

Precise for Sybase

User Guide

Version 10.1

IDERA

Precise for Sybase User Guide

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

Document release version 1.0

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

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

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

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

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

Global Headquarters

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

Third-party legal notices

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

Customer Support

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

Contents

- Introducing Precise for Sybase12
 - About Precise for Sybase12
 - Improved performance management.....12
 - Integrating with other Precise products12
 - Components of Precise for Sybase13
 - Precise for Sybase agents.....13
 - Precise Listener and Precise for Sybase FocalPoint.....14
 - PMDB Loaders.....14
 - How Precise for Sybase works14
 - Where to get more information14
- Precise for Sybase basics15
 - About Precise for Sybase tabs15
 - How most tabs are structured15
 - About the Precise bar.....16
 - About the Main area.....18
 - About the Association area18
 - About drilling down in context18
 - How a drilldown affects the Tab heading.....19
 - About configuring your settings19
 - About configuring general settings19
 - About configuring tuner settings19
 - About configuring display settings19
 - About configuring Tree view settings19
 - About configuring Time Frame settings19

Tasks common to most tabs	20
Switching to a different tab.....	20
Selecting a time frame	20
Selecting an entity.....	21
Filtering data.....	21
Associating entities with data that meet specific criteria	21
Saving a statement.....	22
Focusing on information in overtime graphs	22
Sending an email message	22
Adding, viewing, and deleting Favorites	24
Determining which table columns to display	24
Copying data to the clipboard.....	24
Exporting to the Precise Custom Portal	24
About role management in Precise for Sybase	24
How information is collected.....	25
About duration values.....	25
About session states	25
About Statement identifiers	27
How Precise for Sybase monitors availability	27
Launching Precise for Sybase from the Precise StartPoint.....	27
Usage concepts.....	29
About application performance tuning	29
About viewing Current Activity.....	29
Analyzing instance performance.....	30
Identifying resource consumers.....	30
Associating with offending SQL statements.....	30
Analyzing previous blocking situations	31
About SQL statement tuning	31

- Identifying the most resource consuming statements or batches31
- Collapsing constants into bind variables31
- About viewing an execution plan31
- Understanding an execution plan32
- Analyzing the statement execution plan over time32
- Examining object space usage over time32
- About instance statistics tuning32
 - Observing instance statistics over time.....32
- About getting an overview of your application33
 - About the Dashboard tab33
 - How the Dashboard tab is structured33
 - About the Main area in the Dashboard tab34
 - About the Instance Details area in the Dashboard tab35
 - How the Dashboard tab can help you identify performance problems39
 - Determining which instance consumes the most resources.....39
 - Identifying the highest resource consumers in various entities.....39
 - Examining instance statistics39
 - Monitoring a session activity during the last minute39
- Getting current application information40
 - About the Current tab.....40
 - How the Current tab is structured.....41
 - About the Main area in the Current tab.....41
 - About the Association area in the Current tab.....41
 - About the entities you can examine in the Current tab41
 - About the Instance entity42
 - About the Database entity42
 - About the Session entity43
 - About Statement and Batch entities.....45

- How the Current tab can help you identify performance problems47
 - Examining the resource consumption of an entire instance and database47
 - Observing session activity47
 - Examining a single connection47
- Examining performance over time49
 - About the Activity tab49
 - How the Activity tab is structured49
 - About the Main area in the Activity tab50
 - About the entities you can examine in the Activity tab51
 - About the All Instances entity51
 - About Instance and Collector entities51
 - About Statement, Batch, Collapsed Statement, and Collapsed Batch entities54
 - About getting an overview of Statement, Collapsed Statement, Batch and Collapsed Batch entities55
 - About viewing Statement, Collapsed Statement, Batch and Collapsed Batch text56
 - About viewing In Sybase data for Statement, Collapsed Statement, Batch and Collapsed Batch entities56
 - About viewing Statement, Collapsed Statement, Batch and Collapsed Batch entities locking56
 - About viewing scalability for Statement, Collapsed Statement, Batch and Collapsed Batch entities62
 - Statement and Batch entities Association area62
 - Collapsed Statement and Collapsed Batch entities Association area63
 - About Session and Blocker Session entities64
- How the Activity tab can help you identify performance problems63
 - Examining resource consumption of an entire instance and database63
 - Identifying heavy resource consumers64
 - Examining resource consumption over time64
 - Resolving blocking situations65
 - Identifying offending SQL statements65
 - Examining the scalability of your application or Sybase servers66

Examining the relationships between database objects67

- About the Objects tab67
- How the Objects tab is structured67
 - About the Main area in the Objects tab67
 - About the Association area in the Objects tab67
- About the entities you can examine in the Objects tab67
 - About the Instance entity67
 - About the Database entity68
 - About the Table entity and Tables Contain Same Column entity69
 - About the Index entity71
 - About viewing data storage and fragmentation information for the selected index72
 - About the Column entity73
 - About the View entity and Views Contain Same Column entity74
 - About the User entity74
 - About the Database Device entity75
 - About the Segment entity75
 - About the Data Cache entity76
 - About the Unused Indexes entity76
 - About the Statement entity76
- How the Objects tab can help you identify performance problems80
 - Examining an object's space usage over time80
 - Examining fragmentation of table data and indexes80
 - Examining distribution statistics of indexes and columns80

Examining Sybase instance statistics81

- About the Statistics tab81
- How the Statistics tab is structured81
 - About the Main area in the Statistics tab81
 - About the Association area in the Statistics tab82

About the entities you can examine in the Statistics tab	82
About the Tier entity.....	82
About the Instance entity	83
About the Engine entity	91
About the Device entity	91
About the Data Cache entity.....	92
About the Performance Counter or Counter Instance entity.....	92
About the Wait Events entity.....	92
About the Replication Server entity.....	93
About the Replicate(d) Databases entity	93
How the Statistics tab can help you identify performance problems	93
Examining the kernel utilization of Sybase over time	94
Examining I/O statistics.....	94
Examining network activity.....	95
Examining memory bottlenecks.....	95
Examining a specific counter over time	96
Evaluating explain plans	97
About the SQL tab.....	97
How the SQL tab is structured.....	98
About the View tabs.....	98
About viewing the execution plan of a statement	98
Observing a history of resource consumption	105
Breaking down a batch's access plan into statements.....	106
About the Actions tab	107
Creating a new statement.....	107
Opening an existing statement	108
Editing an existing statement.....	108
Re-Explaining an existing statement	108

Display the Sybase plan using Sybase SQL Advantage and Interactive SQL	108
How the SQL tab can help you identify performance problems	111
Identifying problematic operators in the latest access plan	111
Locating referenced tables that pose potential problems	111
Examining how access path changes affect statement performance	112
Locating the most resource-consuming statements in a batch	112

Introducing Precise for Sybase

This section includes the following topics:

- [About Precise for Sybase](#)
- [Components of Precise for Sybase](#)
- [How Precise for Sybase works](#)
- [Where to get more information](#)

About Precise for Sybase

Addressing the needs of DBAs, system administrators, managers, quality assurance specialists and planners, Precise for Sybase provides a complete view of application performance by capturing, measuring, and correlating performance metrics from all critical system components.

Precise for Sybase offers a wide range of data presentations to enable quick and accurate monitoring and performance analysis.

Current operational data is displayed to enable easy identification of problematic conditions. Recent operational data, for designated time periods, supplements operational data to enable operations staff to detect problematic conditions that occurred in the past and that are likely to occur again and again, if not corrected.

Performance data is archived and a fully detailed display of the historical data is provided for use in performance, QA and capacity planning analysis.

Improved performance management

IT staff frequently search for answers to questions like "Why did the application run so slowly last night?" or "Where is the application spending its time?"

Precise for Sybase can help IT staff find the answer to these questions. The first step is to establish a baseline level of database resource consumption by understanding the behavior of the average application. The behavior of the average application represents the average percentage of the total time that all applications spend consuming or waiting for each resource. These aggregate figures provide a broad picture of the operation of your database. They enable you to identify major bottlenecks in your system, for example, CPU and I/O limitations, Sybase database waits due to locks, remote queries, and so on.

Precise for Sybase also enables you to focus on and drill down to individual applications, SQL statements and database objects in your system that are responsible for heavy consumption of system and database resources and to quickly find the definitive root cause of performance degradation.

Integrating with other Precise products

Precise for Sybase has been integrated with Alerts, Precise Custom Portal, and Report Manager in addition to the PMDB.

Alerts provides a notification system that will alert you to application response time problems and enables you to launch Precise for Sybase in context to view more detailed data.

For more information, see the *Alerts User's Guide*. The Precise Custom Portal, is a lightweight, configurable portal application that provides access to various types of data from different applications. The Precise Custom Portal features a Web-based view. It is highly customizable and extensible, allowing you to build a dashboard for each user or each function within your organization. For example, you can build a dashboard for all the information that a system administrator or director of IT operations would need to constantly follow up on. This component is automatically installed as part of the Precise framework installation.

For more information, see the *Precise Custom Portal User's Guide*.

Report Manager gathers and organizes historical information that enables the IT team to focus on infrastructure hot spots. Report Manager is a reporting tool that resides on a Web server. You can use Report Manager to compare period-to-period performance against a baseline and identify response time problems “at-a-glance”, before they affect your bottom-line. It queries the PMDB and displays the results in reports that are generated on a scheduled basis or on demand. For more information, see the *Report Manager User's Guide*.

The PMDB is the repository of the Precise for Sybase suite which stores application performance data for long-term analysis and trending reports. The data is collected by the Precise product suite. For more information, see the *Precise Administration Guide*.

Components of Precise for Sybase

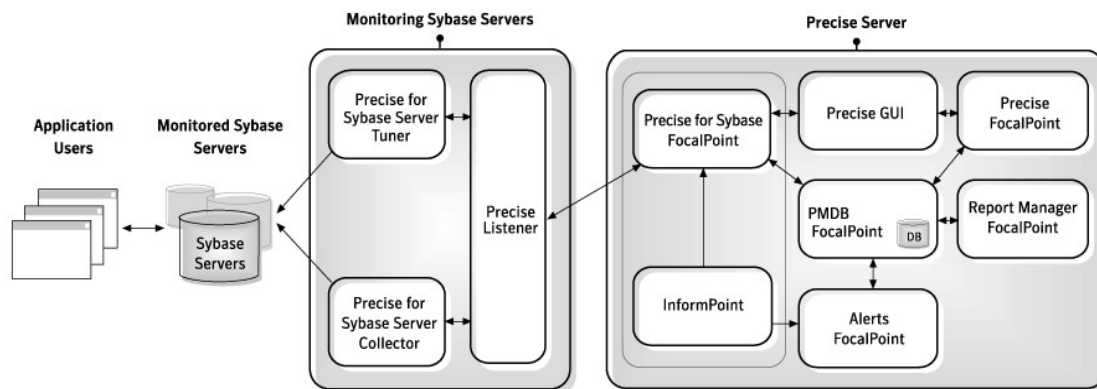
Precise for Sybase employs a client/server architecture to provide high performance database monitoring and tuning. The following figure illustrates the typical architecture of Precise for Sybase.

Precise for Sybase consists of the following components:

- One or more Precise for Sybase Collector agents installed on one or more servers in your application
- Precise for Sybase FocalPoint
- PMDB

The figure below illustrates the typical architecture of an installation.

Figure 1-1 Precise for Sybase architecture



Precise for Sybase agents

The Precise for Sybase Collector agent is installed on the monitoring server machines and communicates using TCP/IP. The Collector agent collects and analyzes Sybase database performance data, which is then stored in flat historical files and later in the PMDB.

The Collector agent monitors and records all Sybase database activity on a 24 x 7 basis. The Collector agent works using a sampling technique and thereby consumes less resources than a trace agent.

The Collector agent collects resource consumption data on the following Sybase data entities: databases, sessions, programs, statements, batches, login names, DB users, and machines.

The Precise for Sybase Collector agent uses the following sources to collect performance data:

- Sysprocesses** This system table provides information regarding the resource consumption in Sybase.
- Monitoring Tables** The Monitor Tables provide the text of the improvised batches, executed in Sybase.

The resource consumption data and the text batches are correlated according to the batch ID of the two resources.

NOTE The text of the Stored Procedures is taken from the Syscomments system table.

Precise Listener and Precise for Sybase FocalPoint

Each Precise product has a FocalPoint that communicates with the Listeners installed on the monitored servers. The Precise Listener communicates with other Precise components on the same server. For example, the PMDB FocalPoint requests to load data into the PMDB are sent to the Precise Listener. Then, the Precise Listener executes the PMDB Loader to aggregate and prepare performance data to be loaded.

PMDB Loaders

PMDB Loaders periodically load data, collected by the Collector and stores it in the historical files, into the PMDB. When the PMDB Loaders run, the flat historical files, written after the last time data was loaded into the PMDB, are identified and opened. The summarized data is also loaded into the PMDB. Data on the last incomplete time slice is only loaded into the PMDB when the next scheduled load process runs.

How Precise for Sybase works

Precise for Sybase stores a complete history of your database activity in the PMDB. This provides you with a performance tracking tool that gives you the ability to assess database performance trends and analyze changes in database activities over long periods of time. This data is vital for identifying patterns in resource consumption and predicting future resource consumption in your Sybase database.

The Precise for Sybase Collector agent runs on your monitoring server machine, using a unique, accurate and comprehensive sampling technology, that requires low overhead. It collects performance data from an instance and then correlates and loads the data into the PMDB tables. This data is then analyzed and summarized by the FocalPoint. You can later access this data using the Precise for Sybase Web browser interface.

Where to get more information

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

Precise for Sybase basics

This section includes the following topics:

- [About Precise for Sybase tabs](#)
- [How most tabs are structured](#)
- [About drilling down in context](#)
- [About configuring your settings](#)
- [Tasks common to most tabs](#)
- [About role management in Precise for Sybase](#)
- [How information is collected](#)
- [Launching Precise for Sybase from the Precise StartPoint](#)

About Precise for Sybase tabs

Precise for Sybase continuously collects high-quality metrics from each supporting tier of your Sybase-based custom applications infrastructure and correlates these metrics to build a clear picture of application performance from the end-user's perspective - that is, response time.

This information is displayed in the Precise for Sybase tabs and provides you with the information necessary to successfully track your database's performance and identify patterns in resource consumption.

