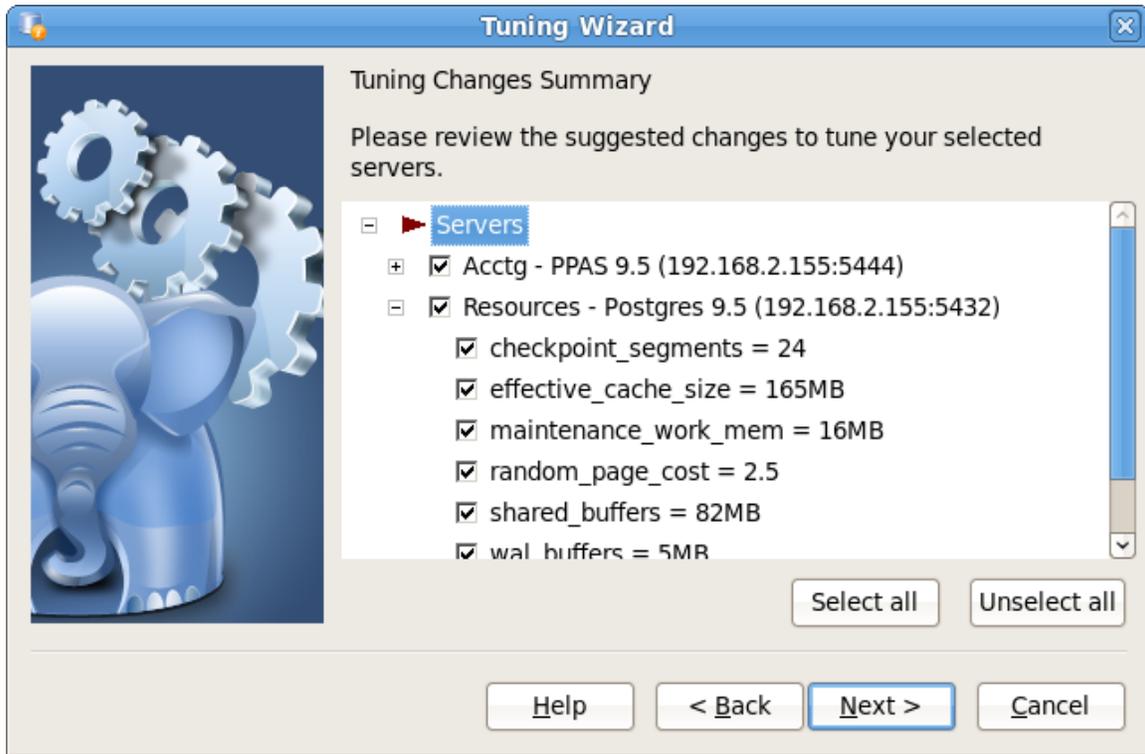


Figure 12.5 - The Tuning Changes Summary dialog.

The tree control on the Tuning Changes Summary dialog displays the servers analyzed by the Tuning Wizard, and the parameter setting modifications recommended for each server. 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.



Figure 12.6 - Schedule changes or generate a report.

Select the `Schedule changes` radio button (see Figure 12.6) to enable the buttons in the `Schedule selected tuning changes` box and 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 the `Configure tuning now` radio button to instruct PEM to apply the recommendations and restart the server immediately. This is the default.
- Select the `Schedule it for some other time` radio button to enable the calendar date and time selectors, and specify a time at which the changes will be applied and the server restarted. Once scheduled, the job can be viewed in the `Scheduled Tasks` dialog.

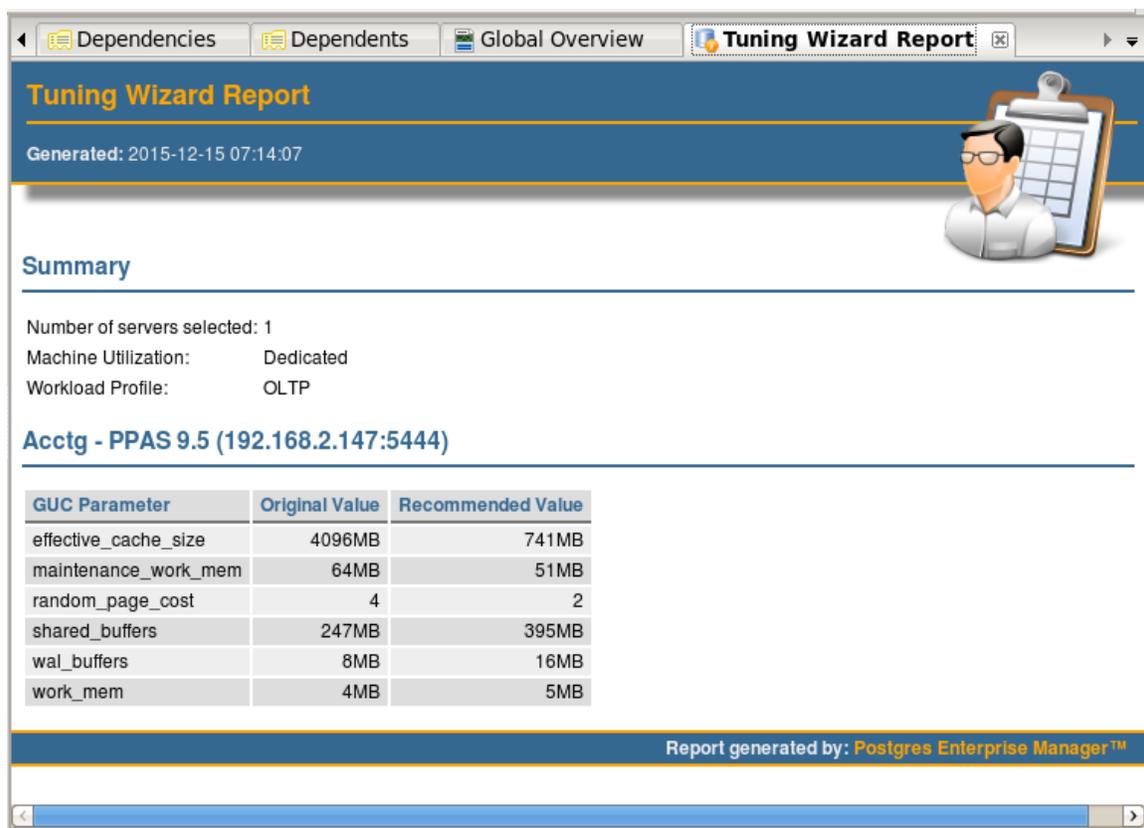
Select the `Generate report` radio button to enable the options in the `Generate tuning recommendation report` box and preview the changes recommended by the Tuning Wizard before applying any modifications. If you select `Generate report`, PEM will create a report that contains a list of the current 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.

Use the options in the **Generate tuning recommendation report** box to specify if PEM should display the report onscreen immediately, or write the report to a file:

- Select the radio button next to **View the report now?** to display the **Tuning Wizard Report** onscreen.
- Select the radio button next to **Save the report to a file:** and specify a file location to instruct PEM to write the report (in html format) to file for review.

Click the **Finish** button to either apply the **Tuning Wizard's** modifications or generate a report immediately (see Figure 12.7) and exit the **Tuning Wizard**.



Tuning Wizard Report
Generated: 2015-12-15 07:14:07

Summary

Number of servers selected: 1
Machine Utilization: Dedicated
Workload Profile: OLTP

Acctg - PPAS 9.5 (192.168.2.147:5444)

GUC Parameter	Original Value	Recommended Value
effective_cache_size	4096MB	741MB
maintenance_work_mem	64MB	51MB
random_page_cost	4	2
shared_buffers	247MB	395MB
wal_buffers	8MB	16MB
work_mem	4MB	5MB

Report generated by: **Postgres Enterprise Manager™**

Figure 12.7 - 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 12.8).

