

Administering the PMDB

This section includes the following topics:

- [About the PMDB](#)
- [About viewing and administering performance processes](#)
- [About viewing the status of the PMDB](#)

About the PMDB

The PMDB helps you identify resource consumption patterns and predict future resource consumption and response times. The following methods perform the identification and predictions:

- Track historical resource consumption trends so you can understand and predict long-term performance behavior.
- Perform period-to-period comparisons so you can analyze performance improvements or performance degradation over time.
- Track load patterns, entity changes, entity statistics, and component parameter changes so you can understand their effects on performance.
- Track data growth and data distribution changes so you can optimize data storage management.
- Proactively detect performance bottlenecks before they turn into problems and issue alerts because performance degrades from established baselines.

By using the PMDB to store this information, you can better manage your applications, make knowledgeable decisions about application changes and hardware upgrades, and improve plans for the future.

The PMDB uses the following methods to store and manage the PMDB data:

- Schedule and run batch processes that load historical data from each instance in each AppTier.
- Manage the summary procedure.
- Provide a common interface for requests coming from the Precise user interfaces.
- Manage and maintain the PMDB by filtering unnecessary data, saving baseline data, summarizing data, and purging old data.

You can control all PMDB settings from AdminPoint, in the following way:

- Edit your Warehouse password from the Warehouse tab. See [Editing your Warehouse password](#).
- Run, schedule, and change the parameter default values of the PMDB processes from the Processes view, under the Warehouse menu. See [Viewing the PMDB processes](#).
- Examine the status of instances and the PMDB processes and check in which PMDB tables instance data is stored from the Status view, under the Warehouse menu. See [About viewing the status of the PMDB](#).
- Manage hour groups from the Hour Group Settings dialog box. See [Configuring hour groups](#), [About the Dashboard tab](#), and [About the Applications view](#).

Editing your Warehouse password

 If your PMDB is SQL-based, see the instructions on how to change the database password manually in the Changing the PMDB SQL Server Collector Password technical note that can be found on the Precise Customer Portal.

To change your Warehouse password

1. In the Warehouse Processes tab in AdminPoint, click **Edit Warehouse Password**. The Edit Warehouse Password dialog box appears.
2. Select whether to update the Precise registry only (set as the default option), or to update both the Precise registry and the warehouse database.
3. As indicated, enter your current password (when enabled), your new password, and confirm your new password.
4. Click **OK**. A confirmation window appears.

 After editing your Warehouse password, there is an automatic restart to FocalPoints.

About viewing and administering performance processes

The Processes view shows information about batch processes performance on the Warehouse work station. The Processes table in the Main area includes the following columns:

Column	Description
Node	The specific Precise framework installation.  A Precise system can have one or more framework deployments.

Technology	The name of the technology that the process gathers data for. Processes that maintain the PMDB itself are listed, but the technology column remains empty.
Process	The process name is a logical name that describes the process function. For example, "Calculate Baselines." Two types of processes gather PMDB data. One process type includes processes that are part of the PMDB FocalPoint and run on the PMDB server, such as the Summarize Data process. The other process type includes processes that are part of another product, such as the Schema Changes process in a Microsoft SQL Server instance that Precise for SQL Server monitors. These processes run on the product's server and on the PMDB FocalPoint, which is in charge of scheduling them. See Batch processes in Precise .
Product	The name of the product the process runs for.
Last Run	The last time the process that is named in this row ran.
Next Scheduled Run	Time when the process that is named in this row is scheduled to run again.
Status	An icon indicating the current status of the process that is named in this row.

The current status of a process can be any of the following statuses:

Icon	Description
	The process is running.
	The process completed successfully.
	The process failed during the run.

Under the Processes view (on the Warehouse workstation), you can run a process, view the execution history of a process, and configure the scheduling and the parameters of a process.

Manually running a process

All processes are scheduled to automatically run at the set interval, but you may also manually run a process whenever required. See [Process parameters](#).

To run a process manually

- In the Process table, select a process and click **Run**.

The process starts and the page refreshes to show the updated status icon.

Viewing the run history of a process

You can examine additional information about the past executions of a process, for example to find out which tables were affected during a specific execution or which component has failed during a run.

To view the run history

1. In the Process table, select a process and click **Run History**.
2. In the Run History dialog box, in the top table, select the execution that you want to focus on, such as the last execution that failed (indicated by the status icon). The bottom table populates with detailed information on the performance processes tables that the execution affected.
3. Investigate all relevant data and click **Close**. For example, if an error occurred, focus on the details in the Error Message column.

