

AdminPoint basics

This section includes the following topics:

- [Launching Precise AdminPoint](#)
- [About the tab selection bar](#)
- [About the Precise bar in AdminPoint](#)
- [About the filtering mechanism](#)

Launching Precise AdminPoint

AdminPoint is a central application that lets you manage your Precise product suite from one location. Through AdminPoint, you can view, control, and configure all major Precise components, such as agents or FocalPoints. You can also define other settings across products, such as environments, SLAs, or Groups. If you ever need to contact Precise Technical Support, AdminPoint provides an easy way of assembling the required support information.

In addition, AdminPoint lets you manage all relevant PMDB process settings. The PMDB is an integrated database that Precise uses to store historical performance and statistics data for all products and instances.

AdminPoint can be launched a number of ways, including via StartPoint. You can also launch to AdminPoint from any of the Precise product screens.

 To launch directly into AdminPoint, use your browser to jump to <http://url:port/admin>.

To launch AdminPoint

- On the StartPoint screen, click **AdminPoint** on the Precise bar. The Precise AdminPoint Dashboard tab opens. See [About the Dashboard tab](#).

About the tab selection bar

The tab selection bar, at the top of the AdminPoint screen, lets you view any of the following tabs:

- Dashboard
- Agents
- Licenses
- Setup
- Warehouse

In each tab, specific command buttons let you perform actions on the kind of information that you are viewing. The AdminPoint user interface includes the following elements:





- Precise bar menu for AdminPoint
- Tab selection bar
- Main area and details area for a selected tab






See the appropriate section for specific details regarding a tab.

About the Precise bar in AdminPoint

The Precise bar lets you configure Favorites, support, Precise settings, or launch to StartPoint. The following table gives a description of the screens available on the Precise bar.

Table 2-1 Precise bar in AdminPoint

Icon	Name	Description
	Back	During a work session, Precise products keep track of where you have navigated to. The Back, Forward, and History buttons enable you to navigate between previously visited views. The Back control displays your previous view.
	Forward	Enables you to navigate to the next view. This button is only enabled if you clicked Back, or if you chose a history option.
	Stop	Stops a request for information from the server.
	Refresh	Updates the data currently displayed.

	Home	Navigates to the highest level entity, usually the instance or AppTier (all instances). The time frame settings remain the same.
	Send	Opens new email message in your email program with the link to the current environment in context.
	Help	Opens the online help in context.
	Favorites	Enables you to add or remove favorites in your Precise Favorites list. See Adding, viewing, and deleting Favorites .
	Settings	Lets you define Precise settings in StartPoint. See Launching StartPoint .

About the Print function

The new Print function includes the following features:

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



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

Recommended print settings

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

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

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

Additional information

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

To print, save, or preview a currently displayed screen

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

Adding, viewing, and deleting Favorites

The Favorites function allows you to save a specific location in Precise and to retrieve the same location later without having to navigate to it. For this purpose, a friendly user interface similar to the familiar Favorites option in the Internet Explorer has been implemented.

About the Favorites function

The new Favorites function includes the following features:

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

To add a new Favorite location

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

To view a Favorites location

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

To delete an existing Favorite location

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



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

Generating a support file

When you contact Precise Technical Support you need to include in your help request a set of log files that assists the Support staff to determine which components cause the current problem. The Gather Support Information dialog box provides a quick and convenient way of generating all required files and collecting them in one zipped support file. You can create an express support file or a full support file.

To generate a support file

1. On the Support menu, in the AdminPoint toolbar, click **Support**.
2. In the Gather Support Information dialog box, select the servers that experience the problem and click the right arrow to move them to the Selected Servers list.
3. Select one of the following options:
 - **Express**. Creates a support file that holds detailed information about the selected server so that Support staff can determine the faulty component on the server.
 - **Full**. The Full support file works exactly the same way as Express, but provides more files and information. The servers are still selected by the user.
4. Click **Create Support File**. A dialog box indicating the progress of the support file creation opens. Creating the file may take several minutes.
5. Open the support file, or save it locally on your computer.
6. Send the file to Precise Technical Support.