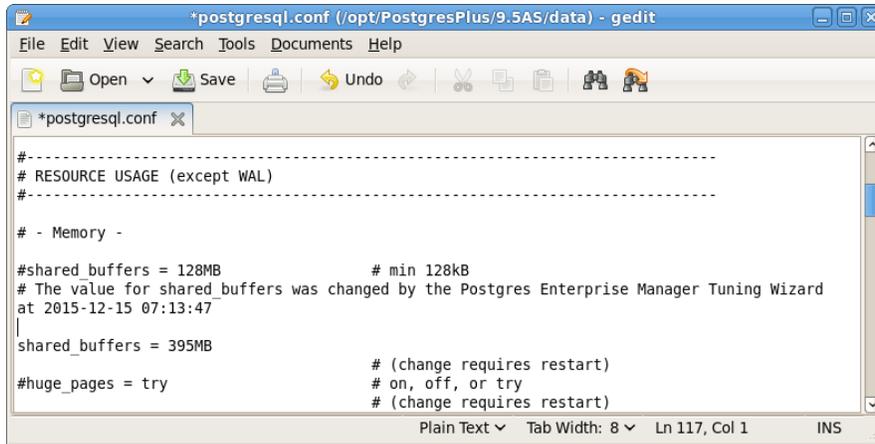


Figure 12.8 - A modified postgresql.conf file entry.

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 (see Figure 12.9) 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.

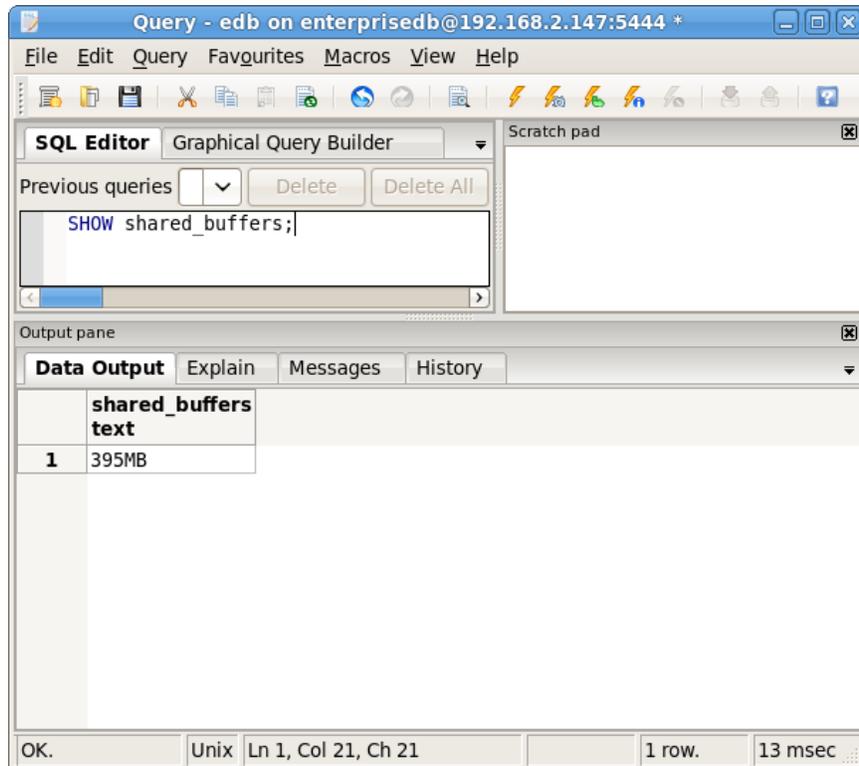


Figure 12.9 - The PEM Query Tool.

13 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.

13.1 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 13.1), click Next to continue.

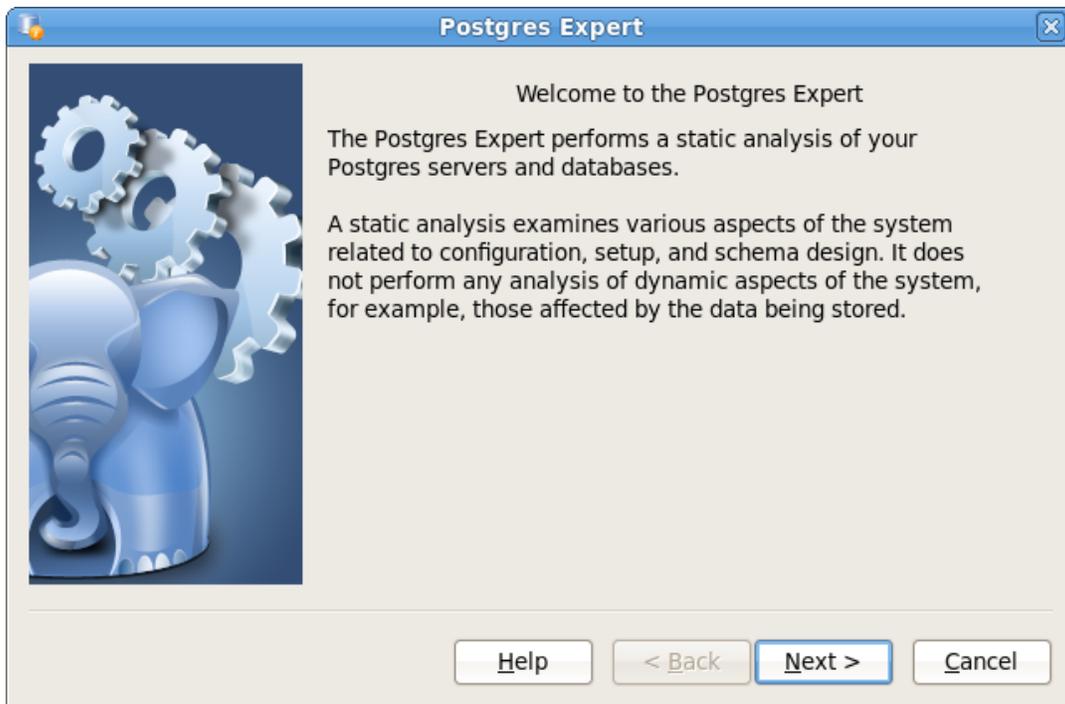


Figure 13.1 - The Postgres Expert Welcome dialog.

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 13.2).



Figure 13.2 - 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 selected item.

You can:

- Check the box next to the name of an expert to select all of the configuration items listed under that node of the tree control.
- Click the `Select all` button to instruct Postgres Expert to review the selected server for all of the items listed in the tree control.
- Click the `Unselect all` button to un-check all of the rules, and 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 13.3).