The table below describes which tasks can be performed in each tab.

Table 2-1 Precise for Sybase tabs

Tab	Description
Dashboard	Allows you to quickly identify which Sybase instance has a performance problem, providing the first glimpse into where the problem lies.
Current	Allows you to solve the performance problems that are happening now.
Activity	Allows you to examine the application's performance over time and helps locate performance bottlenecks in its behavior.
Objects	Allows you to identify which Sybase tables are used inefficiently by the application, and their effect on the application components.
Statistics	Allows you to identify performance bottlenecks in the way Sybase instance resources are used.
SQL	Allows you to tune problematic SQL statements.

How most tabs are structured

Though each tab is structured differently, most tabs consist of two different areas. Each area can include different control elements, such as tabs and view controls, and displays information in various formats, such as tables, graphs, or charts. The various areas are related to each other in that performing an action in one area affects the information displayed in other areas on the page. For example, the Current tab contains an upper and a lower area. The lower area (Association area) shows information in a table format. Each row in the table represents an instance or entity. The upper area (Main area) displays general information on the selected instance or entity, in context.

The entities displayed in the Association area are associated with the selected entity displayed in the Main area. At times, the relationship between the entity displayed in the Main area and those displayed in the Association area is that of parent to child (such as in the case of Instance to Database), and sometimes it merely represents that there is a relationship between the selected entity and the entities displayed in the Association area (such as when Statement is the selected entity displayed in the Main area and the list of Batches associated with the selected statement is displayed in the Association area).

It is possible to click on a row and drill down to an additional level below the selected level or use the filter feature to determine which information is displayed in the table. When you perform an action on an element in the lower area (in this case the Association area) the information displayed in the upper area (in this case the Main area) changes to reflect the action you requested.

You can control the tab by doing any of the following:

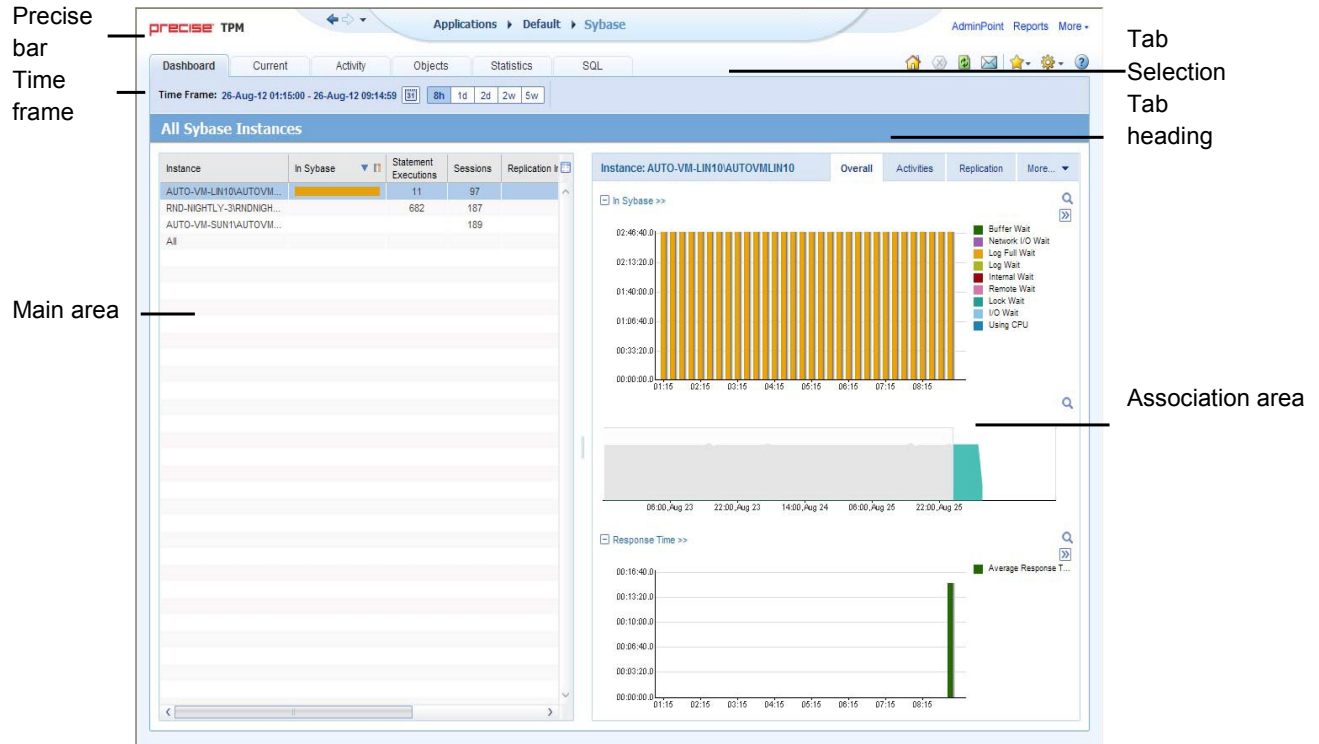
- Keeping the selected entity but changing the associated entities - by selecting a different entity from the Association controls in the Association area.
- Changing the selected entity - by selecting a row in the Association area.
- Changing the view for the selected entity - by selecting a different view in the Main area.

For example, when you want to view information on a specific session in the Current tab, choose Sessions from the Association controls. The Association area changes to display session-related information. The Tab heading and the Main area remain unchanged.

If you want to view detailed information on a specific session, click on the session in the Association area. The Tab heading indicates the newly selected entity; the Main area displays In Sybase information on the session you drilled down to, and the Association area shows information on statements executed by the session you selected.

The following figure provides a description of the common key elements in the Current tab and what happens when you drill down to a specific entity.

Figure 2-1 How a typical tab is structured



About the Precise bar

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

Table 2-2 Precise bar functions

Icon	Name	Description
	Back	During a work session, keeps track of where you have navigated to. The Back 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.
	AdminPoint	Launches Precise AdminPoint.
	Home	Navigates to the highest level entity, usually the instance or Tier (all instances). The time frame settings remain the same.
	Stop	Stops a request for information from the server.
	Refresh	Updates the data currently displayed.
	Favorites	Enables you to add or remove favorites in your Favorites list.
	Send	Opens a new email message in your email program with the link to the current application in context.
	Settings	Opens the General Settings, Transactions & Users Settings, and Time Frame Settings dialog boxes.
	Help	Opens the online help in context.

About the Main area

The Main area displays general information on the selected entity. The structure of this area depends on the selected entity and tab.

Information can be displayed as follows:

- As two over-time graphs that are displayed side-by-side and describe two data series.
- As a pie chart and a table that are displayed side-by-side. The table displays more details regarding a specific slice in the pie chart.
- As Lock trees and a corresponding Lock Breakdown table.

The View Control drop-down menu allows you to choose different view options. You can access this menu by clicking the arrow icon.

In some tabs, you can filter the contents of the view by clicking the Filter button. See [Filtering data](#) on page 21.

About the Association area

The Association area lists entities that have similar attributes and that are related to the selected entity named in the Tab heading. Besides the Tab bar, the Association area is the main navigation tool. It enables you to drill down from one entity to related entities (usually of a different type).

The entities displayed in this area depend on the selected tab and the selected entity. The same entity may be displayed in several tabs, with different data displayed for each one (for example, a Sybase table will be displayed in the Objects tab with data from the Sybase data dictionary, and in the Activity tab with data sampled by the Collector agent).

The Association control, accessed by the arrow icon, allows you to associate with different entities. The last option, More..., opens a dialog box that lets you carry out additional operations, such as to view additional associations, change the sort order, or control the number of returned rows.

The information displayed in the Association area is arranged in a tabular format. Each row in the table represents an entity, and each column describes an attribute of the entity, such as text or a graph. It is possible to change the sort order of the entities in the table by clicking a column title.

The Association tabs located above the table let you display different sets of attributes for the entities displayed in the table. The Association tabs are entity-dependent. For some entities, no tabs are displayed.

A numeric column can optionally display a graph, and at times, it can display a stacked bar graph as well. Switching to a graph view in these columns offers some level of breakdown to the information displayed in the column.

The following icons are used to indicate which format information is displayed.



Displays the information in numerical format.



Displays the information in a bar graph.



Displays the information as a stacked bar graph.



Displays the information in percentage format.

The Association area enables you to drill down to another entity by clicking a row. The information displayed in both the Main area and the Association area will change to reflect your selection. See [Filtering data](#) on page 21. See [Associating entities with data that meet specific criteria](#) on page 21.

About drilling down in context

The term 'in-context' means that you can display additional information on a selected item by drilling down to another tab or view. The filter settings you defined (for example, the selected time frame you chose) and the entity you selected are carried over to the other view or tab, to allow you to continue analyzing your subject from a different perspective. This concept takes on slightly different meanings depending upon where you are attempting to drill down in context from.

For example, when viewing information on an instance in the Dashboard tab, you can click a link in the Instance Details area (right pane) to view additional information on the related tab, in context to your original selection. Alternatively, when viewing a list of alerts for your product in Alerts, you can launch your product in context and continue investigating the factors that led the system to issue that alert.

In short, the information displayed when drilling down in context is always related to your original selection's settings.

How a drilldown affects the Tab heading

The Tab heading displays the name of the currently selected entity on this screen. When you drill down to a new entity in the Association area, the Tab heading changes to reflect the name of the newly selected entity.

About configuring your settings

The following settings allow you to control the appearance and behavior of the user interface. They can be configured from the Settings menu on the Precise bar:

- General settings - Includes the General, Display, and Tree view tabs.
- Time Frame settings

About configuring general settings

On the General tab of the General Settings dialog box, you can adjust the following:

- Number of items that will be displayed for each association. From the list, select the maximum number of items that will be displayed in a table. The default is 50 items. You can control the number of items returned for a specific association by using the More... option in the Association controls. This way, you can display up to 500 items.
- Maximum number of characters to display in a text ToolTip. When you move your cursor over text columns in a table listing Statements or PL/SQLs, a ToolTip appears, displaying more text than can normally fit into a table cell. This field controls the maximum number of characters that can be displayed. The default is 500 characters.
- How often the information in the Current tab is refreshed. If you select the Auto-refresh Current tab every check box, the Current tab automatically refreshes at the specified interval. By default, the tab is not refreshed automatically. If you select this option, the default refresh interval is every five minutes.

About configuring tuner settings

In the Tuner tab you can change the Login name and Password used when launching Sybase SQL Advantage.

About configuring display settings

When you drill down to or choose a different entity, selecting the option on the Display tab of the General Settings dialog box allows you to maintain the selected view or tab when changing from one tab to another. This is sometimes referred to as sticky tabs.

About configuring Tree view settings

(Only applied to Precise products that include a Tree view)

On the Tree View tab of the General Settings dialog box, you can adjust the following:

- Number of categories to display when expanding an item.
- Number of items to display when expanding a category.

About configuring Time Frame settings

You can determine the resolution of the data that is displayed in the overtime graphs using the Time Frame Settings dialog box. By using this dialog box you can define the default time frame to display.

Tasks common to most tabs

The following tasks are commonly performed in most tabs:

- [Switching to a different tab](#)
- [Selecting a time frame](#)
- [Selecting an entity](#)
- [Filtering data](#)
- [Associating entities with data that meet specific criteria](#)
- [Saving a statement](#)
- [Focusing on information in overtime graphs](#)
- [Sending an email message](#)
- [Adding, viewing, and deleting Favorites](#)
- [Determining which table columns to display](#)
- [Copying data to the clipboard](#)
- [Exporting to the Precise Custom Portal](#)

Switching to a different tab

You can easily switch between the different tabs using the Tab Selection bar. When you start your Precise product, the Dashboard tab opens by default. For other Precise products, another tab will open by default. The button of the selected tab is displayed in orange.

To select a tab

- Click a button on the Tab Selection bar to display information on the selected entity in a different tab.

Selecting a time frame

You can configure Precise for Sybase to display transaction performance data for a specific time frame using the predefined time frame options or calendar icons.

Selecting a predefined time frame from the toolbar displays transaction performance data for the selected time period up to the current time. See [Selecting a predefined time frame from the Precise for Sybase toolbar](#) on page 20.

Selecting the time frame using the calendar icon, you can choose to define a time range independent of the current time, or to define a time range up to the current time. See [Selecting a time frame using the calendar icon](#) on page 21. The predefined time frame options are:

- Last 8 hours (8h) (default)
- Last 1 day (1d)
- Last 2 days (2d)
- Last 2weeks (2w)
- Last 5 weeks (5w)

The time frame selected affects all information displayed in Precise for Sybase. Only data that falls within the selected time frame is shown in these areas.

Selecting a predefined time frame from the Precise for Sybase toolbar

To select a predefined time frame

1. From the Precise for Sybase toolbar, select one of the predefined time frames.

Selecting a time frame using the calendar icon

To select a time frame

1. Click the calendar icon. In the dialog box that is displayed perform one of the following:
 - a. To define a time frame independent from the current time, select the 'Time Range' option and select the Start and End dates and times.
 - b. To define a time frame up to the current time, select the 'Last' option and enter the desired time frame.
 - c. To use one of the three previously used time frames, select the 'Recently used' option and from the drop-down menu select the desired time frame.
 - d. To use a previously saved time frame, select **Use a previously saved time frame** and from the drop-down menu select the desired time frame.
2. To save your settings for future access, select **Save these definitions for future use as:** and enter a name in the corresponding field.
3. Click **OK**.

Selecting an entity

You can select the entity that you want to examine in more detail. To select an entity

- On the Main area of the tab, select the entity you want to examine.

Filtering data

You can focus on specific contents of the Dashboard and Activity tabs by showing a subset of the information. This shows the contribution of entities such as programs and users.

You can define flexible selection criteria in the Filter dialog box. To open the dialog box, you need to click the Filter Off icon next to the Time Frame selector. Filter On or Filter Off indicates the current state of the filtering mechanism. Filter Off means all information is shown; Filter On means a subset of information is shown.

When you apply your selections in the Filter dialog box, the information displayed in both the Main area and the Association area is modified to reflect your selections. Also, the filtering continues to apply when you drill down to associated entities.

NOTE You can enter multiple criteria, in which case each criterion is applied using the logical AND operator.