Unless instructed otherwise by Precise Technical support, select express support.

Precise automatically adds support information from related FocalPoint servers even though you have not explicitly selected these servers.

Support files require disk space on the main Precise FocalPoint machine. For an Express support file, 5M is needed per server. For a Full support file, 100M is needed per server.

Sending an email message

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

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

To send an email message

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

Determining which table columns to display

Tables are used to display information about a set of related entities in the Main and Association areas. It is possible to determine which columns to display in the Association area tables.

To determine which columns to display in the Association area

1. Click the Table icon on the upper right-hand side of a table and select **Column Chooser**.
2. In the Table columns dialog box, click the arrows to move the names of the columns that you want to display to the Visible box and the ones that you do not want to display to the Invisible box.
3. Click **OK**.

About Clear All Filters

In StartPoint the filtered table columns can be cleared.

To clear the filters

- Click the Table icon on the upper right-hand side of a table and select **Clear All Filters**.

The table will be refreshed.

Copying data to the clipboard

At times you may want to save data displayed in the table area in a Microsoft Excel spreadsheet for further analysis or save an image of a graph to the clipboard.

To copy data displayed in the Association area to the clipboard

- Click the Table icon on the upper right-hand side of a table and select **Copy to clipboard**.

The table can be pasted into Microsoft Excel or as an HTML file.

To copy a graph to the clipboard

- Right-click a graph and choose **Copy to clipboard**.

You can now paste the image into any application that works with the clipboard.

Exporting to ASD

The Export to ASD (Application Service Dashboard) Portlet feature enables you to export the view of the chosen table or graph and generate a portlet with that view in ASD, so that it will provide you with another way of monitoring your environment.

Prerequisites


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

- View permissions to all AppTiers in the environment

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

Exporting the information

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

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

To export a table view

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

To export a graph view

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

Launching StartPoint

StartPoint can be launched via AdminPoint.

To launch StartPoint

- On the AdminPoint screen, click **StartPoint** on the Precise bar. The Precise StartPoint opens.

About the filtering mechanism

This section includes the following topics:

- [Top *n* filtering](#)
- [Retention of totals and partial filtering](#)

- [Two-phase filtering](#)
- [Ranking formula](#)
- [Technology-specific implementations](#)
- [Additional filtering during summary](#)

Top n filtering

The filtering mechanism is based on top n . We rank all activities and then apply the top n filtering. Subsequently, only the top n activities will be loaded. The rest will be filtered out.


This approach has two major benefits:

- **One threshold fits all.** The same n threshold can fit both production and test systems. Loading top activities ensures that most important activities will not be filtered out.
- **Predictability.** The n threshold limits the number of rows loaded into the PMDB. Thus, it should be a lot easier to predict the PMDB size.

Some technologies collect data and prepare multiple performance tables to be loaded. When necessary, multiple n thresholds are defined to allow a more granular control on the filtering.

All filtering parameters are defined in the PMDB (per technology) Load process. Filtering thresholds (as well as any other PMDB process parameters) can specify a default value per technology plus specific values for selected instances.

Retention of totals and partial filtering

 Retention of totals and partial filtering is not new. It has been in use since version 7.1.

If the filtering mechanism just filters out rows, the sum of all remaining rows will be less than the actual totals. Thus, all filtered rows are aggregated to one row. This row will be loaded into the PMDB and will retain the totals. The value of all identifiers in the filtered row is “_I3OTHER_”.

Most string identifiers are not kept in the statistics (fact) tables. String values are kept in dimension tables. Integer keys are kept in the statistics tables instead. In all dimension tables a zero key is reserved for the “_I3OTHER_” value.