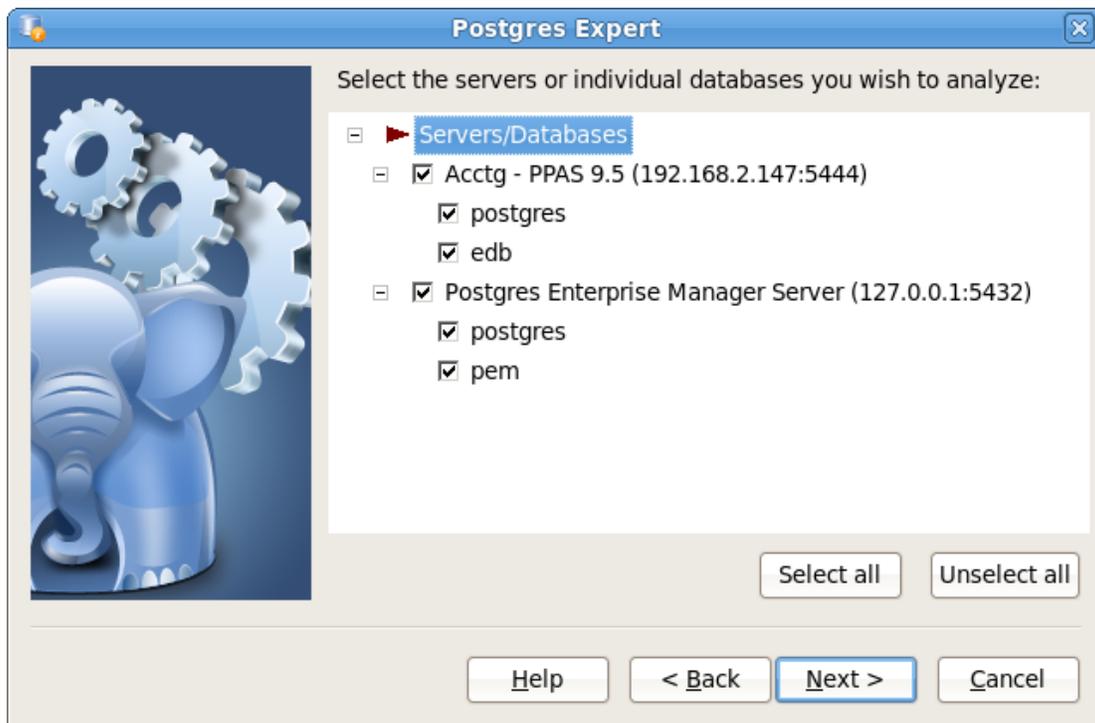


Figure 13.3 - The Servers/Databases dialog.

If you select multiple servers or databases, the resulting report will contain a separate analysis of each target. Select or de-select the servers and databases that you would like Postgres Expert to analyze, and select `Next` to select a report destination (see Figure 13.4).

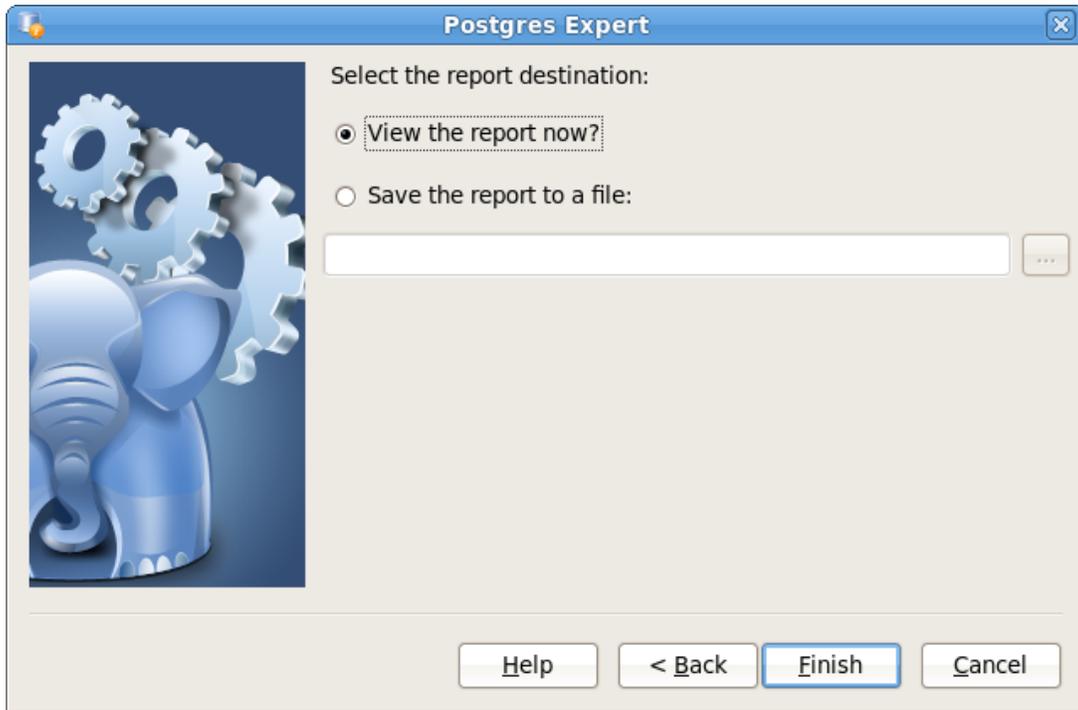


Figure 13.4 - Specify a report destination.

You can select the default option and click `Finish` to view an onscreen report from Postgres Expert, or check the box next to `Save the report to a file` to save a copy of the report to an HTML file for later use. If you wish to save the report to a file, specify a filename in the field provided, or use the button to the right of the field to open a browser window and select a location using the browser.

13.2 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 13.5).

The screenshot displays the 'Postgres Expert Report' interface. At the top, it shows the report title and generation date: 'Generated: 2015-12-15 08:09:45'. A 'Jump to:' dropdown menu is set to 'Postgres Enterprise'. Below this is a 'Summary:' section with a table of settings:

Settings	Value
Number of servers tested:	2
Number of rules checked:	31
Number of High alerts:	2
Number of Medium alerts:	5
Number of Low alerts:	4

The report is divided into sections for different servers. The first section is for 'Server: Postgres Enterprise Manager Server (127.0.0.1:5432)'. It features an 'Advisor: Configuration Expert' section with a table of rules:

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

Below this is an 'Advisor: Schema Expert' section with a table of rules:

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

The second section is for 'Server: Acctg - PPAS 9.5 (192.168.2.147:5444)'. It includes an 'Advisor: Configuration Expert' section with a table of rules:

Rule	Database	Severity
Check checkpoint_completion_target	-	Medium
Check effective_io_concurrency	-	Low

It also features an 'Advisor: Schema Expert' section with a table of rules:

Rule	Database	Severity
Check data and transaction log on same drive	-	High
Check for missing foreign key indexes	edb	Medium
Check for missing primary keys	edb	Low

Finally, there is an 'Advisor: Security Expert' section with a table of rules:

Rule	Database	Severity
Check SSL for improved connection security	-	Medium

The footer of the report states: 'Report generated by: Postgres Enterprise Manager™'.

Figure 13.5 - 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.

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`.

The screenshot shows the 'Postgres Expert Report' interface. At the top, it says 'Generated: 2015-12-15 08:09:45' and 'Jump to: Postgres Enterprise'. Below this is a 'Summary' section with a table of settings and values:

Settings	Value
Number of servers tested:	2
Number of rules checked:	31
Number of High alerts:	2
Number of Medium alerts:	5
Number of Low alerts:	4

Below the summary, it shows the server: 'Postgres Enterprise Manager Server (127.0.0.1:5432)'. The 'Advisor: Configuration Expert' section contains a table of rules:

Rule	Database	Severity
Check checkpoint_completion_target	-	Medium
Check effective_cache_size	-	Medium

The selected rule, 'Check effective_cache_size', has the following details:

- Trigger:** Current value is not equal to recommended value
- Recommended Value:** 741MB
- Description:** When estimating the cost of a nested loop with an inner index-scan, PostgreSQL uses this parameter to estimate the chances that rows from the inner relation which are fetched multiple times will still be in cache when the second fetch occurs. Changing this parameter does not allocate any memory, but an excessively small value may discourage the planner from using indexes which would in fact speed up the query.
- Current Values:**

Setting	Value
effective_cache_size	4096MB
total_ram_memory	988 MB

Figure 13.6 - 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 13.6). 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 the parameter(s).

14 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 (`pemagent.cfg`) located in `/opt/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 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 14.1.