Process parameters

Specific parameters define the execution behavior of performance processes. These parameters function as a filter in that they determine exactly which information is collected and saved in the Warehouse workstation.

You can manually configure these parameters according to the requirements of your environment. You can edit the default values or override them by defining instance-specific values. If a process does not include any configurable parameters, the Parameters button is disabled. See [About configurable process parameters](#).

The following table lists all the processes that you can configure per monitored instance.

Table 5-1 Processes configurable per instance

Technology	Process
DB2	Explain Statement
DB2	Explain New Statements
DB2	Load Data
DB2	Purge Internal Data
J2EE	Load Data
Oracle	Load Data
Oracle	Purge Data
Oracle	Explain Statements
Oracle (SmarTune)	Changes
SQL Server	Load Data You can restrict the number of sessions that performance data is collected for by increasing the In MS-SQL Time value. However, this value applies to the collapsed statement, including all of its children. For example, if you set the In MS-SQL Time value to 2 and the collapsed statement includes five children that spent 1 second each in MS-SQL, this amounts to 5 seconds overall and causes Precise for SQL Server to collect the performance data for the collapsed statement, even though each child spent less than the defined 2 seconds in MS-SQL.
SQL Server	Explain Statements
SQL Server	Explain New Statements
SQL Server	Collect Objects Performance
SQL Server	Collect Operational Statistics
SQL Server	Collect Space Usage
SQL Server	Collect Schema Changes
SQL Server	Purge Internal Data
SQL Server	Perform SmarTune Analysis
Sybase	Explain Statements
Sybase	Collect Space Utilization
Sybase	Explain New Statements
Sybase	Load Data
SAP	Load Data
Microsoft .NET	Load Data
Microsoft .NET (SmarTune)	Perform SmarTune Analysis
Operating System	Load Data
Other	Load Data
Tuxedo	Load Data
WebSphere MQ	Load Data
PMDB	Purge Data
PMDB	Summarize Data
SmartLink	Load Data
Web	Load Data

The following table lists the processes that you cannot configure.

Table 5-2 Processes that cannot be configured

Technology	Process
DB2	Calculate Objects Usage
Oracle	Collect Instance Definitions
Oracle	Collect Instance Statistics
Oracle	Collect Schema Changes
Oracle Storage	Load Data
Oracle	Perform SmarTune Analysis on DB Behavior
SQL Server	Collect Instance Statistics
SQL Server	Collect DB File Statistics
SQL Server	Collect Job Statistics
SQL Server (Storage)	Load Data
Sybase	Collect Instance Statistics
Sybase	Purge Internal Data
Sybase Replication Server	Load Data
SAP	Organizational Mapping
PMDB	Data Maintenance (Daily)
.NET	Load Data
	Summarize Current Data

To configure process parameters

1. In the Main area table, select the process whose parameters you want to configure.
2. Click **Parameters**.
3. In the Parameters dialog box, perform one of the following actions:

To edit the default values for the selected technology	Click Edit Default Values. In the Default Values dialog box, specify the values and click OK. Unless instance-specific values override these values, they apply to all monitored instances.
To edit the values for a specific instance	Select an instance in the table. Click Edit. In the Edit Instance dialog box, clear the Use Defaults check box. Specify the values to be used for this instance. Click OK.

4. In the Parameters dialog box, click **OK**.

 After you change the default of the Minimum In Oracle time to load parameter in the Load Data process, you must stop and then restart the Precise for Oracle FocalPoint. The next step is to restart the Collector agents.

 Changing the "Load Data" value from "300" to "30" for Precise for J2EE can take up to ten minutes to take effect.

 You can also configure the Precise for Web URL mapping rule for a domain in a system: "Treat pages with different titles as different pages." Select to view the URL along with the various titles that belong to it. The collected data includes page URLs and page titles. When cleared, you can only differentiate between different URLs, not between page titles. In this case, the collected data includes the page URL and the first page title for any URL in the system.

Scheduling processes

When the default scheduling settings do not match the requirements of your environment, you can manually schedule processes. Processes on the Warehouse workstation let you enable or disable each process. You can also schedule each process individually and run it at any frequency and interval - every 8 hours, hourly, daily, weekly, or by slice.