Some technologies (mainly databases) use partial filtering, meaning that few identifiers of the filtered rows are kept as is (protected identifiers). Other identifiers are changed to “_I3OTHER_”. The filtered rows are aggregated and loaded into the PMDB. Instead of loading only one row to retain totals, multiple rows of partially filtered rows are loaded.

The amount of partially filtered rows depends on the number of different combinations of values of the protected identifiers. To control the amount of partially filtered rows, version 8.5 introduced [Two-phase filtering](#).

Two-phase filtering

This filtering mechanism works as follows:

1. Aggregate all raw data groups by all identifiers of the table.
2. Sort aggregated rows, ordered by the primary counter, descending.
3. Calculate rows count as the number of aggregated rows.

Calculate instance totals as the instance level totals for the primary counter. Calculate top $n-1$ totals as the primary counter totals for the top $n-1$ rows.

1. If rows count is less than the top n parameter, then exit.
2. If top $n-1$ totals divided by instance totals is greater than 0.9, then skip to step 9.
3. For each aggregated row which is not part of the top $n-1$ rows, change less important identifiers to _I3OTHER_ (partial filtering).
4. Aggregate again to combine _I3OTHER_ rows.

Because we did partial filtering, _I3OTHER_ rows will be aggregated to more than one row. These rows will compete with the top $n-1$ rows in step 9.

1. Sort aggregated rows, ordered by the primary counter, descending once again.
2. Aggregate all rows that are not part of the top $n-1$ rows into one row (the n row).

Example of Precise for Oracle using the filter mechanism

Table PW_ORSO_STMT_OBJ_STATS_T is the most detailed statistics table. It contains I/O and locks counters for combinations of Program, User, Machine, Statement, File, and Object (Table or Index).

Prior to version 8.5, we had to choose between two bad alternatives:

- If we specify the Object identifier as non-filtered, hundreds of partially filtered rows are loaded, no matter how high the filtering threshold is. At least one row per accessed object will be loaded regardless of its access time.
- If we specify the Object identifier as filtered, we may miss a top accessed Object just because it is accessed by many different Statements and Users.

Since version 8.5:

- We will specify that an Object is a non-filtered identifier.

- Objects that are accessed by many Statements and Users will be aggregated per Object. These partially filtered rows will compete with the non-filtered rows. Only the most accessed Objects will make it into the top n .

This way we can protect our most important identifier and at the same time avoid flooding the PMDB with low values of that identifier.

Ranking formula

Products that consider Service Time (Avg) as more important than Service Time (Summed) and also use Service Time (Avg) as the default sort field have additional considerations.

Defining Service Time (Avg) as the primary counter is problematic. Sometimes excessive execution of a highly responsive transaction is the main problem. We do not want to lose that transaction for one-off slow transactions.

Defining Service Time (Summed) as the primary counter is problematic as well. The activities with worst Service Time (Avg) might be filtered out.

Thus, in version 8.5, Precise for Web and Precise for J2EE define the primary counter as a formula. This formula is used to rank all activities by the top n mechanism.

The ranking formula is: Service Time (Avg) * Logy Executions.

The user can influence the weight of Executions as High, Medium, or Low. The filtering mechanism will chose the y base of the log function as follows:

- **High.** No log function. Ranking based on Service time (Summed). If an activity is executed 1,000,000 times, it will be weighted 1,000,000 times more.
- **Med.** Base is 2. If an activity is executed 1,000,000 times, it will be weighted 20 times more.
- **Low.** Base is 10. If an activity is executed 1,000,000 times, it will be weighted 6 times more.

Technology-specific implementations

Oracle

Precise for Oracle loads three sets of data:

- Programs (plus users, machines, modules, actions, and work types)
- Statements (plus PL/SQLs, execution plans, and all the above)
- Objects (plus files, Storage devices, and all the above)

Each combination of program, user, etc. may execute many SQL statements. Each combination of program, user, and statement may access several objects. Thus, the objects table may grow very fast.