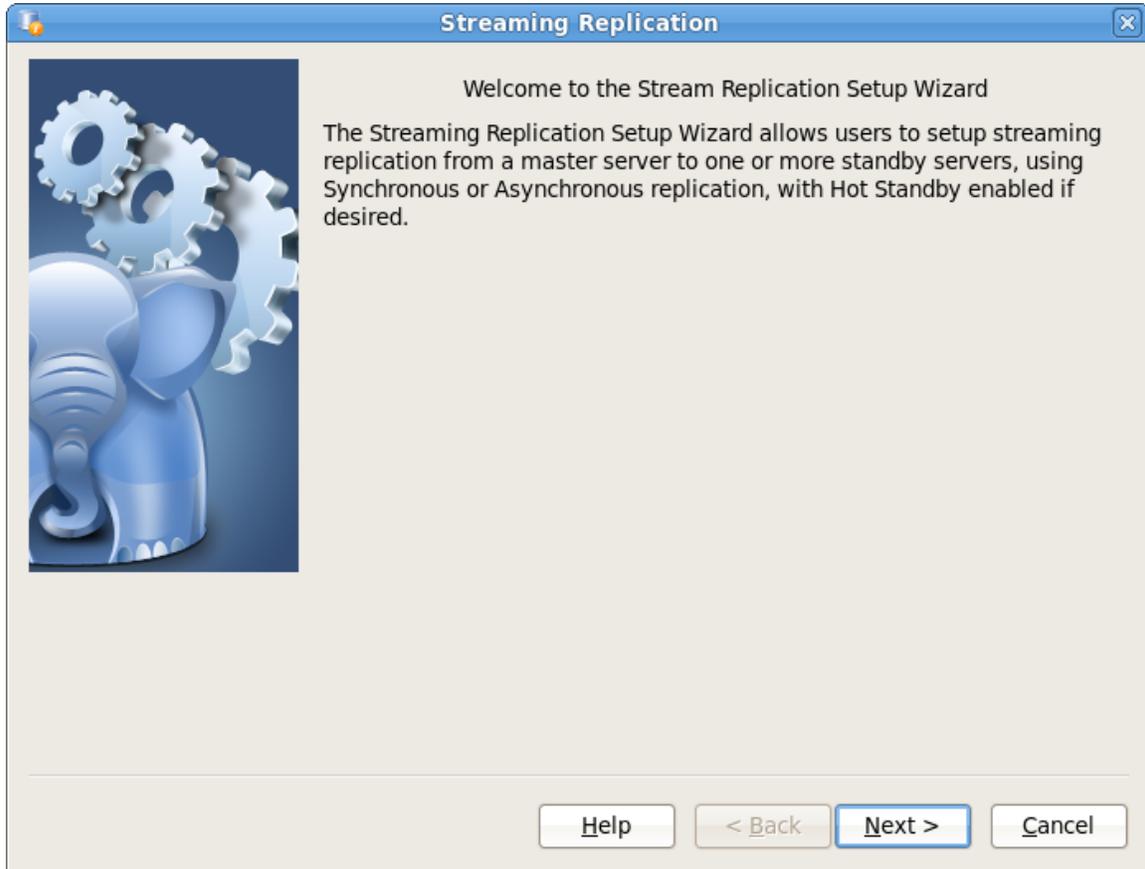


Figure 14.1— The Streaming Replication wizard's Welcome window.

Click `Next` to continue. A popup will open (see Figure 14.2), offering you the option to refresh the package data that is stored on the PEM server about the currently installed packages.

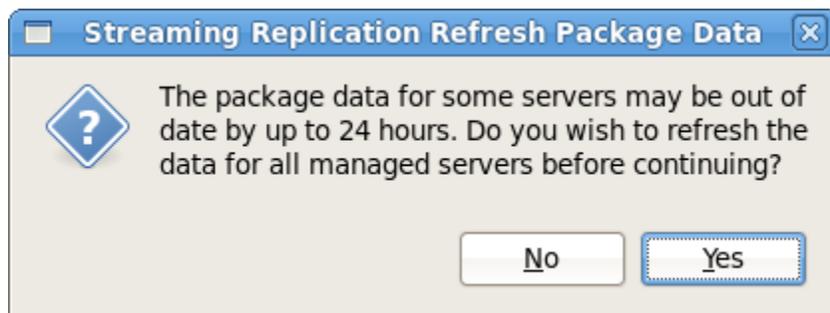
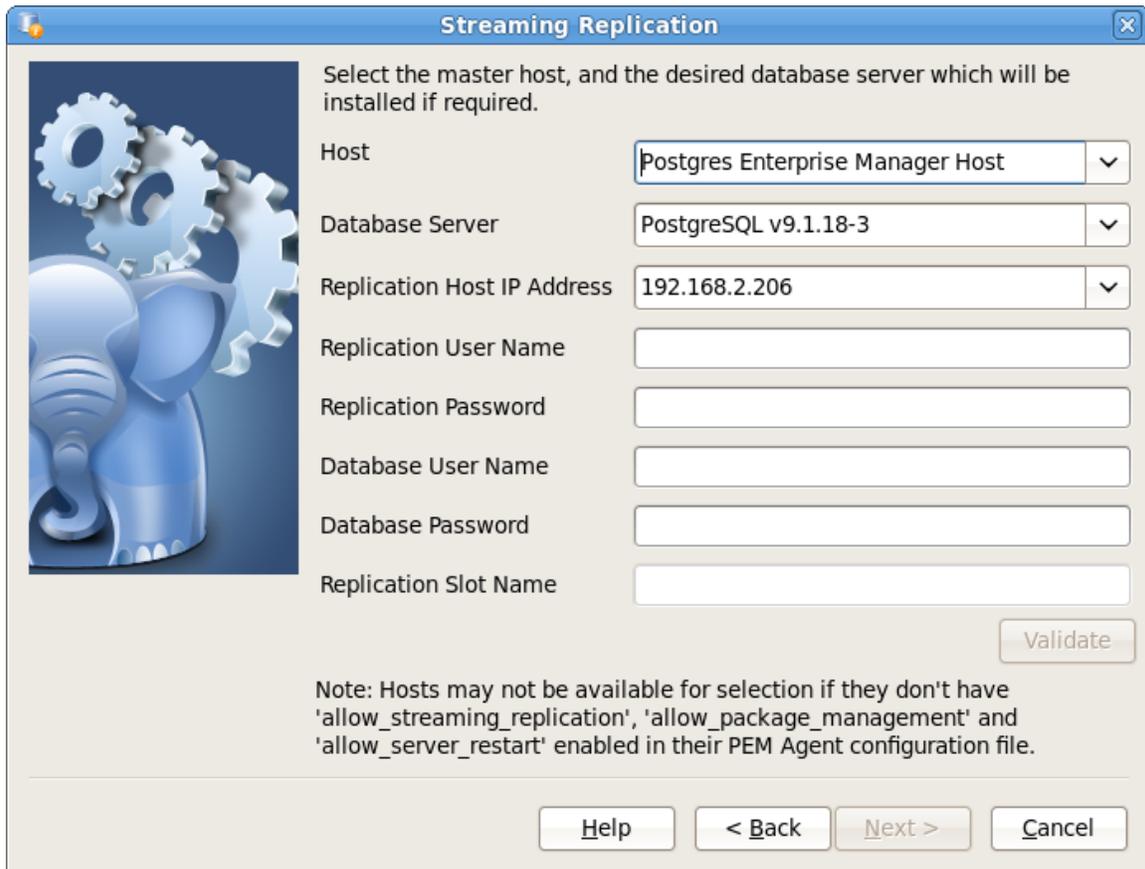


Figure 14.2 – Use the popup to refresh package data.

The PEM `installed_packages` probe retrieves information about the currently installed packages that reside on hosts that are monitored by PEM agents. Select **Yes** on the popup 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.



Streaming Replication

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

Host: Postgres Enterprise Manager Host

Database Server: PostgreSQL v9.1.18-3

Replication Host IP Address: 192.168.2.206

Replication User Name:

Replication Password:

Database User Name:

Database Password:

Replication Slot Name:

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

Figure 14.3 - Specify information about the master node.

Fields on the master node selection dialog (see Figure 14.3) 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 IP Address` field to select the address of the host on which the master node will reside.
- 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.
- Use the `Replication Slot Name` field to specify the name for a replication slot; the PEM server will create the replication slot on the master node during the replication setup process, 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.

Please note that replication slots are supported only on server versions 9.4 and above. For more information about replication slots, see the PostgreSQL Core documentation, available at:

<http://www.postgresql.org/docs/9.5/static/warm-standby.html#STREAMING-REPLICATION-SLOTS>

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. Click `Next` to continue.