To filter data

1. Click Filter is Off/Filter is On.
2. In the Filter dialog box, do the following for each entity you want to filter:

From the left drop-down list	Select an entity.
From the middle drop-down list	Select an operator, such as Like, < >, Not Like, In, Not In.
In the text box	Type the criteria (case sensitive) for the selected entity.

If you select the operator Like or Not Like, you can use the % wildcard character. If you select the operator In or Not In, type a comma to separate values.
3. Click **Ok**.

Associating entities with data that meet specific criteria

You can associate the displayed entity with specific data to focus your analysis.

NOTE The criteria no longer apply when you drill down to another entity.

To associate entities with data that meets specific criteria

1. Click the arrow located on the left of the Association controls and select **More....**
2. In the Associate With dialog box, on the Entries tab, select the entity you want to associate data with from the Populate table with list.

3. In the Sort entries by list, determine according to which criteria you want the information to be sorted and in which order.
4. On the Display top list, select the number of rows to display.
5. On the **Criteria** tab, do the following for each entity you want to associate data with:

From the left drop-down list	Select an entity.
From the middle drop-down list	Select an operator, such as Like, <>, Not Like, In, Not In.
In the text box	Type the criteria (case sensitive) for the selected entity. If you select the operator Like or Not Like, you can use the % wildcard character. If you select the operator In or Not In, type a comma to separate values.
6. Click **OK**.

Saving a statement

Save As allows you to save the statement or batch so that you can examine its access plan at a later date. To save a statement for future use

1. Select the statement you want to save in the Association area. You can click the Save As button only after you drill down to the statement and it moves up to the Main area.
2. On the Precise bar, click the **Actions** tab and click **Save As**.
3. In the Save Statement dialog box, enter the following properties:

Instance	The name of the instance that the statement belongs to.
Cabinet	The name of the cabinet that the statement is saved in (creates a cabinet if it does not already exist).
Folder	The name of the folder that the statement is saved in (creates a folder if it does not already exist).
Name	The name of the statement (generated randomly).
Database	The name of the database where the statement is running.
User	The statement's parsing user.
4. Enter the Sybase text for the statement in the Text frame.
5. Click **OK**.

The new statement is saved in the PMDB in a logical cabinet and folder hierarchy.

Focusing on information in overtime graphs

Some entities display an overtime graph. The overtime graph displays entity statistics over a specified time period. Depending on the number of points displayed in the graph, you may need to zoom in or out. The text displayed on the x-axis varies according to the time frame. If there is a year or day change, x-axis labels will display accordingly.

To focus on information in overtime graphs, select the desired time frame on the overtime graph, click and drag the left or right handle to increase or decrease the time range. The small zoom (spyglass) icon will display on the upper right of the selected time range, and above the overtime graph legend. Click the zoom icon to zoom in according to the selected time frame.

Sending an email message

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

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

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

Adding, viewing, and deleting Favorites

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

About the Favorites feature

The new Favorites feature includes the following options:

- 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 AutoComplete.
- Auto Naming – The Favorites dialog generates item names based on the current location.

UI description

A Favorites menu has been added to the Precise bar in each product including StartPoint.

An Add/Delete Favorites option under the Favorites menu allows you to save the current location or delete an existing one.

To add a new Favorite location

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

To view a Favorites location

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

To delete an existing Favorite location

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

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

Determining which table columns to display

Tables are used to display information about the 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**.

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

1. 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

1. 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 the Precise Custom Portal

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

Prerequisites

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

- View permissions to all Tiers in the application

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

Exporting the information

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

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

To export a table view

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

To export a graph view

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

About role management in Precise for Sybase

The Roles management feature in Precise for Sybase lets you assign roles to users and control whether they will be able to do the following operations:

- Gen_Estimated_Plan (explain)—lets you generate a new execution plan.
- Launch_Sybase_Tool—lets you launch external Sybase tools.

It is possible to grant permissions to all roles or to revoke permission from individual roles. Permission is a combination of permission type (such as: INSTALL, ADMIN, MONITOR) and operation (such as: VIEW, UPDATE).

Certain types of permission enforce permission dependencies. For example, if you grant a TUNE.RECOMMEND permission it will automatically grant TUNE.EXPLAIN. In the same manner, if you revoke TUNE.EXPLAIN then the TUNE.RECOMMEND permission is also revoked. You can configure permissions and roles in AdminPoint.

For more information on configuring permissions and roles, see the *Precise Administration Guide*.

How information is collected

Precise for Sybase collects and displays performance data that enables quick and accurate performance monitoring and analysis. To understand the information displayed in the various graphs, tables and views, it is important to understand how information is collected.

About duration values

Precise for Sybase displays duration values for several entities.

The table below describes the duration value for each type of entity:

Table 2-3 Duration values for entities

Entity	Description
Current Session	The time that has elapsed since the session started (login time). This is not the time the session was active executing statements. For example, if a session was connected to Sybase for one hour without executing any statement, the duration will still be one hour.
Programs (Logins, Machines, DB Users, Databases, Time Units, and so on.)	The elapsed time of the sessions executing the program. The time is not measured from the login time (as it is in the Current session entity), but from the moment the Collector agent started. The average duration counter displays the average time each program has been running
Statement (batch)	The time that elapsed while the statement was executed. This is very similar to the duration of programs (see above); but since the statement is almost always active, this duration also indicates the time the statement was active. The average duration counter displays the average execution time of the statement.

About session states

Sybase provides reporting on many different types of events. Precise for Sybase groups events according to logical relationships to enable you to easily pinpoint heavy resource consumers.

For example, when the Collector agent samples the Sybase, it assigns a state for each connection. An application may wait for many events in Sybase. To simplify analysis, and help you to easily identify the source of resource consumption, Precise for Sybase groups various wait types into wait groups or states and defines the status of sessions currently connected to Sybase by their states. The state is divided into two types of states: In Sybase and Non-In Sybase. The session states are described in the following table:

Understanding In Sybase session states

The In Sybase session states signify that the session is performing activities in Sybase. The table below describes the In Sybase session states .

Table 2-4 In Sybase session states










Icon	In Sybase State	Description
	Using CPU	The session is currently executing and is not waiting for any other event. For versions lower than version 12.5.0.3, the Using CPU state is usually identified by <code>sysprocesses status = Runnable/Running</code> . For version 12.5.0.3 and higher, this state is identified by the <code>monProcess WaitEventID</code> field containing a value in <code>Wait Class 1</code> (waiting to be scheduled). Usually this state is an ideal state, but verify that the amount of CPU usage is reasonable compared to the transaction being performed. In addition monitor this state over time and verify that a growth in the CPU usage is a result of a growth in load and not a result of any change that caused the performance to degrade.





Table 2-4 In Sybase session states

Icon	In Sybase State	Description
	I/O Wait	This state indicates the session is waiting for an I/O operation to terminate. For versions lower than version 12.5.0.3, this state is usually identified by sysprocesses status = sleeping, unless the cmd field = COMMIT/ROLLBACK/LOG SUSPEND. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value in Wait Classes 2 or 3 (waiting for a disk read / write to complete).
	Lock Wait	The session is waiting for a lock to be released. The lock is held by another session. For versions lower than version 12.5.0.3, this state is identified by sysprocesses status = lock sleep. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value in Wait Class 5 (waiting to take a lock).
	Log Full Wait	The session is waiting for an operation of the log file to terminate. This state is generally encountered during a transaction log full. For versions lower than version 12.5.0.3, this state is identified by sysprocesses cmd = LOG SUSPEND. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value of 259.
	Log Wait	The session is waiting for an operation of the log file to terminate. This state is generally encountered during a commit or rollback operation. For versions lower than version 12.5.0.3, this state is identified by sysprocesses status = sleeping, and the cmd field is COMMIT or ROLLBACK. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value of 283.
	Buffer Wait	The session is waiting to access a user log cache. For versions lower than version 12.5.0.3, this state is identified by sysprocesses status = PLC sleep. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value in Wait Class 6 (waiting for memory or a buffer).
	Latch Wait	The session is waiting for an internal lock to be released. Latches are short term light weight synchronization objects. Latches are not held for the duration of a transaction. Latches are used for datapage and datarow locking, including OAM and allocation pages, but not for allpages locking. They mostly occur in SMP applications. For versions lower than version 12.5.0.3, this state is identified by sysprocesses status = latch sleep. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value in Wait Class 9 (waiting for internal system event).
	Remote Wait	The session is waiting for a remote query to terminate. For all versions, this state is identified by sysprocesses status =remote I/O.
	Network I/O Wait	The session is completing a network I/O operation. Network I/O occurs if results from a large result set are transferred to the client using shared memory netlibs or TCP/IP. For versions lower than version 12.5.0.3, this status is identified by a status of send sleep or rcv sleep. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value in Wait Classes 7 or 8 (waiting for network input/output).

Understanding Non-In Sybase session states

The table below describes the Non-In Sybase session states.

Table 2-5 Non-In Sybase session states

Icon	State	Description
	Request Wait	The session is waiting for the client to issue a request. For versions lower than version 12.5.0.3, this state is identified by sysprocesses cmd = AWAITING COMMAND. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing a value of 250.
	Other Wait	Sessions waiting in states not described here are assigned the Other wait state.
	Waitfor command	The session is executing the WAITFOR DELAY command. For versions lower than version 12.5.0.3, this state is identified by sysprocesses status = alarm sleep. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing values of 260-263.
	Parallel Sync Wait	The session is waiting for a synchronization message from another process in the family. For versions lower than version 12.5.0.3, this state is identified by sysprocesses status = sync sleep. For version 12.5.0.3 and higher, this state is identified by the monProcess WaitEventID field containing values of 152, 307 or 308.

About Statement identifiers

Each statement is identified by the following parameters:

- Text** The actual statement text.
- Database** Each statement is associated with a single database.
- User** Each statement is associated with a single user (also called the parsing user).

For example, the statement "select * from sysobjects" executed in the master and pubs3 databases is represented by two different statements in the product. The ability to roll-up the performance of those statements to one is also important. This is supplied using the collapsed statement entity.

How Precise for Sybase monitors availability

The Precise for Sybase Collector agent monitors Sybase every second and collects data regarding resource consumption. If the instance status (Up/Down) was changed, a message code is sent by the Collector agent and causes the current availability of the instance saved in the PMDB to be changed. In addition, the availability of the instance is saved, over time, in the PMDB. During a planned downtime period the instance availability is shown as it was available during this period.

NOTE If the Precise for Sybase Collector agent is down, the instance availability status reported is "Unknown."

Launching Precise for Sybase from the Precise StartPoint

Precise is a Web-based application. You can access the Precise user interface using the Internet Explorer browser, version 6.0, or later.

The syntax of the Precise URL address is: `http://<server>:<port>`, where `<server>` refers to the Precise FocalPoint and `<port>` refers to the port number used by the GUI Web server. By default, the port number is 20790. For example: `http://beanotown:20700`.

Launching Precise for Sybase from the Precise StartPoint

This URL provides secure access to the StartPoint using authorized roles. From here, you can launch all Precise products. It gives you a quick overview of the status of your applications and access to the AdminPoint, where you can perform various management tasks (For more information, see the *Precise Administration Guide*).

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

To launch Precise for Sybase using Precise StartPoint

1. Type the address of the StartPoint user interface into the Address bar of your browser and click **Enter**.
The Precise login page opens. The login page provides secure access to Precise and to your specific product.
2. Specify your authorized role name and password.
By default, both role name and password are `admin`. For more information about role names, see the *Precise Administration Guide*.
3. Click **Login**.
The StartPoint page opens. This is the Precise home page.
4. On the Product Selection bar, from the drop-down list, select the product you want to launch.

Usage concepts

This section includes the following topics:

- [About application performance tuning](#)
- [About viewing Current Activity](#)
- [About SQL statement tuning](#)
- [About instance statistics tuning](#)

About application performance tuning

Precise for Sybase allows you to proactively monitor, analyze, and tune Sybase instances. By capturing, measuring, and correlating performance metrics from all critical system components, Precise for Sybase provides a complete view of your application's performance. It detects business performance problems and helps you find suitable and comprehensive solutions thereby ensuring that your business applications always perform at peak efficiency.

Detailed information on the resource consumption of an application is extensive and includes data on CPU Consumption, I/O Waits, Internal Database Waits, and so on. Still, with Precise for Sybase you can easily identify the top resource consumers at different levels and at different times.

For example, you may want to start at the Instance level and then drill down to different Login names and the SQL statements they executed. You can view the status of current sessions, or focus your analysis on a particular time period, and examine their consumption patterns during the last minute, or further back in the past, such as yesterday, last week, or last year. You can also examine changes in resource consumption patterns over time in time slices of 15 minutes, hours, four hours, days, weeks, or months.

About viewing Current Activity

The Current tab displays near real-time information on all sessions that were connected to your Sybase instance during the last minute. Use the Current tab to solve bottlenecks, such as those caused by locked sessions or runaway processes. The Current tab displays instance activity information up to the last time slice of activity.

You can select a Sybase instance and examine its resource consumption during the last minute, such as CPU Usage, I/O Wait, and Lock Wait.

The application resource consumption is displayed on the instance level or as a summary of the session level of resource consumption.

If you identified a blocking situation that is currently in progress (for example, there is a session waiting for a lock), select the Locks view to view the current sessions that are blocking the other sessions. The information is displayed in a tree format, which helps you easily identify which are the locked sessions, who they are being locked by, and what is the lock level. You can drill down to the sessions involved in the lock to obtain more information about what is causing the lock. For example, you can drill down to a locking session to determine which SQL and which program are currently being executed.

Analyzing instance performance

While the Current tab displays near real-time data, it is recommended to begin the actual tuning process in the Activity tab.

The Activity tab allows you to analyze performance behavior over time or when the problem is reported (since in many cases the performance problem may no longer be evident from the data displayed in the Current tab). Using the Activity tab, you can analyze what happened at a specific point in time, identify problematic time periods, and drill down to analyze a bottleneck that caused a performance problem. You can begin by selecting an Instance and viewing the In Sybase overtime graph to identify problematic time periods.

In addition, long-term analysis is also required for understanding performance patterns and predicting future resource consumption. You can use the Activity tab to detect resource consumption patterns and suspicious deviations in past behavior.

The Activity tab can also be referred to when planning your system's capacity, performing period-to-period comparisons, and determining how best to optimize storage. Use the Activity tab to examine resource consumption of your Sybase Instance, over time, as well as the resource consumption of your databases, programs, logins, and so on.