By default, Precise for Oracle will load only top 300 programs, top 500 statements, and top 700 objects per 15 minutes. Top items are ranked according to In Oracle (Summed).

In Oracle (Avg) is not taken into consideration at all. The focus here is on top resource consumers.

To change the defaults

- Go to **AdminPoint>Warehouse>Processes** and select **Load Data process of Oracle technology>Parameters**.

Precise for Oracle uses the [Two-phase filtering](#) technique as follows:

- **Program.** After the first aggregation phase, all records that are not part of the top 300 are partially filtered by protecting all programs, modules, and users. All other identifiers are converted to `_I3OTHER_` and then the second aggregation phase takes place. The partially filtered records compete with the original top 300 records. When an ERP extension is installed, the user is overridden by the ERP user. In this case it will not protect the user any more. If many SAP users executed the same SAP transaction, after second aggregation phase all of them will be aggregated into one record. So, this protected SAP transaction is promoted and may make it into the top 300 programs.
- **Statement.** After the first aggregation phase, all records that are not part of the top 500 are partially filtered by protecting all statements and execution plans. All other identifiers are converted to `_I3OTHER_` and then the second aggregation phase takes place. The partially filtered records compete with the original top 500 records.

If many users and/or programs executed the same statement, after the second aggregation phase all of them will be aggregated into one record. So, this protected statement is promoted and may make it into the top 500 statements.

- **Object.** After first aggregation phase, all records that are not part of the top 700 are partially filtered by all objects' name, schema, and type. All other identifiers are converted to `_I3OTHER_` and then the second aggregation phase takes place. The partially filtered records compete with the original top 700 records.

If many statements and/or programs access the same object, after the second aggregation phase all of them will be aggregated into one record. So, this protected object is promoted and may make it into the top 700 objects.

DB2

The Precise for DB2 load mechanism works the same as Precise for Oracle. The only difference is: Precise for DB2 collects and loads only the first two sets of data: Programs and Statements.

The "Calculate Objects Usage" PMDB process calculates the Objects data daily based on the Statements performance and execution plan.

SQL Server and Sybase

Precise for SQL Server and Precise for Sybase implement a similar top n filtering mechanism. Both technologies generate and filter only one set of records: each record contains a combination of database, user, program, login, machine, statement, and batch.

The [Two-phase filtering](#) technique is used as follows:

- After the first aggregation phase, all records that are not part of the top 500 are partially filtered by protecting all databases, users, and programs. All other identifiers are converted to `_I3OTHER_` and then the second aggregation phase takes place. The partially filtered records compete with the original top 500 records.
- If a user and/or program executed many statements, after the second aggregation phase all of them will be aggregated into one record. So, this protected combination of database, user, and program is promoted and probably makes it into the top 500 statements.

As specified above, by default only top 500 statements will be loaded. To change the default: go to **AdminPoint>Warehouse>Processes** and select **Load Data process of the SQL Server or Sybase technology>Parameters**.

Web

Precise for Web has both pre-aggregation and post-aggregation filtering.

- Pre-aggregation filtering
- Post-aggregation filtering

Pre-aggregation filtering

By default, a client IP address is not loaded into the PMDB. The original client IP address is converted to Location (if defined), Country, State, and City (if enabled and if found in the Precise IP addresses database). Then, the original IP address is replaced with the "N/A" value.

When the raw data will be aggregated, the thousands of distinct IP addresses will be aggregated into a few dozen combinations of: location, country, state, and city.

If the monitored application is a local intranet application serving less than 100 distinct client IP addresses, you may consider enabling Collect Client IPs in **AdminPoint>Warehouse>Processes** and select **Load Data process of the Web technology>Parameters**.

Post-aggregation filtering

Precise for Web loads two different sets of data:

- **Web server.** Major entity is URL.
- **Web client.** Major entity is Web Page.

For each of the two sets you can configure the maximum number of rows to be loaded. Top records will be loaded based on a formula that weights both average service time and number of executions as described in [Ranking formula](#).