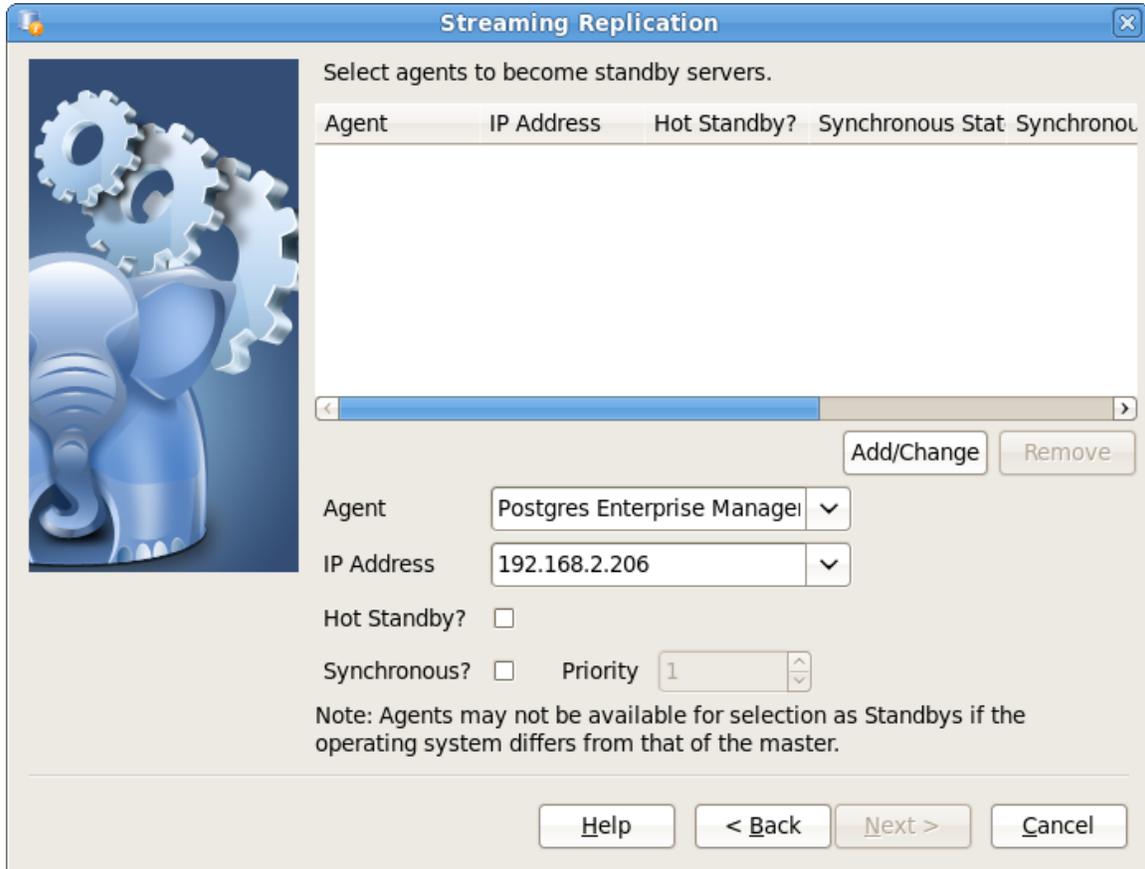


Figure 14.4 - Select the standby servers.

Fields on the standby selection dialog (see Figure 14.4) prompt you to provide properties of one or more standby nodes:

- Use the `Agent` drop-down listbox 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.
- Use the `IP Address` drop-down listbox to select the IP address of the standby node.
- Check the box next to `Hot Standby` if the standby node should be used for read-only queries while acting as a standby node in the replication scenario.

- Check the box next to `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 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.

Click the `Add/Change` button to add a standby node to the list of servers, or to modify the values associated with a server in the list. When you've defined the standby servers in the replication scenario, click `Next` to continue.

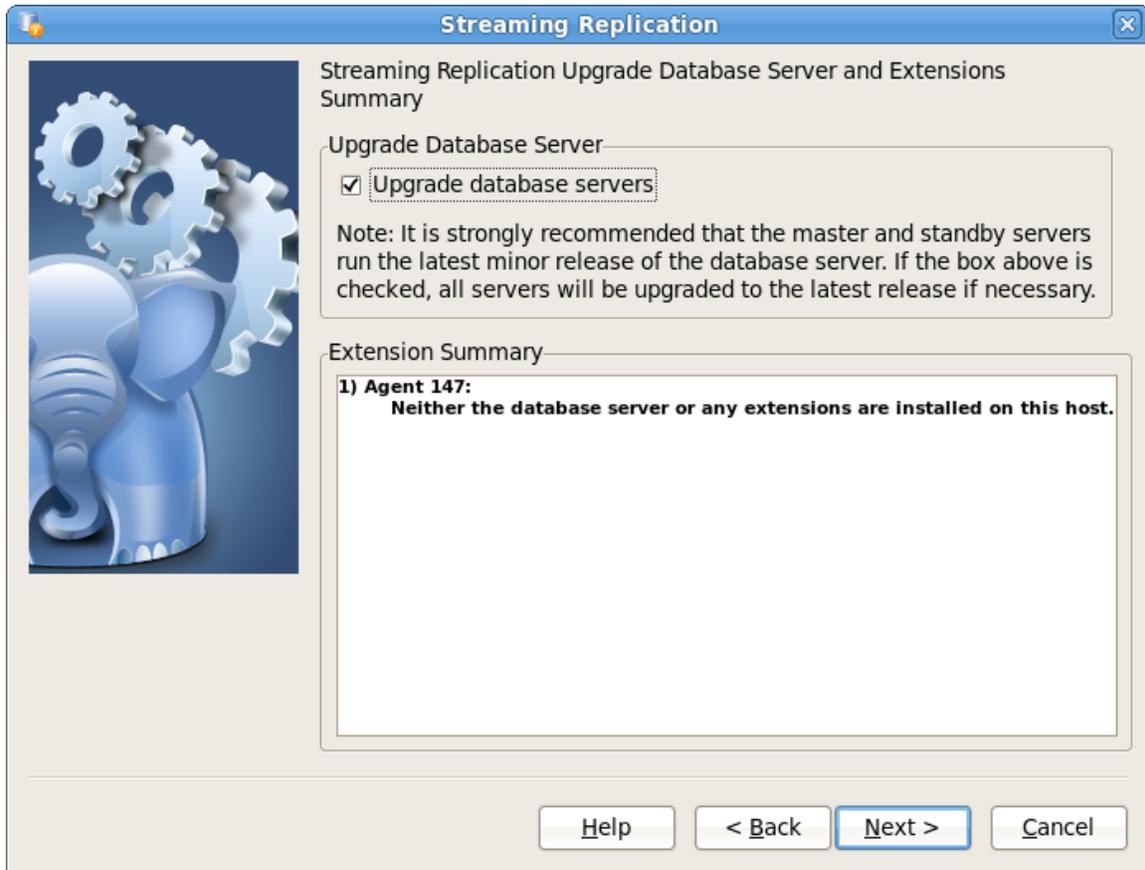


Figure 14.5 - The wizard's upgrade and extension window.

Check the box in the `Upgrade Database Server` panel (see Figure 14.5) to indicate that the server should be upgraded during the configuration process. Please note that the upgrade process will only upgrade to the most recent minor version of the selected server version available (for example, from version 9.3.0 to 9.3.9).

The `Extension Summary` panel 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.