In environments with several instances, each instance creates a different amount of data and has its own frequency of activity. Therefore, it is sometimes necessary to schedule each instance according to the specific instance's behavior. Processes on the Warehouse workstation let you schedule some processes per instance.

The following table provides information on the default execution intervals of all Precise batch processes.

Table 5-3 Default intervals of Precise processes

Default Interval	Technology	Process
Weekly	PMDB	Purge Data
Weekly	PMDB	Warehouse Maintenance (Weekly)
Weekly	PMDB	Calculate Baselines
Weekly	Oracle	Collect Instance Definitions
Weekly	Oracle	Collect Schema Changes
Weekly	Oracle	Purge Data
Weekly	Oracle	Identity Changes (SmarTune)
Weekly	SQL Server	Collect Schema Changes
Weekly	SQL Server	Purge Internal Data
Weekly	Sybase	Purge Internal Data
Daily	Oracle	Explain Statements
Daily	Oracle	Collect Objects Statistics Changes
Daily	SQL Server	Explain Statements
Daily	SQL Server	Collect Space Utilization
Daily	Sybase	Explain Statements
Daily	Sybase	Collect Space Utilization
Daily	SAP	Organizational Mapping
Daily	PMDB	Warehouse Maintenance (Daily)
Hourly	J2EE	Load Data
Hourly	Oracle Storage	Load Data
Every 8 hours	PMDB	Summarize Data
Hourly	Oracle	Collect Instance Statistics
Hourly	Oracle	Analyze Database Behavior (SmarTune)
Hourly	SQL Server	Collect DB File Statistics
Slice	Oracle	Load Data
Slice	Web	Load Data
Slice	SQL Server	Load Data
Slice	SQL Server	Collect Instance Statistics
Slice	SQL Server	Explain New Statements
Slice	Sybase	Load Data
Slice	Sybase	Collect Instance Statistics
Slice	Sybase	Explain New Statements
Slice	Sybase Replication Server	Load Data

Slice	SAP	Load Data
Slice	WebSphere MQ	Load Data
Slice	Tuxedo	Load Data
Slice	Oracle Applications	Load Data
Slice	Network	Load Data
Slice	Operating System	Load Data
Slice	Insight SmartLink	Load Data

The following table lists processes that cannot be scheduled per instance.

Table 5-4 Precise processes not scheduled per instance

Technology	Process
PMDB	Calculate Baselines
SAP	Organizational Mapping
WebSphere MQ	Load Data
Tuxedo	Load Data
Oracle Applications	Load Data
Other	Load Data

To schedule a process

1. Select a process in the Main area table and click **Scheduling**.
2. In the Scheduling dialog box, perform any of the following tasks:
 - Clear the **Run this process in the PMDB** check box to disable the process. This action removes the process from the PMDB Scheduler, but you can still run the process manually. See [Manually running a process](#).
 - Change the time that the process is scheduled to run. Each process runs at a predefined interval by default: Every 15 minutes (slice), hourly, daily, or weekly. You cannot change this interval, but you can set the exact time that the process runs, for example every 2 hours, 15 minutes after the hour.
 - Schedule the process per instance (only available for some processes). This mechanism lets you perform the following tasks:
 - Disable a process for a specific instance by clearing the relevant check box in the first table column.
 - Set the interval of an instance within the process. For example, if you set the Run Every... column to **2**, the instance executes every 2 minutes (or other) run of the process.
3. Click **OK**.

Batch processes in Precise

The following sections provide a description of all existing Precise batch processes. The following table lists all batch processes on the Warehouse workstation.

Table 5-5 Batch processes on the Warehouse workstation

Process	Description
Summarize Data	Aggregates the loaded data into daily, weekly, and monthly summary tables.
Purge Data	Deletes old and unnecessary data from the statistics tables in the Warehouse workstation.  Although you can set the purge settings per instance, it is recommended to use the same settings for all instances in the same AppTier. Otherwise, a 'No Data Found' message appears when you launch from one product to another, and inconsistent information may be displayed in cross-AppTier products like Insight and Report Manager.
Warehouse Maintenance	Maintains the Warehouse workstation database. This process deletes old entries from the job history table and analyzes statistics and summary tables to update the database statistics. It also rebuilds and defrags indexes and maintains the normalization tables.