You can also identify dominant resource consumption by selecting the Overview view.

The Overview view provides a breakdown of the In Sybase states, such as, CPU Usage, I/O Waits, Lock Waits, and Remote Waits. Examining this data allows you to pinpoint the type of resource causing the bottleneck. If for example you identify that Sybase is suffering from an I/O wait problem, upgrading your server to a much more powerful machine will probably fail to provide you with the expected results.

NOTE A high CPU rate exceeding 90% may indicate that indexes are not being used.

Identifying resource consumers

Precise for Sybase lets you drill down to additional levels to determine which entities are responsible for a particular performance problem. In the Association area you can observe different entities, such as databases, login names, users, programs, and client machines, to pinpoint the application component that consumes the most database resources.

For each entity (as, for example, the program entity), you can access the following information:

- Entities that are related to the entity (such as, the logins or machines that executed the program).
- Resources consumed by the entity. This enables you to determine if the entity is CPU or I/O bound, or waiting for a lock.
- Resource consumption patterns, which allows you to observe and identify trends in resource consumption. You can drill down even further to other entities associated with the selected entity in the Association area.

For example, once you have identified a problem in a specific program you can drill down to the batches or statements the program issued, to determine the source of the problem.

Associating with offending SQL statements

Having identified a major resource-consuming entity, you can further drill down to the statements and batches executed by that entity. These may include long-running SQL statements or SQL statements that use few resources but are executed frequently.

For each statement or batch, you can access the following information in Precise for Sybase:

- Entities that executed the statement or batch.
- Resources consumed by the statement or batch, which enables you to determine if the statement or batch is CPU or I/O bound, or waiting for a lock, and so on.
- Resource consumption patterns that allow you to observe if the statement or batch execution is affected by excessive use of the same resource, by another statement or batch.

Analyzing previous blocking situations

One of the main reasons that statements possess long execution times is that they are waiting for a resource that is locked by another session.

Using Precise for Sybase, you can easily identify the following:

- Statements waiting for a locked object
- Locked objects (that is, contended objects)
- Sessions causing locks

In Precise for Sybase, a statement waiting for the lock on a resource to be freed is in a Lock Wait. The session holding the lock is called a Blocker session.

About SQL statement tuning

Most problems occur in SQL statements or batches. Precise for Sybase lets you to easily identify the most resource-consuming SQL statement or batch and view a detailed explanation of the execution plan of the SQL statement or batch. It can also help you resolve the problem by recommending which index would provide a better execution plan and reduce resource consumption.

Identifying the most resource consuming statements or batches

Use the Activity tab to identify the most resource-consuming SQL statements or batches over a selected time period. By focusing in on a specific time period you can apply your efforts to tuning a problematic time frame, such as a batch window or end-of-the-month processing.

In the Association area of the Activity tab, you can associate to a specific database entity and then from there to its associated SQL statements.

When selecting an SQL statement, you can view a breakdown of its resource consumption over time, its text, and all locking contentions that it participated in (time spent waiting for a resource).

Collapsing constants into bind variables

When an application uses constant values and literals instead of bind variables (parameter markers), many different SQL statements actually constitute the same SQL statement with different values.

When each of these similar SQL statements is treated as a different statement, it is hard to summarize their resource consumption to identify the best candidate for SQL tuning.

Precise for Sybase provides a unique algorithm that ignores constant values and groups all similar statements into collapsed statements. This lets you identify the most resource-consuming collapsed statement.

The performance data for the collapsed statement is a summary of all performance data for all statements in the group. Most of the time, all the statements that were grouped together into one collapsed statement have the same execution plan. However, different values in the WHERE clause may lead the optimizer to choose a different execution plan.

Precise for Sybase enables you to select a Collapsed statement and associate to all the different execution plans of the SQL statements that were grouped together.

About viewing an execution plan

To view the execution plan of the statement and understand how the optimizer chooses to perform the statement, launch to the SQL tab. Each branch in the execution plan represents one operation. Each operation is displayed using the following format:

```
[<execution order>] <operationtype> (<options>) <accessed object>[<cost in percentage>]
```

When examining the execution plan, always start by exploring it from right to left and top to bottom (according to the execution order in a branch on the tree). The first operation to be performed has execution order 1, followed by 2, and so on.

Understanding an execution plan

By observing the access plan of a selected statement or batch, you can analyze the access path chosen by the Sybase Optimizer. Click the Statistics tab to compare the costs of the execution plan operators and quickly locate the most resource-consuming operation in the execution plan.

Clicking the icon on the left-most column in the Statistics tab will highlight the corresponding row in the execution plan tree.

The information displayed in the Operations tab helps you understand the nature of the execution plan, its major operations, and the major access types to objects (index scan, index seek, table scan, and so on).

Analyzing the statement execution plan over time

Precise for Sybase allows you to track Access Plan changes and compare them with the resource consumption of the statement to understand how the changes in the Access Plan affected its performance and execution plan.

Examining object space usage over time

Another way to tune a database object is to examine its space usage. The Collect Space Utilization process collects table and index space information daily (by default).

Space usage information allows you to examine space allocated versus space used by a table or an index. This enables you to easily keep track of object growth over time and decide how to allocate memory and connect objects to database devices in the specified database.

About instance statistics tuning

The Precise for Sybase Collector Instance Statistics process collects a wide range of performance counters from the operating system, some which belong to Sybase and some to the operating system (OS) itself. These performance counters are sampled every minute and displayed in the Current and Statistics tab. When loaded into the PMDB, the counters are aggregated and kept at a slice level. The counters are then aggregated to daily, weekly, and monthly levels.

Observing instance statistics over time

The Statistics tab provides many performance counters grouped into several predefined overtime graphs that enable you to locate various aspects of performance problems, such as CPU, paging, I/O, and network. The statistics tab can be used to monitor your system's current state as well as previous states. It is also possible to monitor a specific counter associated with a particular performance group by selecting the required counter.

About getting an overview of your application

This section includes the following topics:

- [About the Dashboard tab](#)
- [How the Dashboard tab is structured](#)
- [How the Dashboard tab can help you identify performance problems](#)

About the Dashboard tab

The Dashboard tab provides a preview of the information displayed in the other Precise for Sybase tabs. The Dashboard is the starting point for your in-depth analysis since it displays an overview of all the instances available in the selected application and information that is also available in some of the other tabs.

The Dashboard tab displays statistics and resource consumption of all Sybase instances during the selected time frame. This tab can be used to monitor your system's current state by viewing all the available instances, a list of all active sessions and resource consumption for the last minute (Current tab). In addition you can view statistics over time, resource consumption over time and top resource consumers during the selected time frame (in the Overall, Statistics, and Activities tabs).

The Dashboard tab provides answers to the following types of questions: "Which instance is consuming the most resources?" or "Which program is consuming the most resources for the specified instance and across all instances?"

When you select the All row, in the All Sybase Instances table, located in the Main area (left pane), a summary of the activities of all the instances and a comparison of application identifiers is displayed in the various tab pages, viewed in the Instance Details area (right pane).

In general, launching a tab from the All row, focuses on the Tier entity, which summarizes the data collected for all instances. Tabs that do not have an Tier entity, open on the default instance.

See [About Precise for Sybase tabs](#) on page 15. See [About the Current tab](#) on page 40.

See [About Precise for Sybase tabs](#) on page 15. See [About the Activity tab](#) on page 49.

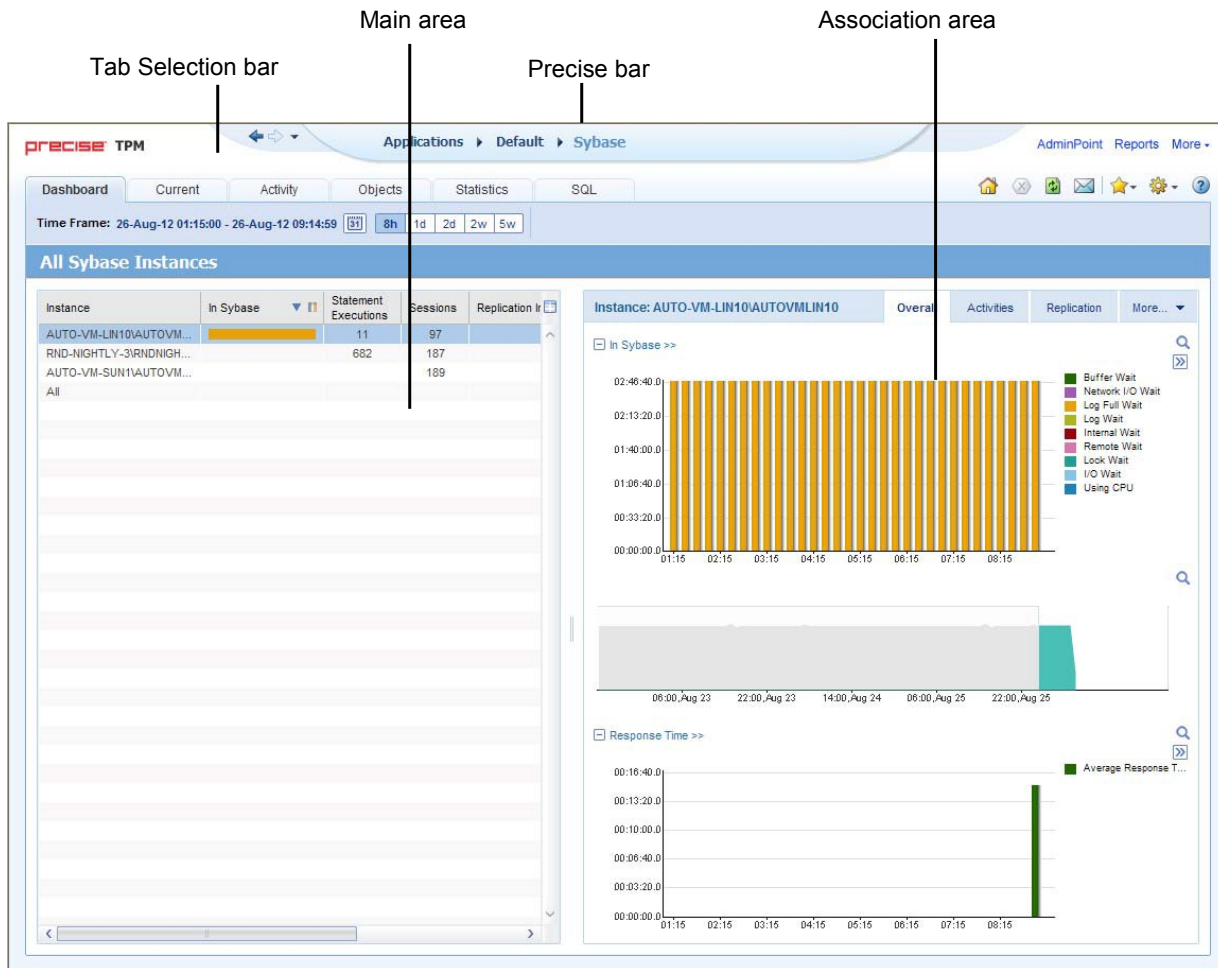
See [About the Statistics tab](#) on page 81. See [About the SQL tab](#) on page 97.

How the Dashboard tab is structured

The Dashboard tab is divided into two areas - the Main area and the Instance Details area. The Main area lists all the instances that are monitored by the Precise for Sybase Tier. The Instance Details area provides comparative information regarding the selected instance.

The figure below shows the structure of the Dashboard tab.

Figure 4-1 Dashboard tab structure



About the Main area in the Dashboard tab

The Instance table displayed in the Main area lists all the instances that are monitored by the Precise for Sybase Tier. Each row corresponds to an instance. The All row enables you to view information for the entire Tier.

The table below shows the information displayed in the Instance table.

Table 4-1 Instance table

Column	Description
Instance	Name of the instance.
In Sybase	Breakdown of In Sybase duration of the instance during the selected time frame.
Statement Executions	Total statement executions during the selected time frame.
Sessions	Total number of sessions during the selected time frame.
Replication in Sybase	Total in Sybase resulting from the replication program in the selected time frame.
RepAgent Cmds	Number of Commands sent to replication by the selected ASE in the selected time frame.
Outbound Backlog	Number of commands delayed in the outbound processing phase in the selected time frame. This column is calculated as follows: sum (SQMR Cmds Read - SQM Cmds Written).

About the Instance Details area in the Dashboard tab

This area displays overall activity, active sessions and statistical information for the selected instance. Clicking on the hyperlink for specific data items launches, in-context, the related tab, from which the data was derived.

NOTE If you select the All row, the Current tab will not be available.

Gaining an overview of your system's performance

The Overall tab displays overall activity of the selected instance during the selected time frame.

All items in this display launch the Activity tab with an instance or Precise for Sybase Tier context. The table below describes the information displayed in the Overall tab.

Table 4-2 General instance data

General instance data	Description
In Sybase	Displays a breakdown of the total In Sybase duration over time, for the selected instance during the selected time frame, into states such as, Using CPU, I/O Wait, Lock Wait, and so on.
Load	Displays the total number of sessions over time, for the selected instance during the selected time frame.
Response time	Displays the average response time over time, for the selected instance during the selected time frame.

Observing top resource-consuming databases, programs and batches

The Activities tab displays the top resource-consuming databases, programs and batches, during the selected time frame.

All entities in this display launch the Activity tab, in-context of an instance or Tier. The table below describes the information displayed in the Databases table.

Table 4-3 Databases table

Column	Description
Database	Displays the names of the top resource-consuming databases.
Sessions	Displays the total number of sessions opened, during the selected time frame.
In Sybase	Displays the resource consumption of the database broken down into states such as, Using CPU, I/O Wait, and Lock Wait states. Click the respective icons to display the data in this column in a stacked graph or numerical format.

The table below describes the information displayed in the Programs table.

Table 4-4 Programs table

Column	Description
Program	Displays the names of the top resource-consuming programs.
Sessions	Displays the total number of sessions opened, during the selected time frame.
In Sybase	Displays the resource consumption of the program broken down into states such as, Using CPU, I/O Wait, and Lock Wait states. Click the respective icons to display the data in this column in a stacked graph or numerical format.

The table below describes the information displayed in the Batches table.