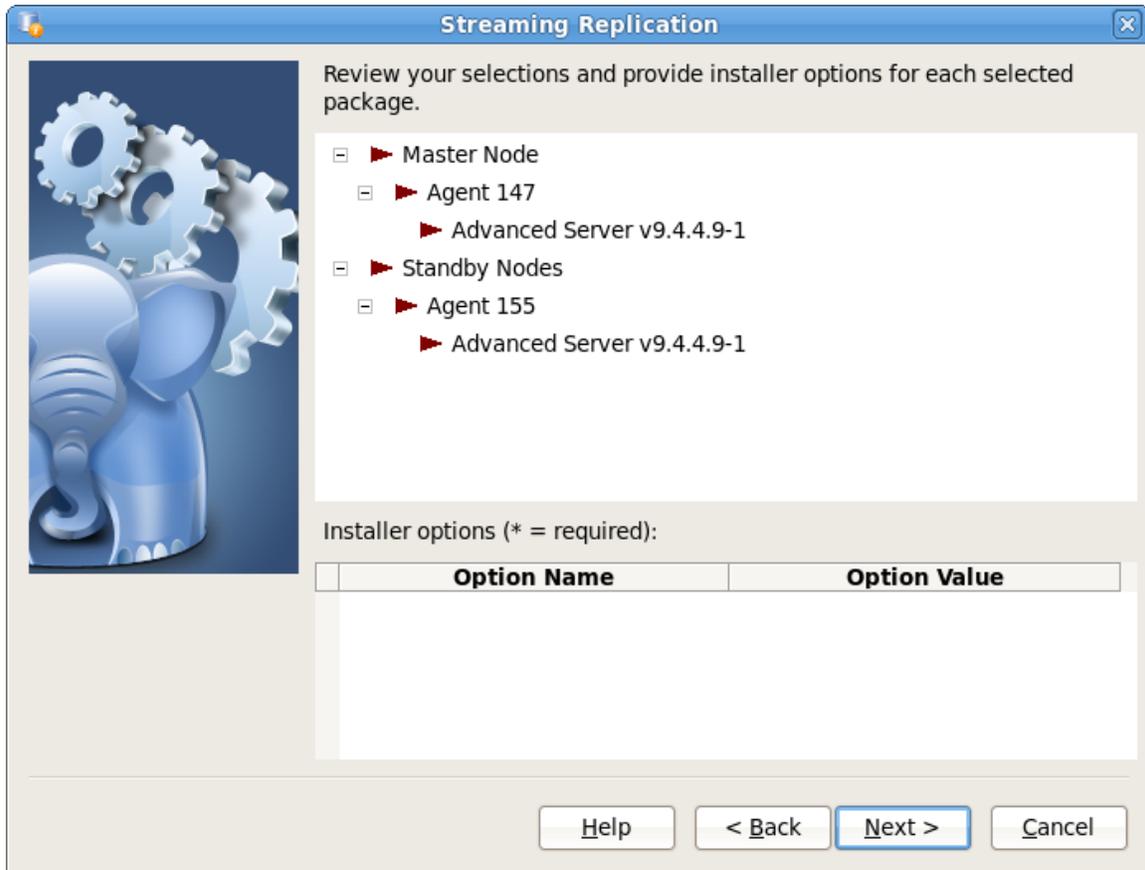


Figure 14.6 - Specify installation options for the master and standbys.

If PEM is installing new servers, the Streaming Replication wizard opens to a tree control (see Figure 14.6) that provides an overview of the master and standby nodes and allows you to specify installation properties for each server in the replication scenario. To review or modify the installation properties, highlight the name of a node in the tree control, and provide values in the Option value field. Please note that you must provide a value for any option marked with an asterisk (*).

After providing any required options, click `Next` to continue.

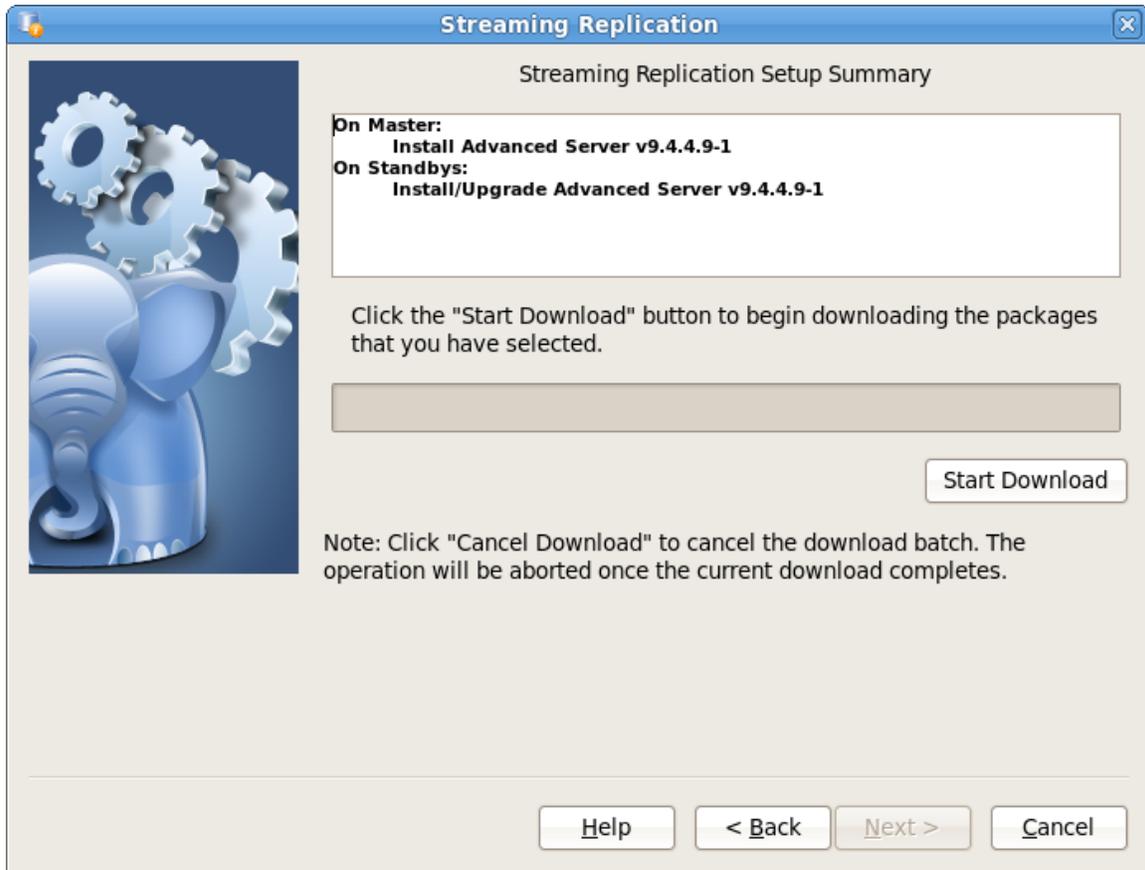


Figure 14.7 - The Streaming Replication wizard's setup summary window.

The Streaming Replication Setup Summary (see Figure 14.7) displays a list of the servers that will be part of 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.