Calculate Baselines	Tracks the exceptions of currently measured counters that are compared to their historical values.
Load Data	Loads the data that is collected by a product into the database tables. This information lets you understand the behavior of your applications by examining their performance over time, identifying problematic time periods, and enabling you to drill down to the cause of the performance problem. Each product has its own Load Data process.
Organizational Mappings	A daily process that maps each organization, as it appear in the SAP system, to the physical site where its employees are located. Each organization can be mapped to more than one location. This mapping data is used to show SAP workload statistics according to locations.

The following table lists all Precise for Oracle batch processes.

Table 5-6 Precise for Oracle batch processes

Process	Description
Explain Statements	The PMDB can explain SQL statements that are loaded from the Precise for Oracle short-term history. This process explains new statements that have not yet been explained, or statements that have changed one of their dependent jobs. An object is considered changed either if its definition changes (for example when an index is created) or when its statistical data changes (for example the number of 'dirty' blocks in a table that is increased after running the Oracle Analysis command).
Collect Instance Definitions	Collects the changes in the instance definition, such as changes of the init.ora file or database parameters.
Collect Instance Statistics	Collects data on activity statistics of the Oracle server.
Collect Schema Changes	Collects the schema definitions and statistics from the Oracle database, compares them to the definitions that are collected in the previous execution of the process, and stores changes in definitions.
Collect Objects Statistics Changes	Collects the statistics changes for tables, indexes, and columns (identified based on DBMS_STATS execution). Also collects space utilization information for tablespaces.
Identify Changes (SmarTune)	Analyzes the changes to schema or instance settings to identify changes that affected the resource consumption or the response time (or both), either on the application level or even on the instance level.
Analyze Database Behavior (SmarTune)	Analyzes the instance behavior to detect places where your instance spends more time than expected.
Purge Data	Deletes old and unnecessary data that is related to instance definitions and schema changes.

The following table lists all Precise for SQL Server batch processes.

Table 5-7 Precise for SQL Server batch processes

Process	Description
Explain statements	Generates the execution plans for the statements that are collected by the Precise for SQL Server Collector agent and saves the execution plan information in the Warehouse workstation to enable you to tune these statements. This way, you can understand the execution plan that SQL Server chose for the statement, identify problematic operations in the execution plan, and analyze the objects that is referenced in this statement.
Collect Instance Statistics	Loads the Instance Statistics counters that are sampled by the Precise for SQL Server Collector agent into the Warehouse workstation. These performance counters enable you to analyze various aspects of the performance of the SQL Server instance and the overall system (such as CPU, paging, I/O, Network) over time.
Collect DB File Statistics	Samples SQL Server and collects I/O statistics for database files and saves the information that is sampled in the Warehouse workstation. This information lets you analyze the I/O activity of the database files over time and discover which files show heavy I/O activity so that you can make a sound judgment as to how these files should be dispersed amongst the various disks.
Collect Space Utilization	Samples SQL Server to collect information on the space that is allocated and used by the tables, indexes, and database files and saves this information in the Warehouse workstation. This information lets you easily keep track of object growth over time and decide how to disperse database files over various disks and how to relate objects to FILEGROUPS in the specified database.

Collect Schema Changes	Identifies the instance and database configuration changes, schema changes, and job changes that are performed in the SQL Server instance and saves the changes in the Warehouse workstation. This information lets you track and examine the changes that are performed on the instance and its entities and to understand how the changes affected the application's performance.
Purge Internal Data	Purges data as part of maintenance, such as purging an Availability log file.
Explain New Statements	Similar to Explain Statements, but it is executed every 15 minutes and generates execution plans for a few of the most resource-consuming statements. This way, the execution plan information becomes available as soon as possible, and not only after the Explain Statements have been carried out (which would make the Execution plan information only available after a one-day delay).
Collect Job Statistics	Samples SQL Server job runs and saves this information in the Warehouse workstation. This information lets you examine job runs over time, identify problematic time periods, and view the general trends in the job runs.

The following table lists all Precise for Sybase batch processes.