Table 4-5 CMDBatches table

Column	Description
Batch	Displays the IDs of the top resource-consuming batches.
Duration (Avg)	Displays the average duration of the batch, during the selected time frame.
In Sybase	Displays the resource consumption of the database broken down into states such as, Using CPU, I/O Wait, and Lock Wait states. Click the respective icons to display the data in this column in a stacked graph or numerical format.

Examining Replication for an instance

The Replication tab provides predefined graphs displaying the activity results from replication in the selected ASE, the outbound processing performed for the selected ASE, and the commands sent for replication by the selected ASE.

The table below describes the information displayed in the Replication tab.

Table 4-6 Replication graphs

Statistical data	Description
Total In_Sybase vs. Replication activity (outbound)	Displays the total In_Sybase of the instance versus the In_Sybase resulting from replication activity in the selected time frame. The In_Sybase resulting from replication activity is displayed as a stack bar which represents the various states. This graph can give an indication of how much the replication activity is taking out of the entire instance activity. In addition, you can see a change in the activity trend. You can click on the graph title to launch to the Activity tab to see more details regarding the activity resulting from replication.
Outbound Load	Displays counters related to outbound processing that can indicate latency in one of the outbound components during the selected time frame. The following counters are displayed in this graph: <ul style="list-style-type: none"> • SQM Cmds Written • SQMR Cmds Read • DSIE Cmds Read • DSIE Cmds Applied Clicking the graph title launches the Statistics tab, allowing further investigation.
Inbound Load	Displays the number of commands sent by the selected ASE for replication during the selected time frame. <ul style="list-style-type: none"> • RepAgent Cmds Sent Clicking the graph title launches the Statistics tab, allowing further investigation.

Examining statistics for an instance

The Statistics tab displays statistical information for the selected instance, as reported by Sybase and collected by the Precise for Sybase Collect Instance Statistics Process during the selected time frame.

All items in this display launch the Statistics tab with an Instance or Precise for Sybase Tier context. The table below describes the information displayed in the Statistics tab.

Table 4-7 Instance statistics

Statistical data	Description
Engine utilization	Displays the Engine utilization for the selected instance during the selected time frame.
Committed transactions	Displays the number of committed transactions during the selected time frame.
Connections opened	Displays number of connections opened during the selected time frame.

Observing top resource-consuming entities

The Activators tab displays the top resource-consuming machines, logins and users during the selected time frame. All entities in this display launch the Activity tab, in-context of the instance.

The table below describes the information displayed in the Machines table.

Table 4-8 Machines table

Column	Description
Machine	The names of the top resource-consuming machines.
Sessions	Total number of sessions opened during the selected time frame.
In Sybase	The resource consumption of the machine broken down into states such as, Using CPU, I/O Wait, and Lock Wait states. Click the respective icons to display the data in this column in a stacked graph or numerical format.

The table below describes the information displayed in the Logins table.

Table 4-9 Logins table

Column	Description
Login	The names of the top resource-consuming logins.
Sessions	Total number of sessions opened during the selected time frame.
In Sybase	The resource consumption of the login, broken down into states such as, Using CPU, I/O Wait, and Lock Wait states. Click the respective icons to display the data in this column in a stacked graph or numerical format.

The table below describes the information displayed in the Users table.

Table 4-10 Users table

Column	Description
User	The names of the top resource-consuming users.
Sessions	Total number of sessions opened during the selected time frame.
In Sybase	The resource consumption of the user, broken down into states such as, Using CPU, I/O Wait, and Lock Wait states. Click the respective icons to display the data in this column in a stacked graph or numerical format

Observing your system's current activity

The Current tab displays the overall activity and active sessions that occurred within the last minute for the selected instance. All items in this display launch the Current tab, in the context of an instance. The table below describes the information displayed in the Sessions Activity area.

Table 4-11 Sessions activity

Item	Description
Sessions	Total number of sessions connected to the selected instance.
Active sessions	Total number of sessions currently active.
Locked sessions	Total number of locked sessions.

The table below describes the information displayed in the Current Activity graphs

Table 4-12 Current activity graphs

Item	Description
Overall Activity (Last minute)	Provides a graphical display of the overall activity that occurred within the last minute, for the selected instance, broken down into the following states: Waitfor command, Synchronized Wait, Other Wait, Request Wait and In Sybase.
In Sybase (Last minute)	Provides a graphical display of In Sybase activity that occurred within the last minute, for the selected instance, broken down into states, such as, Using CPU, I/O Wait, and Lock Wait.

The table below displays a list of all the active sessions that occurred within the last minute, for the selected instance.

Table 4-13 Active sessions (Last minute)

Item	Description
Icon	Represents the current session's state (such as, Lock Wait, Request Wait).
SPID	The session ID.
Program	The name of the program that opened the session.
State In Sybase	The current state of the session. Total In Sybase entities for the session, occurring within the last minute, broken down into states (such as, Using CPU, I/O Wait, and Lock Wait).

How the Dashboard tab can help you identify performance problems

You can identify a performance problem by doing one or more of the following:

- [Determining which instance consumes the most resources](#)
- [Identifying the highest resource consumers in various entities](#)
- [Examining instance statistics](#)
- [Monitoring a session activity during the last minute](#)

Determining which instance consumes the most resources

Precise for Sybase enables you to determine which are the most resource-consuming instances, indicating a potential performance problem.

To determine which instance consumes the most resources

1. In the All Sybase Instances table, observe the total In Sybase statistics and statement executions for an instance during the selected time frame.
2. View its resource consumption over time to identify problematic time periods and get an idea of the general trend of resource consumption.

Identifying the highest resource consumers in various entities

Precise for Sybase lets you determine which are the most resource-consuming application components (such as, Machines, Logins, Users, Databases, Programs, and Batches).

To identify the highest resource consumers in various entities

1. Identify the most resource-consuming instances.
2. Select the Activators and Activities tabs in the Instance Details area to determine which are the most resource-consuming application components.
3. Launch the Activity tab to continue the performance tuning process on the selected component.

Examining instance statistics

You can examine instance statistics over time to confirm that Engine utilization and Load (Committed Transactions and Connections Opened) are within normal ranges. A continually high Engine utilization may indicate the need to improve the application or to add more engines to the machine to enhance performance.

To examine instance statistics

1. Select the Statistics tab in the Instance Details to view Engine utilization during the selected time frame.
2. Launch the Statistics tab to observe additional instance statistics that can help you locate performance problems in various other aspects of your system, such as I/O and Network.

Monitoring a session activity during the last minute

Monitoring a session's activity during the last minute in Sybase, can help you locate anomalous activities such as Locking, Internal Waits, and so on.

To monitor a session's activity during the last minute

1. Select the Current tab in the Instance Details area to view all the active sessions and an In Sybase breakdown into states, such as, Using CPU, I/O Wait, and Lock Wait.
2. Launch the Current tab to view additional information on the current sessions.

Getting current application information

This section includes the following topics:

- [About the Current tab](#)
- [How the Current tab is structured](#)
- [About the entities you can examine in the Current tab](#)
- [How the Current tab can help you identify performance problems](#)

About the Current tab

The Current tab displays near real-time information on all sessions that were connected to your Sybase database during the last minute. If you have recognized an immediate performance issue and are looking for an indication as to where to begin your analysis, start with the Current tab. It is also worthwhile to start with the Current tab if your system appears to hang and there is no throughput.

You can use this tab to monitor your system's current activities and resource consumption and to identify performance problems, such as current locking problems and sessions involved in locks, as they occur.

The Current tab provides answers to the following types of questions: "How many sessions are currently running?" or "Why doesn't my application work when I click OK?" The Current tab also shows you which sessions are currently active, and which machines are connected to which users.

Using the Current tab, you can monitor current connections to your Sybase instance, check the state of each connection, and identify the resources currently utilized by each session. Further investigation of each session provides a thorough understanding of the SQL statement being executed, the login name executing the statement, the program, and the resource utilization for the session. In addition, you can view the statements and batches executed by the selected session within the last time slice.

In the Current tab, data on a current session is continuously updated to provide information on what happened during the last minute or time slice. Because the displayed session information is derived from a dynamic application, it is recommended to regularly refresh the screen, either manually or automatically, to ensure that the information displayed is up-to-date.

NOTE The Current tab displays information on the last minute of activity. For example, if you have worked at your screen for ten minutes and then refresh the information displayed on your screen, you will be able to view new information that occurred during the last minute of activity.

See [About the Dashboard tab](#) on page 33.

How the Current tab is structured

The Current tab displays information on a selected entity and its associated entities. For example, if the parent entity is a Session, its associated entities would be Statements or Batches.

When you open the Current tab from the Dashboard tab, the selected entity is by default Instance, meaning that information is displayed on the instance level. When you open the Current tab from the Activity tab, the tab is launched with the selected instance or entity (in the case of an entity, this is only valid for the database entity), meaning that the information is displayed on the instance or entity level. If you open the Current tab from another tab, the historical settings (meaning those settings which were selected when you left the tab, such as a drill-down to a Locked Objects entity) are taken into account and the information displayed the last time you viewed this tab is displayed (similar to clicking the History button and returning to a previous tab).

The selected entity is always reflected in the Tab heading, which serves as a point of orientation. The highest-level entity you can view information for in the Current tab is Instance. You can view another instance by selecting it from the Instance list.

About the Main area in the Current tab

The Main area shows comprehensive information on the selected entity, such as the In Sybase data for the last minute, and the number of active sessions.

About the Association area in the Current tab

The Association area provides corresponding information on the entities associated with the selected entity (displayed in the Main area). For example, you can associate to the Active Sessions related to the selected instance. The selection you make is reflected in the Association area only; the Main area remains unchanged. For some entities, tabs above the Association area enable you to view additional information. For example, the tabs displayed for a database are: General, In Sybase and Statistics. Clicking a tab displays different table columns with the same type of child entities.

From the Association area, you can also drill down to another entity by clicking a table row. A drilldown affects the entire tab. When you drill down to another entity, the Tab heading displays the new selection, the Main area displays information on the newly selected entity, and the Association area displays the entities associated with the selected entity.

For example, when you want to view information on a specific session, in the Current tab, choose Sessions from the Association controls. The Association area changes to display sessions-related information. Note that the Tab heading and the Main area remain unchanged.

In the Association area, click the row of the session that you want to view detailed information for. The Tab heading indicates the newly selected entity; the Main area displays detailed information on the session you drilled down to, and the Association area shows information on the statements associated with this session.

See [About Precise for Sybase tabs](#) on page 15. See [About the Dashboard tab](#) on page 33.

About the entities you can examine in the Current tab

The Current tab displays information on different entities. This section provides an overview of all entities, their meaning, and their views.

The following entities can be examined in the Current tab:

- Instance entities
- Database entities
- Session entities
- Statement and Batch entities

About the Instance entity

The Instance entity displays information on the resource consumption of the entire instance and allows you to focus on how it is performing, during the last minute.

About getting an overview of current Instance entities

The Overview displays information on the activity of the instance, such as number of sessions and the resource consumption breakdown of the instance.

The table below describes the information displayed in the Instance overview.

Table 5-1 Instance overview

View area	Description
Sessions	<p>Displays the following session-related information:</p> <ul style="list-style-type: none"> Sessions - indicates the number of sessions (different SPIDs) that were connected to the Sybase instance during the last sample. Active Sessions - number of sessions (different SPIDs) connected to the Sybase instance during the last sample that were not in a state that is not part of the In Sybase group.
Statistics	<p>Displays information on the following statistical parameters:</p> <ul style="list-style-type: none"> CPU Time - cumulative CPU time for all sessions connected to the instance during the last minute. Physical I/O - the total number of physical I/O requests performed by all sessions connected to the instance during the last minute.
In Sybase (Last minute)	<p>Displays the resource consumption breakdown of the instance during the last minute (that is, 60 samples of the Collector agent). In addition to the resource distribution, the view displays the number of sessions in each state during the last sample.</p> <p>It is important to remember that since this view only displays last minute's activities, you will always be viewing the activities that occurred during the last minute and not the activities that occurred since the last refresh action.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> State - displays the In Sybase state. In Sybase - graphical representation of the Time column. Time - amount of time the selected entity was in this state. % - percentage of time the selected entity was in this state. Sessions (Last Sample) - current number of sessions in the specified state. For example, there are currently three sessions experiencing a Lock Wait state. <p>NOTE: All the information displayed in this table represents information collected during the last minute, with the exception of the Sessions column, which represents the number of sessions, in the last sample, that were in a particular state.</p>

See [About session states](#) on page 25.

About the Database entity

The Database entity displays information on the activity and resource consumption of the database, including the number of sessions connected to the database during the last sample taken, various session-related statistics and the database resource consumption breakdown for the last minute.

About getting an overview of current database activities

The Overview displays information on the activity and resource consumption of the database, including the number of sessions connected to the database during the last sample taken, various session-related statistics and the database resource consumption breakdown for the last minute.

The table below describes the information displayed in the Database overview.

Table 5-2 Database entity overview

View Area	Description
Session	<p>Displays the following session-related information:</p> <ul style="list-style-type: none"> Sessions - number of sessions (different SPIDs) that were connected to the database during the last sample. Active Sessions - number of sessions (different SPIDs) connected to the database during the last sample that were not in Request Wait state.
Statistics	<p>Displays information on the following statistical parameters:</p> <ul style="list-style-type: none"> CPU Time - cumulative CPU time for all sessions connected to the database during the last minute. Physical I/O - total number of physical I/O requests performed by all sessions connected to the database during the last minute.
In Sybase (Last minute)	<p>Displays the resource consumption breakdown of the database during the last minute (that is, 60 samples of the Collector agent). In addition to the resource distribution, the view displays the number of sessions in each state during the last sample.</p> <p>It is important to remember that since this view only displays the last minute's activities, you will always be viewing the activities that occurred during the last minute and not the activities that occurred since the last refresh action.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> State - displays the In Sybase state. In Sybase - graphical representation of the Time column. Time - amount of time the selected entity was in this state. % - percentage of time the selected entity was in this state. Sessions (Last Sample) - current number of sessions in the specified state. For example, there are currently three sessions experiencing a Lock Wait state. <p>NOTE: All the information displayed in this table represents information collected during the last minute, with the exception of the Sessions column, which represents the number of sessions in the last sample that were in a particular state.</p>

See [About session states](#) on page 25.