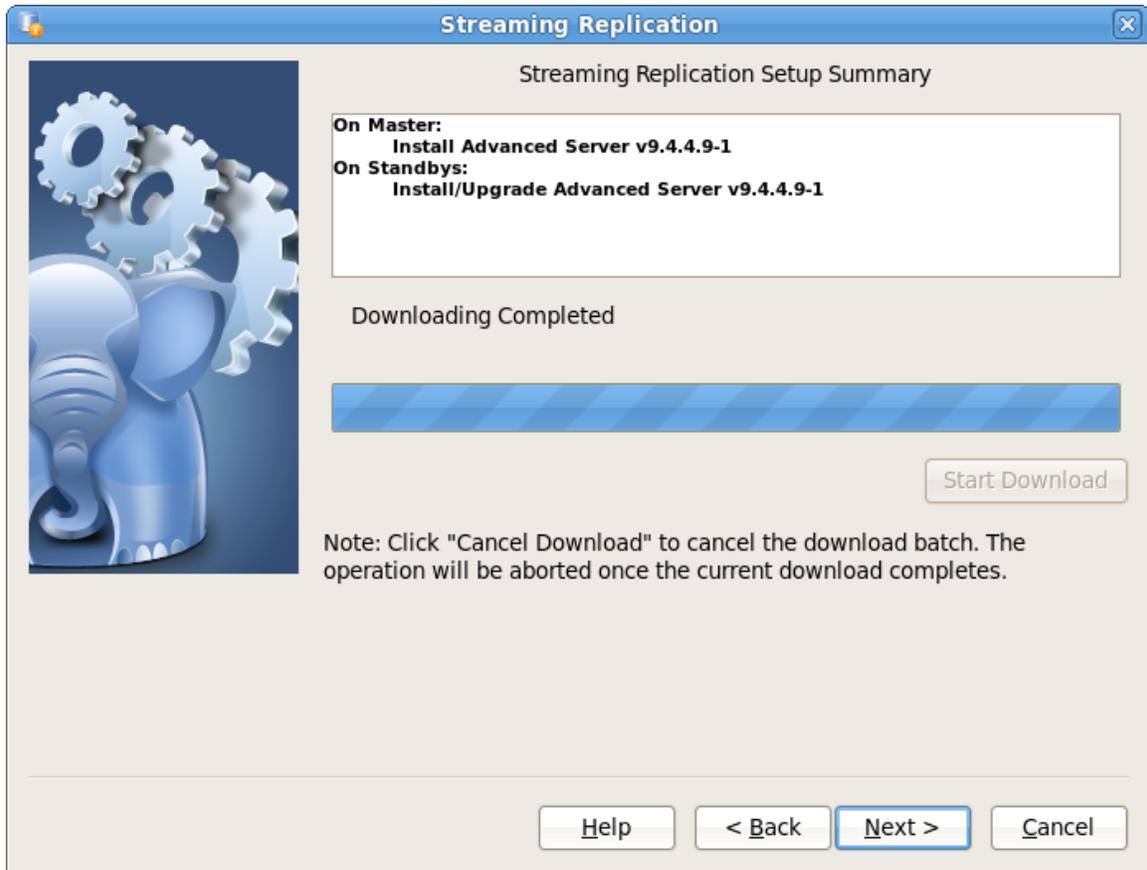


Figure 14.8 - The Installer download has completed.

When the download completes (see Figure 14.8), click `Next` to continue; the streaming replication wizard will open a dialog that allows you to schedule streaming replication setup.

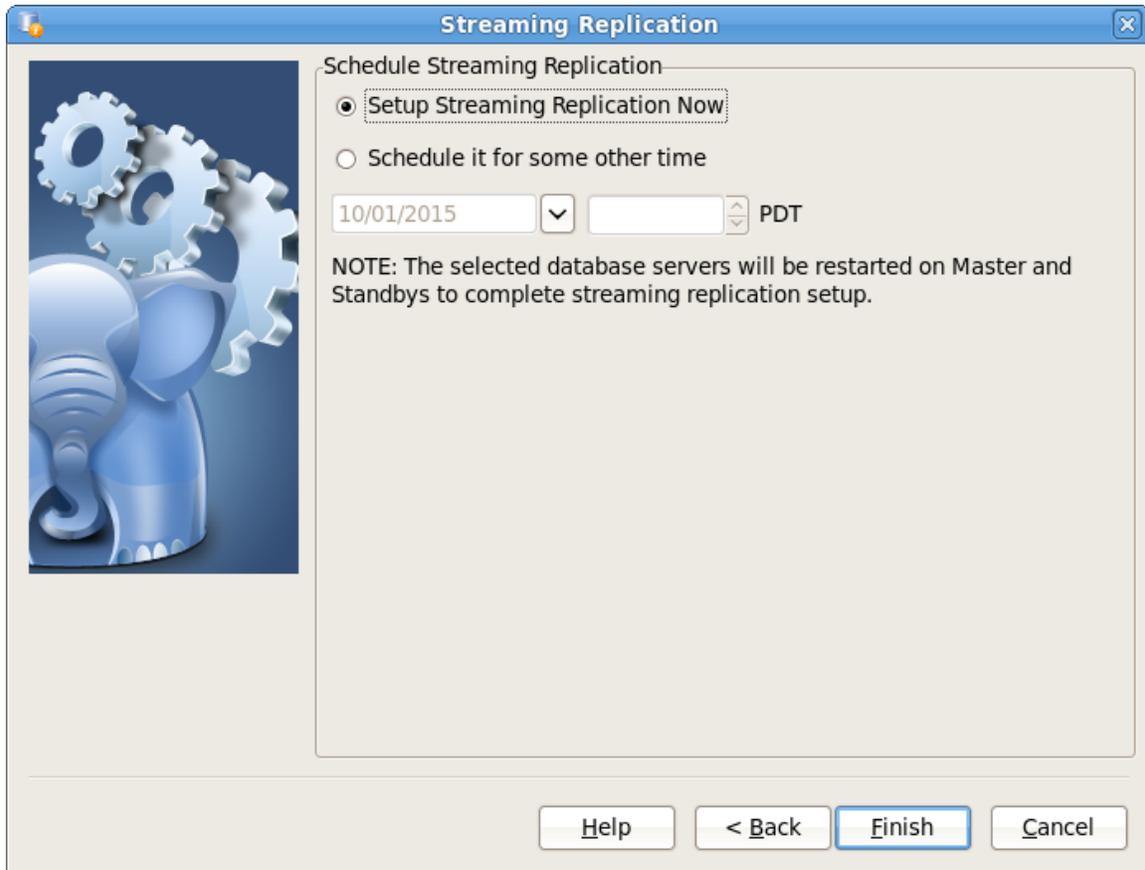


Figure 14.9 - Select a time to configure replication.

Use fields on the `Schedule Streaming Replication` dialog (see Figure 14.9) to specify the most convenient time for the server to configure the replication scenario:

- Click the radio button next to `Setup Streaming Replication Now` to instruct PEM that it should install and configure streaming replication immediately.
- Click the radio button next to `Schedule it for some other time` 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 in the Task Manager; to open the Task Manager dialog, highlight the name of the PEM agent for which you wish to review the job queue and select `Scheduled Tasks...` from the context menu.

When the installation and setup completes, you can define a server connection to the master or standby nodes in the PEM client, and monitor the new 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.

To view the Streaming Replication Analysis dashboard for the master node of a replication scenario, you must enable the following probes:

```
Streaming Replication
WAL Archive Status
```

To view the Streaming Replication Analysis dashboard for the slave node of a replication scenario, you must enable the following probes:

```
Streaming Replication Lag Time
```

Then, to open the Streaming Replication dashboard (see Figure 14.10), 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.