By default, top 1000 URL records and top 100 Page records will be loaded each 15 minutes. In both sets, the weight of executions is **MEDIUM**. If you want to focus on the worst average response time, change the weight of executions to **LOW**. If you want to focus on the top resource consumers, change the weight of the executions to **HIGH**.

All the above can be configured for all Web instances or per instance using **AdminPoint>Warehouse>Processes** and select **Load Data process of the Web technology>Parameters**.

OS, MQ, Tuxedo, OA, and Other

Insight technologies have post-aggregation filtering.

Insight loads several different sets of data including OS, MQ, Tuxedo, OA, and Other. For each of these sets you can configure the maximum number of rows to be loaded. Top records for MQ, Tuxedo, OA, and Other will be loaded based on a formula that weighs both the average service time and the number of executions, as described in the Ranking formula. Top records for OS will be loaded based on a formula that weights both average memory consumption and CPU consumption, similar to the method described above for the other technologies.

By default, the top 200 records will be loaded every 15 minutes. In all sets, the weight of executions is **MEDIUM**. If you want to focus on the worst average response time, change the weight of executions to **LOW**. If you want to focus on the top resource consumers, change the weight of the executions to **HIGH**.

All the above can be configured for all instances or on a per instance basis using **AdminPoint>Warehouse>Processes** and selecting **Load Data process of the OS | MQ | Tuxedo | OA | Other technology>Parameters**.

Microsoft .NET, J2EE and TPM

Precise for J2EE has two data filtering mechanisms, one that takes place inside the client JVM (where it can be defined as activated or deactivated), and one that takes place within the PMDB. For information about data filtering inside the client JVM, see "About configuring data filtering" on page 105.

Version 9.5 introduces an additional data aggregation interval for Precise for J2EE and Precise SmartLink - the Short Time Slice (STS), that can be defined as either 30 or 300 seconds. To configure this aggregation interval, see "Modifying the Precise Load Data slice size" on page 77.

Additional filtering during summary

All the above mentioned filtering refers to filtering the performance data prior to the load into the PMDB. This section deals with additional filtering applied after loading data into the PMDB, when summarizing data into daily, weekly, and monthly tables.

The following example illustrates the need for additional filtering.

For example, we load the top 500 SQL statements per instance, every 15 minutes. The daily summary process aggregates the entire day (96 time slices of 15 minutes each) into one row per SQL statement that contains its daily activity.

The best case scenario is when the application is static and the same top 500 statements are loaded every 15 minutes. In this case, the daily summary will only load 500 statements to the daily tables.

The worst case scenario is when a completely different set of top 500 statements are loaded. In this case nearly 50,000 (500*96) rows will be loaded into the daily tables per instance per day.

In real life, 500 rows times 96 slices will be aggregated to something between 5,000 to 20,000 rows (depending on the application).

With additional filtering, while aggregating the rows from the time slice tables, only the top *n* rows will be loaded into daily tables (plus one row per hour group that retains totals).

The same mechanism is applied when summarizing daily tables into weekly tables and monthly tables. By default, the *n* threshold is defined per summary level as follows:

- **Daily.** Top 2,000 rows will be loaded into daily tables.
- **Weekly.** Top 3,000 rows will be loaded into weekly tables.
- **Monthly.** Top 3,000 rows will be loaded into monthly tables.

The Precise administrator can modify these three thresholds per technology or per instance, using the **Parameters** button when the PMDB Summarize Data process is selected.

In version 9.0, the new STS (short time slice) table is introduced for Precise for J2EE and Precise for SmartLink. The Precise administrator cannot modify the top *n* threshold value when summarizing a J2EE or SmartLink tree-like table from the STS level to the regular, cube-like, T (15 minute) level. This threshold is defined as the technology's default daily value, and this summary filters the top *n* trees (each tree is an activity with multiple rows), into a daily table of rows.