About the Session entity

The Session entity displays information on a single connection with Sybase. Precise for Sybase identifies a session by the following attributes:

- System Process ID (SPID) - a unique integer assigned to each user connection when the connection is made.
- Host Process (PID)

When one of these attributes changes, a new session is recorded and displayed.

About getting an overview of a session entity

The Overview displays information on the session identifier's connection details, session-related statistics and a breakdown of session resource consumption, during the last minute.

About the entities you can examine in the Current tab

The table below describes the information displayed in the Session overview.

Table 5-3 Session overview

View Area	Description
Session tab	<p>Displays information on the following session identifiers:</p> <ul style="list-style-type: none"> • State - current state of the session. • In Sybase - graphical representation of the Time column. Text -Text of the statement which is currently running. • Status - current status of the session as reported by Sybase. The Session status can be one of the following: <ul style="list-style-type: none"> • Alarm sleep - waiting for alarm to wake process up. • Background - a process, such as a threshold procedure, run by Adaptive Server rather than by a user process. • Infected - server has detected a serious error condition; extremely rare. • Latch sleep - waiting on a latch acquisition. • Lock sleep - waiting on a lock acquisition. • PLC sleep - waiting to access a user log cache. • Recv sleep - waiting on a network read. • Remote I/O - performing I/O with a remote server. • Runnable - in the queue of runnable processes. • Running - actively running on one of the server engines. • Send sleep - waiting on a network send. • Sleeping - waiting on a disk I/O, or some other resource (often indicates a process that is running, but doing extensive disk I/O). • Stopped - stopped process. • Sync sleep - waiting on a synchronization message from another process in the family. • Last Batch Time - If the session is currently executing a batch, this field displays the time the batch started; otherwise it displays the time the previous batch finished executing. Idle Time -Time elapsed since last batch. • Parallel Sessions - Current number of sub-threads used to execute a batch in parallel. This counter is calculated from the fid column in the sysprocesses. It enables you to verify that Sybase uses the best execution plan for the current statement. • Blocked By - If the session is blocked by another session, displays the SPID of the blocker session.
Connection tab	<p>A connection includes session identifiers, such as database, login, machine and user, as reported by the sysprocesses.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> • Session status - Current status of the session as reported by Sybase. • Last Batch Time - If the session is currently executing a batch, this field displays the time the batch started; otherwise it displays the time the previous batch finished executing. • Idle Time -Time elapsed since last batch.
In Sybase (last minute)	<p>Displays the resource consumption breakdown of the session during the last minute (that is, 60 samples of the Collector agent).</p> <p>It is important to remember that since this view only displays the last minute's activities, you will always be viewing the activities that occurred during the last minute and not the activities that occurred since the last refresh action.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> • In Sybase - graphical representation of the Time column. • Time - amount of time the selected entity was in this state. • % - percentage of time the selected entity was in this state.

About the entities you can examine in the Current tab

View Area	Description
Locks tab	Displays information on the current session locks that exists in the instance. Contains the following information on each session: <ul style="list-style-type: none"> Blocked by - If the session is blocked by another session, displays the SPID of the blocker session. Blocker user/program Locking chain - for more than one locked session. Lock wait - lock duration.
Statistics tab	Displays information on the following statistical parameters: <ul style="list-style-type: none"> CPU Time - CPU time for the sessions during the last minute. Physical I/O Operations - total number of physical I/O requests performed by the session during the last minute. Memory Pages Used - number of pages in the procedure cache that are currently allocated to the SPID.

See [About session states](#) on page 25.

About viewing the current statement's text

The Text view displays the text of the current statement. If the session is not active and executing a statement, a message notifying you that the session is not active and that no text is available, is displayed.

About Statement and Batch entities

The Statement and Batch entities display information on a SQL statement that was executed by the selected session during the last time slice.

About getting an overview of current Statement and Batch entities

Displays general information on the statement or batch, including overall statement properties, the breakdown of the statement in Sybase during the last time slice, and some statistics regarding the resource consumption of the statement.

When the selected entity is a Batch, the information area displays the entire text of the improvised batch or stored procedure. When the selected entity is a Statement, the information area only displays the part of the batch text that forms the statement.

The table below describes the information displayed in the Statement or Batch overview.

Table 5-4 Statement or Batch overview

Parameter	Description
Statement/Batch ID	Internal statement hash value (SHV) calculated by the Collector agent.
Database	Identifies the database on which the statement or batch is run. This information is used by Precise for Sybase to display the statement's execution plan.
User	Displays the parsing user who runs the statement or batch. This information is used by Precise for Sybase to display the statement's execution plan.
Executions	Number of times the Collector agent identifies that the statement or batch was executed.
Duration	Total time required to execute the selected statement or batch during the last time slice.
Duration (Avg)	Average duration of each statement execution during the last time slice.
CPU Time	Cumulative CPU time for the statement during the last time slice.

Table 5-4 Statement or Batch overview

Parameter	Description
Physical I/O Operations	Total number of physical I/O requests during the execution of the statement or batch during the last time slice.
Parallel Sessions (Min)	Minimum number of threads used to execute the statement or batch in parallel. This counter is calculated from the fid column in sysprocesses. This enables you to verify that Sybase is using the best execution plan for the current statement.
Parallel Sessions (Max)	Maximum number of threads used to execute the statement or batch in parallel. This counter is calculated from the fid column in sysprocesses. This enables you to verify that Sybase is using the best execution plan for the current statement.
In Sybase (Last minute)	<p>Displays the resource consumption breakdown of the statement or batch during the last minute (that is, 60 samples of the Collector agent). In addition to the resource distribution, the view displays the number of sessions in each state during the last sample.</p> <p>It is important to remember that since this view only displays the last minute's activities, you will always be viewing the activities that occurred during the last minute and not the activities that occurred since the last refresh action.</p> <p>The following information is displayed:</p> <ul style="list-style-type: none"> • State - displays the In Sybase state. • In Sybase - graphical representation of the Time column. • Time - amount of time the selected entity was in this state. • % - percentage of time the selected entity was in this state.

See [How information is collected](#) on page 25.

About viewing current Statement or Batch text

Displays text of the selected statement or batch.

About the Statement and Batch Association area

It is possible to view statement or batches in the Association area while drilling down from a single session. The table includes only statements or batches that were executed during the last time frame. A time frame is a 15-minute interval. If a statement or batch was executed in an earlier time frame, this statement or batch no longer appears in the Current tab. Instead, you can view information on this statement or batch in the Activity tab.

The table below describes the information displayed in the Association area, when viewing a statement or batch entity.

Table 5-5 Statement/Batch Association table

Association Tab	Function
General	Shows general information on the displayed statements or batches.
In Sybase	Shows information on performance and statistics of the displayed statements or batches.
Statistics	Shows statistics of the displayed statements or batches.

See [About session states](#) on page 25.

See [Understanding In Sybase session states](#) on page 25.

See [Understanding Non-In Sybase session states](#) on page 27.

How the Current tab can help you identify performance problems

You can identify a performance problem by doing one or more of the following:

- See [Examining the resource consumption of an entire instance and database](#) on page 47.
- See [Observing session activity](#) on page 47.
- See [Examining a single connection](#) on page 47.

Examining the resource consumption of an entire instance and database

The Current tab provides an almost real-time picture of the performance and behavior of your Sybase instance. It displays information on the last minute of activity.

To examine the resource consumption of an entire instance and database

1. Examine the activities of the entire instance level to determine what is happening in your system.
2. Examine the current activity at an instance level of display to see how the system is currently behaving and determine if there is an overall problem that affects all sessions or if there is a problem with specific sessions.
3. Start your investigation in the Overview view of an Instance or Database entity. This will show you a breakdown of resource consumption for the last minute.

Observing session activity

In the Association area you can see which sessions are currently connected to the Instance, or view which sessions are currently using a specific database or holding a specific object.

To observe session activity

1. Observe the current state of the session and determine if it is active or waiting for a resource.
2. Identify the session by different identifiers such as its login, the machine it is running on, or the programs it is running. The current active SQL code is displayed for active sessions.
3. If you have hundreds of sessions running, and you want to focus your investigation on the sessions that were active in Sybase during the last minute, select Active Sessions from the Association Controls in the Association area.

Examining a single connection

You may need to explore a single running SPID, to obtain more information for it. To examine a single connection

1. Drill down to the specific SPID in question. Continue examining its resource breakdown, its recent statements, or the batches in which the session was executed, during the last few minutes.
2. Check if the session is blocked and by whom.

Examining performance over time

This section includes the following topics:

- [About the Activity tab](#)
- [How the Activity tab is structured](#)
- [About the entities you can examine in the Activity tab](#)
- [About Statement, Batch, Collapsed Statement, and Collapsed Batch entities](#)
- [How the Activity tab can help you identify performance problems](#)

About the Activity tab

The Activity tab allows you to analyze performance behavior over time or within a given time period. When investigating a performance issue, the Activity tab is an excellent place to start. You can easily identify the performance profile of your Sybase instance and determine who is executing which application, what they are waiting for and how many resources they are consuming. This is the critical information you need to identify the top consumers of your Sybase Instance and its major latencies. By focusing on the top consumers and latencies (wait states) you can maximize the tuning impact on your system, while minimizing the time it takes to obtain results.

The Activity tab is also very useful for a scale-up exercise. Too frequently companies add CPUs to an existing server in their attempt to improve performance. Although occasionally successful, step one of any scale-up activity should be to identify the major latencies in the Sybase Instance and resolve them. Only once the major wait states are resolved, will a scale-up actually be capable of improving performance. Otherwise, you may be simply wasting money and resources.

The information displayed in the Activity tab is derived from historical data and enables you to access information for different time intervals at varying levels of detail. You can analyze what happened at a specific time, pinpoint problematic time periods, and drill down to analyze a bottleneck that caused a performance problem.

The Activity tab was designed to answer questions such as, "What happened yesterday, as well as a month ago and a year ago?" and "What has been the performance level of my applications over time".

How the Activity tab is structured

The Activity tab displays information on various entities and their associated entities. The information displayed when you launch the Activity tab depends upon where you are launching from.

For example, when you launch the Activity tab from the Dashboard tab, one of the following occurs, depending upon what you selected in the All Sybase Instances table:

If you selected	the Focus in Activity tab is on
an instance row	the selected instance
the All row	the Tier level

When you access the Activity tab from other tabs, the entity, time frame, views, and associations that were in place the last time you visited the tab, are displayed (similar to clicking the Back button in the Precise toolbar).

When you launch another tab, from within the Activity tab, the entity time frame you currently selected is taken into account. The tab you launched will automatically display information in reference to the entity and time frame (if the entity exists in the tab you launched). This is called in-context navigation and enables you to focus on other aspects of the selected entity to further your analysis.

The information on the selected entity and its associated entities are displayed in two areas of the tab - the Main area and the Association area.

About the entities you can examine in the Activity tab

When you perform an action on an element in the lower area (the Association area) the information displayed in the upper area (the Main area) will change to reflect the action you requested. Accordingly, you can control the tab by:

- Keeping the entity in the Main area, but selecting a different set of entities to view in the Association area, by selecting an entity from the Association controls list.
- Switching to a different entity by drilling to another entity from the Association area.
- Switching to a different view in the Main area, thereby displaying additional information on the selected entity.

About the Main area in the Activity tab

The Main area shows comprehensive information on the selected entity. You can choose from several views to examine the entity from different angles. For example, you can focus on an In Sybase over time graph for an Instance entity, to observe its resource consumption over time or you can focus on the Scalability view to investigate the scalability of the instance over time.

See [About drilling down in context](#) on page 18.

See [About the Main area](#) on page 17.

See [About the Association area](#) on page 17.

Navigating in the Activity tab

The name of the entity you selected is displayed in the Tab heading, which serves as a point of orientation. The highest level entity that you can view information for in the Activity tab is the Tier.

Moving the pointer over the Tab heading displays a ToolTip with navigational information that traces the path you have already taken in your investigation

See [About drilling down in context](#) on page 18. See [About the Main area](#) on page 18. See [About the Association area](#) on page 18.

About the Association area in the Activity tab

The Association area displays relevant information on the entities associated with the selected entity (displayed in the Tab heading in the Main area) in a table format. For example, it is possible to associate to Databases, Statements and Logins that are related to a specific Instance, by selecting an entity from the Association Controls list. Notice that some entities, show additional information on the associated entities in different tabs. The tabs are located above the Association area table. Clicking on a tab displays different table columns showing different information for the associated entities. For example, the following tabs show additional information for the Statement entity:

- Performance
- In Sybase
- Statistics
- Plan

If you want to view additional information on an associated entity, drill down to it, by selecting the entity's row. A drilldown affects the entire tab. When you drill down to another entity, the Tab heading reflects the new selection, the Main area displays information on the selected entity, and the Association area displays the entities that are now associated with the newly selected entity.

For example, when you want to drill down to an associated entity, such as a specific program, to view additional information, choose Programs, from the Association controls. The Association area changes to display a list of programs. Note that the Tab heading and the Main area remain unchanged.

In the Association area, click the row of the program you want to view detailed information for. The Tab heading indicates the newly selected program; the Main area displays In Sybase information on the program you drilled down to, and the Association area shows information on statements executed by this program.

See [About drilling down in context](#) on page 18. See [About the Main area](#) on page 18. See [About the Association area](#) on page 18. See [How most tabs are structured](#) on page 15. See [Tasks common to most tabs](#) on page 20. See [About the Dashboard tab](#) on page 33.

About the entities you can examine in the Activity tab

The Activity tab displays information on different entities. This section provides an overview of all entities, their meaning, and their views. By default, all entities open in Sybase view, with the exception of the Statement, Batch, Collapsed Statement and Collapsed Batch entities, which open by default in the Overview view.

The following entities can be examined in the Activity tab:

- All Instances
- Instance and Collector entities
- Statement/Batch/Collapsed Statement/Collapsed Batch Entities
- Session and Blocker Session Entities

NOTE The entities listed above have been grouped together logically since they can be described similarly. They appear separately when viewed in the GUI.