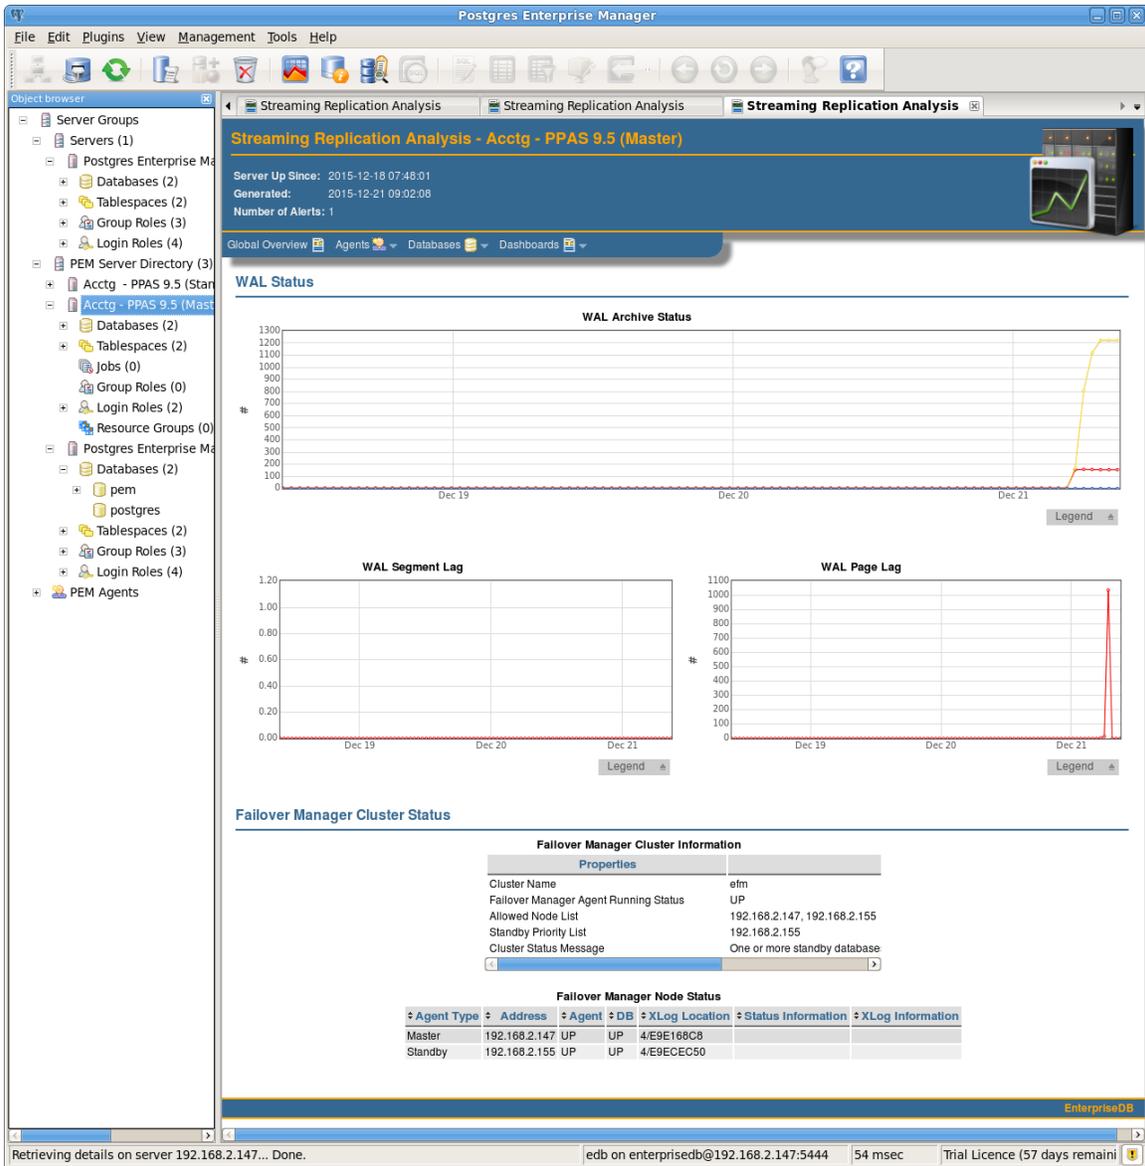


Figure 14.10 The Streaming Replication dashboard.

14.1 Monitoring Replication and Failover

The PEM client can display status information about one or more Failover Manager clusters on the Streaming Replication dashboard (see Figure 14.11).

Failover Manager Cluster Status

Failover Manager Cluster Information	
Properties	Values
Cluster Name	efm
Failover Manager Agent Running Status	UP
Allowed Node List	192.168.2.27, 192.168.2.93
Standby Priority List	192.168.2.93
Cluster Status Message	

Failover Manager Node Status						
Agent Type	Address	Agent	DB	XLog Location	Status Information	XLog Information
Master	192.168.2.27	UP	UP	0/C410D5C0		
Standby	192.168.2.93	UP	UP	0/C410D5C0		

Figure 14.11 - The Failover Manager cluster status report.

Before configuring PEM to monitor a Failover Manager cluster, you must install and configure Streaming Replication and Failover Manager on the cluster. For more information about installing and configuring Streaming Replication and Failover Manager, please see the EnterpriseDB Failover Manager Guide, available at:

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

To configure PEM to monitor a Failover Manager cluster, open the PEM client, log on to the PEM Server, and create a server definition for the master node of the Failover Manager cluster. In addition to using the tabs on the New Server Registration dialog to specify general connection properties for the master node, use fields on the Advanced tab to specify information about the Failover Manager cluster:

- Use the `EFM Cluster Name` 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 `EFM Installation Path` field to specify the location of the Failover Manager binary file. By default, the Failover Manager binary file is installed in `/usr/efm-2.0/bin`.

After saving the server definition, the master node will be included in the list of servers under the `PEM Server Directory` in the PEM client Object browser tree, and will be displayed on the Global Overview dashboard. To view the Streaming Replication Analysis dashboard and the status of the Failover Manager cluster,

right click on the name of the master node in the Object browser tree control and navigate through the Dashboards menu to select Streaming Replication Analysis.

14.2 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 start the failover process, select the `Replace Cluster Master` menu selection on the `Management` menu. When you click the `Replace Cluster Master` menu option, a dialog opens, asking you to confirm that you wish to replace the current master node (see Figure 14.12).

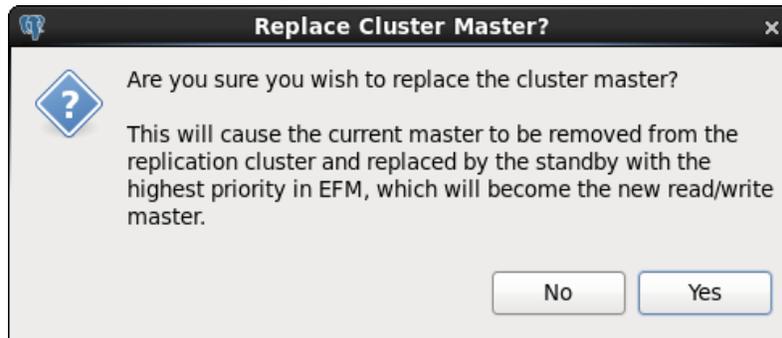


Figure 14.12 – Replacing the Master node of a cluster.

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 14.13).

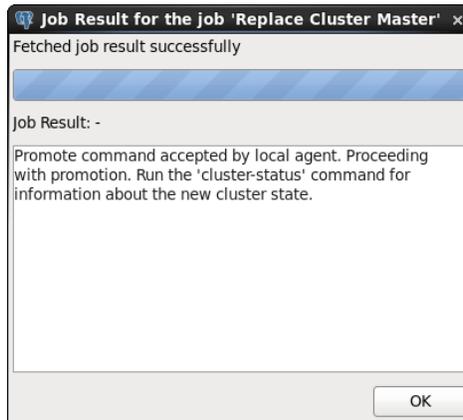


Figure 14.13 – Confirmation of the promotion.

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.

15 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.

15.1 About EnterpriseDB

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.

EnterpriseDB understands that adopting a new database is not a trivial task. You have questions that need answers, schedules and budgets to keep, and processes to follow. We have helped thousands of organizations like yours through the steps to investigate, evaluate, prove, develop, and deploy their PostgreSQL solutions.

To make your work easier and faster, we have special self-service sections on our website dedicated to assisting you in each of the steps. Additionally, visit

<http://www.enterprisedb.com/why-enterprisedb/postgres-plus-evaluation-production>

Getting Started – access to free downloads, installation guides, demos, starter tutorials, and more to help get familiar with the database.

Evaluations and Pilots – learn how EnterpriseDB has helped hundreds of Oracle users cut costs and MySQL users improve operations.

Development – EnterpriseDB employs more PostgreSQL experts, developers and community members than any other company, and offers key application development resources and services.

Deployment – information on how to scale a PostgreSQL application, add Qualities of Service (QoS) like high availability or security, or get a health check.

If you would like to discuss training, consulting, or enterprise support options, please contact EnterpriseDB directly. [EnterpriseDB](#) 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)

EnterpriseDB, Postgres Plus, Postgres Enterprise Manager, and DynaTune are trademarks of EnterpriseDB Corporation. Other names may be trademarks of their respective owners. © 2016.