Table 5-8 Precise for Sybase batch processes

Process	Description
Explain statements	Generates the execution plans for the statements that are collected by the Precise for Sybase Collector agent and saves the execution plan information in the PMDB to enable you to tune these statements. This way, you can understand the execution plan that Sybase chose for the statement, identify problematic operations in the execution plan, and analyze the objects that are referenced in this statement.
Collect Instance Statistics	Loads the Instance Statistics counters that are sampled by Sybase Adaptive Server Enterprise (ASE) into the PMDB. These performance counters enable you to analyze various aspects of the performance of the Sybase instance (such as CPU, paging, I/O, Network) over time.
Collect Space Utilization	Samples Sybase to collect information on the space that is allocated and used by the tables, indexes, segments, and database devices and saves this information in the PMDB. This information lets you easily keep track of object growth over time and decide how to disperse database files over various disks and how to relate objects to devices in the specified database.
Purge Internal Data	Purges internal data as part of maintenance.
Explain New Statements	Similar to Explain Statements, but it is executed every 15 minutes and generates execution plans for a few of the most resource-consuming statements. This way, the execution plan information becomes available as soon as possible, and not only after the Explain Statements have been carried out (which would make the Execution plan information only available after a one-day delay).
Sybase Replication Server: Load Data	Loads the outbound processing Statistics counters that are sampled by the Sybase Replication Server Collector agent into the PMDB. These counters enable you to analyze the outbound processing phase of the replication system, and locate the specific component(s) which may be causing the latency in the replication system.

The following table lists Precise SmartLink processes

Table 5-9 Precise SmartLink processes

Process	Description
Load Data	Loads the data that is collected by a product into the database tables. This information lets you understand the behavior of your applications by examining their performance over time, identifying problematic time periods, and enabling you to drill down to the cause of the performance problem.  The Load Performance parameter for Precise SmartLink does not have a "Technology" value in the Display Warehouse Processes Information table in AdminPoint>Warehouse. To change the Load Data slice size, see Modifying the Precise Load Data slice size.

Modifying the Precise Load Data slice size

Most process are specific to one technology. The Load Data process that does not have a Technology value is shared by SmartLink and J2EE. All changes to this value affect the slice size not only in SmartLink tabs, but also in all J2EE instances in that environment.

To modify the Precise SmartLink/J2EE Load Data Performance slice size

1. Go to **AdminPoint>Warehouse>Processes**.
2. From the table, select the Load Data process with an empty Technology value. Click **Parameters**.
3. Select the relevant environment and click **Edit**.
4. Mark the **Use Defaults** check box to enable editing.

5. From the drop-down menu, select the desired slice size.
6. Click **OK>OK**.



Changing the **Load Data** value from "300" to "30" may take up to 10 minutes to take affect in Precise for J2EE.

About the PMDB Summarize Data process parameters

The table below describes the configurable process parameters and default values in the PMDB Summarize Data process.



The following set of parameters is defined for each technology, server, and instance which is installed in the Precise system.

Table 5-10 Configurable process parameters and their default values

Parameter	Default value
Summarize to daily level only top <i>n</i> rows	2000 rows
Summarize to weekly level only top <i>n</i> rows	3000 rows
Summarize to monthly level only top <i>n</i> rows	3000 rows

About viewing the status of the PMDB

Under the Warehouse menu, the view shows information on the PMDB load status of each instance in your Precise installation. The Status table in the Main area includes the following columns:

- **Server.** The name or IP address of the server on which the instance runs.
- **Technology.** The technology type of the instance.
- **Instance.** The name of the instance.
- **Unmonitored.** Indication by an icon when the instance is unmonitored.
- **First Loaded.** The date and time the first instance information was loaded into the PMDB.
- **Most Recently Loaded.** The date and time the latest instance data was loaded into the PMDB.
- **Status.** Icon indicating the status of the last load.

Using the Warehouse command buttons, you can get more detailed information on the actual PMDB tables that hold information on the specific instance. You can select any instance from the list and click Details to receive First Load and Most Recent Load information about specific tables in the PMDB.

You can also use the Warehouse Status Command selector to refresh the information that is displayed on the screen.

Viewing PMDB tables

You can examine additional information about the PMDB tables that are loaded for each instance, for example to understand their summary level.

To view the PMDB tables

1. Select an instance in the Main area table and click **Tables**.
2. In the PMDB Tables dialog box, investigate all relevant data.
3. Click **Close**.

Viewing the PMDB processes

You can examine additional information about the processes that are involved in loading information for this instance, for example to find out which tables were affected during the last execution of a process, which process failed during its last run, or when the next process is scheduled.

To view the run history

1. Select an instance in the Main area table and click **Processes**.
2. In the PMDB Processes dialog box, the top table lists all processes involved in loading information for this instance. Select the process you want to focus on, such as a process that failed during its last run (indicated by the status icon).
The last run of this process affects the bottom table populated with detailed information on the PMDB tables. If an error occurred, focus on the details in the **Error Message** column.
3. Investigate all relevant data. Then click **Close**.