About the All Instances entity

The Tier entity displays the resource consumption of all the instances in the application. The following views are available:

- Overview
- In Sybase
- Instance grouping

About viewing a summary of the performance and resource consumption of all instances

The Overview displays a summary of the performance and resource consumption of all instances in the application, as follows:

Instances	Displays the number of Instances defined in the Tier.
In Sybase breakdown (Table)	Displays the resource consumption breakdown of all the instances during the selected time frame. The table includes the following information: <ul style="list-style-type: none">• State - In Sybase state.• In Sybase - displays a graphical representation of the Time column.• Time - amount of time the selected Tier was in this state.• % - percentage of time the selected entity was in this state <p>NOTE: The table only displays the states the selected entity was in. For example, if the selected entity was not waiting for an I/O Wait during the selected time frame, the table will not display I/O Wait parameters.</p>

About viewing In Sybase data for all instances

Displays the overtime graph of the total In Sybase counter for all instances in the Sybase Tier.

About viewing Instance grouping for all instances

Displays instance grouping information for the selected Sybase Tier. It is possible to categorize instances into groups and observe their tuning counters, such as their In Sybase breakdown, according to the groups you defined. This option is useful if there are many instances in the same application. For an explanation on how to define instance groups, see the *Precise Administration Guide*.

About Instance and Collector entities

The Instance entity and the common Collector entities display the resource consumption of the entire instance (or the entity selected). When focusing on an instance, or collector entity, such as a program, you can examine how it is performing during the selected time frame.

About the entities you can examine in the Activity tab

The table below describes which Collector entities can be examined in the Activity tab.

Table 6-1 Collector entities that can be examined in the Activity tab

Collector Entity	Description
Database	Displays information on an activity in reference to a specific database.
Program	Displays information on an activity generated from a specific program. Displays 'N/A' if the program_name column in sysprocesses is null.
Login	Displays information on an activity generated by a specific login. Displays 'N/A' if the suid column in sysprocesses holds a number for which the user name is null.
User	Displays information on an activity generated by a specific user in the database in use.
Machine	Displays information on an activity generated from a specific workstation. If the hostname column in sysprocesses is null, displays the IP Address of the workstation for instances of version 12.5 or later, or 'N/A' for instances of earlier versions.

See [Precise for Sybase agents](#) on page 13.

The following views are available for the Instance and Collector entities:

- Overview
- In Sybase
- Locking
- Overall Activity
- Scalability

About getting an overview Instance and Collector data

The Overview displays information on the performance and resource consumption of the collector entities. It is divided in the following areas.

Load

Displays information on the following session-related items:

- Sessions - number of sessions that occurred during the specified time period of the selected collector entity. For example, if you are viewing the Login entity, this value represents the number of sessions that are connected to the selected Login.
- Duration (Avg) - average duration for a single session, in the selected collector entity during the selected time frame, calculated by dividing the Duration by the number of Sessions.
- Duration (Summed) - total amount of time that elapsed for all of the selected collector entity's sessions during the selected time frame.
- Statement Executions - number of statements executed during the selected time frame that were generated by the selected collector entity's sessions.
- Statement Duration (Avg) - average duration for a single statement execution, in the selected collector entity during the selected time frame. This value is calculated by dividing the Duration (Summed) parameter by the number of statement executions.

Statistics

Displays the following session-related statistics:

- CPU Time - cumulative CPU time for the collector entity or instance during the selected time frame.
- Physical I/O Operations - total number of physical I/O operations performed by all sessions connected to the database during the selected time frame.

In Sybase Breakdown

Displays collector entity resource consumption (table format) and shows instance breakdown during the selected time frame. The table displays the following information:

- State - displays the In Sybase state.
- In Sybase - displays a graphical representation of the Time column.
- Time - amount of time the selected collector entity was in this state.
- % - percentage of time the selected collector entity was in this state

See [About session states](#) on page 25.

About viewing Instance and Collector In Sybase data

The In Sybase view displays the resource consumption of the selected collector entity during the selected time frame. This graph enables you to analyze performance trends for the collector entity, over time.

See “Understanding In Sybase session states” on page 25. See “Focusing on information in overtime graphs” on page 22.

About viewing Instance and Collector locking data

The Locking view displays the blocking tree and the lock wait duration broken down into lock levels (such as, Table, Page, and Key) for the selected entity, over the selected time period. The following information is displayed:

Locks	Displays which sessions were involved in a blocking situation (either locking or blocking). It is important to understand which sessions are involved in each blocking chain, as there may be several sessions involved at the same time. It is then important to determine which session is blocking others, which object is being blocked, and what type of lock has been applied. A yellow lock indicates that the session is currently blocking another session but is not itself blocked. A red lock indicates the session is currently blocked by another session. Note that a session may also block other sessions while being blocked by a different session.
Lock wait graph	The Lock Wait graph shows a breakdown of the Lock Waits and when they occurred over the selected time period.

See [Focusing on information in overtime graphs](#) on page 22.

About viewing Instance and Collector overall activity

The Overall Activity view displays the duration of the activity performed by the selected collector entity during the selected time frame, broken down into the following components:

- Waitfor Command - the amount of time applications were waiting for a Waitfor command.
- Parallel Sync Wait - the amount of time applications were waiting on a synchronization message from another process in the family.
- Request Wait - the amount of time applications were in a Request Wait state, defined as the state wherein the server is waiting for a client request.
- Other Wait - the amount of time applications were in Other states.
- In Sybase - represents total amount of time spent in Sybase.

See [Focusing on information in overtime graphs](#) on page 22.

About viewing Instance and Collector scalability

The Scalability view displays an over time graph that allows you to examine the scalability of the selected collector entity, as follows:

In Sybase (Avg) vs. Sessions	Compares the average In Sybase time of the selected collector entity with the number of sessions during the selected time frame. For example, if the selected collector entity is Login, the graph will show the number of sessions connected with the selected login versus the average In Sybase time of each session over time. Use this graph to analyze if the average session In Sybase time increases when the number of sessions increase.
-------------------------------------	--

See [Focusing on information in overtime graphs](#) on page 22.

Instance and Collector entities Association area

Additional information on performance and In Sybase activity for the instance and collector entities is displayed in the Association area tabs.

Clicking a tab shows additional information on the entities displayed in the Association area table. The rows of the table remain the same; the columns change to display additional details on the entities listed in the table. The information displayed in each tab is described below.

About Statement, Batch, Collapsed Statement, and Collapsed Batch entities**About viewing performance data for Instance and Collector entities**

The table below describes the additional information displayed in the Performance tab in the Association area.

Table 6-2 Precise information displayed in the Performance tab

Column	Description
Entity Name (such as Instance, Database)	Displays the name of the entity.
In Sybase	Displays entity resource consumption (default: stacked bar chart) and shows entity breakdown during the selected time frame.
%	Displays the percentage of the total In Sybase.
Sessions	Number of sessions that occurred during the selected time frame of the selected entity.
Duration (Avg)	Average duration for a single session, in the selected entity during the selected time frame, calculated by dividing the Duration by the number of Sessions.
Statement Executions	Number of statements executed during the selected time frame that were generated by the selected entity's sessions.

Table 6-2 Precise information displayed in the Performance tab

Column	Description
CPU Time	<p>Cumulative CPU time for the entity during the selected time frame.</p> <p>This value is calculated by the Collector agent as follows (this is an approximate value):</p> <p>The <code>CPU</code> field in <code>sysprocesses</code> is sampled at the start point and end point. The value sampled at the start point is subtracted from the value sampled at the end point resulting in the CPU Time reported.</p>
Physical I/O Operations	<p>Cumulative number of physical I/O requests performed for the entity during the selected time frame.</p> <p>This value is calculated by the Collector agent as follows (this value is an approximation):</p> <p>The <code>physical_io</code> field in <code>sysprocesses</code> is sampled at the start point and the end point. The value sampled at the start point is subtracted from the value sampled at the end point resulting in the Physical I/O reported.</p>

About viewing In Sybase data for Instance and Collector entities

The In Sybase tab contains the In Sybase counter and the breakdown of the In Sybase into its states. The In Sybase is displayed with a normalized view while all the other counters display the time spent in each state.

See [Understanding In Sybase session states](#) on page 25.

About Statement, Batch, Collapsed Statement, and Collapsed Batch entities

When an application component (such as Database or Login) consumes resources in Sybase, it is actually executing a statement that is consuming resources. Tuning performance problems typically involves finding and tuning offending statements.

The Statement, Batch, Collapsed Statement and Collapsed Batch entities display the performance and resource consumption of their respective entities.

About Statement, Batch, Collapsed Statement, and Collapsed Batch entities

The table below describes which performance and resource consumption information is displayed for the statement, batch, collapsed statement and collapsed batch entities.

Table 6-3 Statement, Batch, Collapsed Statement and Collapsed Batch entities

Entity	Description
Statement	Displays information on the activity of a specific TSQL statement. A statement is always part of a batch. When an encrypted stored procedure is encountered the statement will display Encrypted Text.
Batch	Displays information on the activity of a specific batch. A batch can either be a stored procedure, stored function, trigger, or a detailed batch. If it is a stored object, the name of the stored object is displayed.
Collapsed statement	Represents similar statements that differ in their parameters (literals). In a collapsed statement, the literals are replaced with parameter markers. The collapsed statement lets you analyze the overall resource consumption of similar statements.
Collapsed Batch	Similar to collapsed statement a collapsed batch represents similar batches that differ in their parameters (literals). In a collapsed batch, the literals are replaced with parameter markers. The collapsed batch lets you analyze the overall resource consumption of similar batches.

See [How information is collected](#) on page 25.

The following views are available for the Statement, Batch, Collapsed Statement and Collapsed Batch entities:

- Overview
- Text
- In Sybase
- Locking
- Scalability

About getting an overview of Statement, Collapsed Statement, Batch and Collapsed Batch entities

The Overview displays information on the entity's activity, such as its number of executions and average statement duration. In addition the resource consumption breakdown and the text of the statement or batch are also displayed, as follows:

Load	Displays information on the following session-related items: <ul style="list-style-type: none"> • Database - identifies the statement's database. • User - identifies the statement's parsing user. • Statement Executions - number of times the Collector agent identifies that the statement was executed during the selected time frame. • Statement Duration (Avg) - displays the average duration of each statement execution during the selected time frame, calculated by dividing the Duration (Summed) by the number of Statement Executions. • Statement Duration (Summed) - total elapsed time required to run all executions of the selected statement or batch during the selected time frame.
-------------	--

About Statement, Batch, Collapsed Statement, and Collapsed Batch entities**Statistics**

Displays the following statement-related statistics:

- **CPU Time** - cumulative CPU time for the entity during the selected time frame. This value is calculated by the Collector agent as follows (this is an approximate value) - the CPU field in sysprocesses is sampled at the start point and end point. The value sampled at the start point is subtracted from the value sampled at the end point resulting in the CPU Time reported.
- **Physical I/O Operations** - total number of physical I/O requests during statement execution during the selected time frame.
- **Parallel Sessions (Min)** - certain statements are executed using parallelisms. This counter displays the minimum number of threads used to execute the statement and is calculated by examining the fid column in sysprocesses. Check this counter to verify that the Sybase is implementing the best execution plan for the statement.
- **Parallel Sessions (Max)** - certain statements are executed using parallelisms. This counter displays the maximum number of threads used to execute the statement and is calculated by examining the fid column in sysprocesses. Check this counter to verify that the Sybase is implementing the best execution plan for the statement.

In Sybase breakdown

Displays resource consumption broken down into states, such as, Using CPU, I/O Wait, Lock Wait, and so on, within the specified time period. This allows you to identify the highest resource consumers.

Text

Displays statement or batch text (formatted for easy readability).

When viewing a batch entity, the entire text of the improvised batch or stored object is displayed. When viewing a statement, only the part of the batch that generates the statement is displayed. When viewing a collapsed statement or batch, the literals are replaced with parameter markers.

About viewing Statement, Collapsed Statement, Batch and Collapsed Batch text

The Text view displays the SQL Text of the statement or batch.

About viewing In Sybase data for Statement, Collapsed Statement, Batch and Collapsed Batch entities

The In Sybase view displays the resource consumption of the selected entity during the selected time frame. This graph enables you to analyze performance trends for the entity, over time.

See [Focusing on information in overtime graphs](#) on page 22.

About viewing Statement, Collapsed Statement, Batch and Collapsed Batch entities locking

The Locking view displays the blocking tree and the lock wait duration broken down into lock levels (such as, Table, Page, and Key) for the selected entity, over the selected time period. The following information is displayed:

Locks

Displays which sessions were involved in a blocking situation (either locking or blocking). It is important to understand which sessions are involved in each blocking chain, as there may be several sessions involved at the same time. It is then important to determine which session is blocking others, which object is being blocked, and what type of lock has been applied.

A yellow lock indicates that the session is currently blocking another session but is not itself blocked.

A red lock indicates the session is currently blocked by another session. Note that a session may also block other sessions while being blocked by a different session.

Lock wait graph

The Lock Wait graph shows a breakdown of the Lock Waits and when they occurred over the selected time period.

See [Focusing on information in overtime graphs](#) on page 22.

About viewing scalability for Statement, Collapsed Statement, Batch and Collapsed Batch entities

The Scalability view displays an over time graph that allows you to examine the scalability of the selected collector entity during the selected time frame. For example, you can determine whether the average statement In Sybase time of an entity's activity increased as a result of an increase in number of executions.

The following over time graph is displayed:

In Sybase (Avg) vs. executions

Compares the average In Sybase time of the selected entity with the number of executions during the selected time frame. For example, if the selected entity is a Statement, the graph will show the number of executions of the selected statement versus the average In Sybase time of each execution over time. Use this graph to analyze if the average statement In Sybase time increases when the number of executions increase.

Statement and Batch entities Association area

The following tabs are displayed in the Association area of the Statement and Batch entities:


- Performance tab (Statement and Batch)
- In Sybase tab (Statement and Batch)
- Statistics tab (Statement and Batch)
- Plan tab (Statement and Batch)

Clicking a tab shows additional information on the entities displayed in the Association area table. The rows of the table remain the same; the columns change to display additional details on the entities listed in the table. The following tables describe the information displayed in each tab.

About viewing performance data for Statement and Batch entities

The table below describes the additional information displayed in the Performance tab in the Association area.

Table 6-4 Information displayed in the Performance tab

Column	Description
	Click to launch the SQL tab with the selected entity, to continue your analysis of the selected statement or batch.
Entity Name	Displays the name of the entity (statement or batch).
In Sybase	Displays entity resource consumption (default: stacked bar chart) and shows entity breakdown during the selected time frame.
%	Displays the percentage of the total In Sybase.
Statement Executions	Number of statements executed during the selected time frame by the selected entity.
Statement Duration (Avg)	Average duration for a single statement, in the selected entity during the selected time frame, calculated by dividing the Statement Duration by the number of executions.
Text	Displays the statement's text.


About viewing In Sybase data for Statement and Batch entities

The In Sybase tab contains the In Sybase counter and the breakdown of In Sybase into its states. In Sybase is displayed with a normalized view while all the other counters display the time spent in each state.

About viewing Statistical data for Statement and Batch entities

The table below describes the additional information displayed in the Statistics tab in the Association area.


Table 6-5 Information displayed in the Statistics tab

Column	Description
	Click to launch the SQL tab with the selected entity, to continue your analysis of the selected statement or batch.
Entity Name (such as Statement)	Displays the name of the entity.
CPU Time	Cumulative CPU time for the entity during the selected time frame.
Physical I/O Operations	Total number of physical I/O requests performed by the selected entity during the selected time frame.
Parallel Sessions (Min)	Minimum number of threads used to execute the statement or batch in parallel. This counter is calculated from the fid column in sysprocesses. This enables you to verify that Sybase is using the best execution plan for the current statement.
Parallel Sessions (Max)	Maximum number of threads used to execute the statement or batch in parallel. This counter is calculated from the fid column in sysprocesses. This enables you to verify that Sybase is using the best execution plan for the current statement.

About viewing Plan data for Statement and Batch entities

The table below describes the additional information displayed in the Plan tab in the Association area.

Table 6-6 Information displayed in the Plan tab

Columns	Description
	Click to launch the SQL tab with the selected entity, to continue your analysis of the selected statement or batch.
Entity Name (such as Statement)	Displays the name of the entity.
Database	Identifies the database on which the statement or batch is run.
User	Displays one of the Sybase users by whom this statement was executed. If this statement is not part of a stored procedure, this user will be used as parsing user when explaining this batch.
Most Recent Plan	Indicates the last time the access path of the statement or batch was changed.
Last Show Plan	Indicates the last time the statement or batch was explained.
Cost	Indicates the total estimated cost of the statement's or batch's execution.
Table Scan	Indicates whether a table scan was performed in the execution plan (applicable to Statement entity only).
Clustered Index Scan	Indicates whether a clustered index scan was performed in the execution plan (applicable to Statement entity only).
Clustered Index Seek	Indicates whether a clustered index seek was performed in the execution plan (applicable to Statement entity only).
Index Scan	Indicates whether an index scan was performed in the execution plan (applicable to Statement entity only).
Index Seek	Indicates whether an index seek was performed in the execution plan (applicable to Statement entity only).
Table Lookup	Indicates that a table lookup operation was performed in the execution plan (applicable to Statement entity only).
Sort	Indicates whether a sort operation was performed in the execution plan (applicable to Statement entity only).

How the Activity tab can help you identify performance problems

Columns	Description
Merge Join	Indicates whether a merge join operation was performed in the execution plan (applicable to Statement entity only).
Join	Indicates that a join (other than merge join) operation was performed in the execution plan (applicable to Statement entity only).
Nested Loops Join	Indicates whether a nested loops join operation was performed in the execution plan (applicable to Statement entity only).
Parallel Access	Indicates whether a parallel access operation was performed in the execution plan (applicable to Statement entity only).
Show Plan Error	Shows the error that occurred during the last explain process.

Collapsed Statement and Collapsed Batch entities Association area

The following tabs are displayed in the Association area of the Collapsed Statement and Collapsed Batch entities:

- Performance tab (Collapsed statement and collapsed batch)
- In Sybase tab (Collapsed statement and collapsed batch)
- Statistics tab (Collapsed statement and collapsed batch)

Clicking a tab shows additional information on the entities displayed in the Association area table. The rows of the table remain the same; the columns change to display additional details on the entities listed in the table. The information displayed in each tab is described below.

About viewing performance data for Collapsed Statement and Collapsed Batch entities

The table below describes the additional information displayed in the Performance tab in the Association area.

Table 6-7 Information displayed in the Performance tab

Column	Description
Entity Name	Displays the name of the Collapsed Statement or Collapsed Batch.
In Sybase	Displays entity resource consumption (default: stacked bar chart) and shows entity breakdown during the selected time frame.
%	Displays the percentage of the total In Sybase.
Occurrences	Number of different statements possessing the selected collapsed form.
Statement Executions	Number of statements executed during the selected time frame by the selected entity.
Statement Duration (Avg)	Average duration for a single statement, in the selected entity during the selected time frame, calculated by dividing the Statement Duration by the number of executions.
Text	Displays the statement or batch text, in the collapsed form.

In Sybase tab (Collapsed Statement and Collapsed Batch)

The In Sybase tab contains the In Sybase counter and the breakdown of In Sybase into its states. In Sybase is displayed with a normalized view while all the other counters display the time spent in each state.

About viewing statistical data for Collapsed Statement and Collapsed Batch entities

The table below describes the additional information displayed in the Statistics tab in the Association area.

Table 6-8 Information displayed in the Statistics tab

Column	Description
Entity Name (such as Collapsed Statement)	Displays the name of the entity.
Database	Identifies the database on which the collapsed statement or batch is run.
CPU Time	Cumulative CPU time for the entity during the selected time frame.

How the Activity tab can help you identify performance problems

Column	Description
Physical I/O Operations	Total number of physical I/O requests performed by the selected entity during the selected time frame.
Parallel Sessions (Min)	Minimum number of threads used to execute the statement or batch in parallel. This counter is calculated from the fid column in sysprocesses. This enables you to verify that Sybase is using the best execution plan for the current statement.
Parallel Sessions (Max)	Maximum number of threads used to execute the statement or batch in parallel. This counter is calculated from the fid column in sysprocesses. This enables you to verify that Sybase is using the best execution plan for the current statement.

About Session and Blocker Session entities

The Session entity displays performance, statistics and locking information on the selected session. This entity differs from other entities in the Activity tab in that session information is retrieved from the Collector agent's temporary files and not from the PMDB.

The Blocker Session entity displays information on the holding locks that cause other sessions to wait for resources to be free.

The following views are available:

- Overview
- In Sybase
- Locking
- Overall Activity

About getting an overview of Session and Blocker Session entities

The Overview displays general details on the selected session or blocker session, such as, Machine, Login, and Client Process ID. A breakdown of the selected session's or blocker session's resource consumption is also displayed.

About viewing In Sybase data for Session and Blocker Session entities

The In Sybase view displays the resource consumption of the selected session during the selected time frame. This graph enables you to analyze performance trends for the session, over time.

NOTE The graph displays data in time slices.

See [Understanding In Sybase session states](#) on page 25. See [Focusing on information in overtime graphs](#) on page 22.

About viewing Session and Blocker Session locking

The Locking view displays the blocking tree and the lock wait duration broken down into lock levels (such as, Table, Page, and Key) for the selected entity, over the selected time period. The following information is displayed:

Locks	<p>Displays which sessions were involved in a blocking situation (either locking or blocking). It is important to understand which sessions are involved in each blocking chain, as there may be several sessions involved at the same time. It is then important to determine which session is blocking others, which object is being blocked, and what type of lock has been applied.</p> <p>A yellow lock indicates that the session is currently blocking another session but is not itself blocked.</p> <p>A red lock indicates the session is currently blocked by another session. Note that a session may also block other sessions while being blocked by a different session.</p>
Lock wait graph	The Lock Wait graph shows a breakdown of the Lock Waits and when they occurred over the selected time period.

NOTE The graph displays data in time slices.

About viewing overall activity for Session and Blocker Session entities

The Overall Activity view displays the duration of the activity performed by the selected session within the specified time period, broken down into the following components:

- Waitfor Command - the amount of time applications were waiting for a Waitfor command.
- Parallel Sync Wait - the amount of time applications were waiting on a synchronization message from another process in the family.
- Request Wait - the amount of time applications were in a request wait state, defined as the state wherein the server is waiting for a client request.
- Other Wait - the amount of time applications were in other states.
- In Sybase - represents total amount of time spent in Sybase.

NOTE The graph displays data in time slices.

See [Focusing on information in overtime graphs](#) on page 22.

Session and Blocker Session entities Association area

The following tabs are displayed in the Association area of the Instance and Collector entities:

- Performance tab (Session and Blocker Session)
- In Sybase tab (Session and Blocker Session)
- Statistics tab (Session and Blocker Session)

Clicking a tab shows additional information on the entities displayed in the Association area table. The rows of the table remain the same; the columns change to display additional details on the entities listed in the table. The following table describes the information displayed in each tab.

About viewing performance data for Session and Blocker Session entities

The table below describes the additional information displayed in the Performance tab in the Association area.

Table 6-9 Information displayed in the Performance tab

Column	Description
SPID	Displays the system process ID. The SPID is a unique integer assigned to each user connection when the connection is made.
Program	Displays the program that opened the session.
In Sybase	Displays entity resource consumption (default: stacked bar chart) and shows entity breakdown over the specified time period.
Database	Displays the database used by session.
Login	Displays the Sybase login that opened the session.
Stat Time	Displays when the session was started during the selected time frame.
Completed At	Displays when the session was completed during the selected time frame.
Machine	Displays the machine the program ran on.
User	Displays the database user that opened the session.
Duration (Summed)	Displays the time that has elapsed since the login time.
Client Process ID	Displays the process ID of the program.

About viewing In Sybase data for Session and Blocker Session entities

The In Sybase tab contains the In Sybase counter and the breakdown of In Sybase into its states. In Sybase is displayed with a normalized view while all the other counters display the time spent in each state.

How the Activity tab can help you identify performance problems

About viewing statistical data for Session and Blocker Session entities

The table below About viewing statistical data for Session and Blocker Session entities.

Table 6-10 Information displayed in the Statistics tab

Column	Description
SPID	Displays the system process ID. The SPID is a unique integer assigned to each user connection when the connection is made.

Table 6-10 Information displayed in the Statistics tab

Column	Description
CPU Time	Cumulative CPU time for the entity during the selected time frame.
Physical I/O Operations	Total number of physical I/O requests performed by the selected session during the selected time frame.
Memory Pages Used	Displays the memory consumed (in pages) by the selected session during the selected time frame.

See [Understanding In Sybase session states](#) on page 25.

How the Activity tab can help you identify performance problems

You can identify a performance problem by doing one or more of the following:

- [Examining resource consumption of an entire instance and database](#)
- [Identifying heavy resource consumers](#)
- [Examining resource consumption over time](#)
- [Resolving blocking situations](#)
- [Identifying offending SQL statements](#)
- [Examining the scalability of your application or Sybase servers](#)

Examining resource consumption of an entire instance and database

When performing a tuning audit it is very important to analyze and understand the performance behavior of your instances and databases. You may have already drilled down to analyze the performance of a single query or program. However, examining the entire instance's behavior can alert you to the health of your system. The Instance (or Database) represents the performance of the average application. This helps you answer questions such as: "This query suffers from Log wait, but does the entire system suffer from lock wait?" If, for example, you discover that your system suffers from I/O wait and you tune a query that performs memory operations, you most likely will not improve the entire system's I/O wait problem. So examining the entire instance can provide a quick overview of the dominant resources consumed.

To examine the resource consumption of an entire instance and database

1. In the Instance list, choose the item you want to analyze.
2. In the Time Frame list, choose the period of time you want to analyze.
3. On the View controls in the Main area, click **Overview**. Examine the entire instance and determine which are the dominant resources that are consumed by your system.
4. On the Association controls in the Association area, click **Databases**.
5. On the Performance tab, place the cursor in the In Sybase column of the top database and view the information displayed on its resource consumption.

NOTE As a rule of thumb, a healthy system should have a high Using CPU value, a 10 - 15% I/O Wait, a Log Wait based on the nature of the application, and the remaining states should show values that are as minimum as possible.

Identifying heavy resource consumers

Precise for Sybase lets you drill down to application components (such as Logins, Databases, Programs, Machines, and Statements) to determine which application component consumes the most resources. This is accomplished by clicking on an application component in the Association area. This process is iterative and you may continue to drill down until you discover the application component that you want to tune.

There are several reasons to drill down and focus on different application components:

- During the tuning process, you drill down to determine which are the top resource-consuming components. Tuning a component (such as a Database, Login, User, or Machine) that consumes most of the resources, will free them to other entities. The tuning process is an iterative process. You will continue to list all components, examine their resource consumption and focus in on several of them, until you have completed the tuning process.
- When you try to analyze user experience you need to isolate one user's activity from the other users. To achieve this you can focus in on a single login or machine.
- When trying to determine which application component to drill down to, take into consideration not only the total In Sybase of the application, but also its number of sessions, executions and average duration. If you do not take these elements into consideration you may try to tune a program that runs once a year, consuming 20 hours, instead of tuning a program that consumes 2 seconds but runs every time a user logs on to the system.

To identify heavy resource consumers

1. In the Instance list, choose the item you want to analyze.
2. In the Time Frame list, choose the period of time you want to analyze.
3. On the View controls in the Main area, click **In Sybase** to view a breakdown of the In Sybase time.

NOTE When trying to determine which application component to drill down to, take into consideration not only the total In Sybase of the application, but also its number of Sessions or Executions and Duration (Avg). If you do not take these elements into consideration you may try to tune a program that runs once a year, consuming 20 hours, instead of tuning a program that consumes 2 seconds but runs every time a user logs on to the system.

4. On the Association controls in the Association area, click **Programs**.
5. On the Performance tab, place the cursor in the In Sybase column of the top program entity and view the information displayed regarding its resource consumption.
6. Drill down on the entity to focus on additional components.
7. Go back to the overview of the instance you originally selected.
8. In the Association area, use the Association controls to select different entities associated with the instance, and analyze the average duration of each entity.

Examining resource consumption over time

When you analyze and tune the performance of an application component (such as, database, program, or query) it is important to take into consideration its performance over time. Precise for Sybase lets you view the component's performance over a selected time period.

There are several reasons to study a component's performance over time:

- When analyzing a performance problem you will want to determine if it is a one-time only problem or has happened before.
- You will want to identify if there are any patterns or trends. For example, does a particular lock problem always happen on Tuesday at around 10 AM?
- You will want to determine how your system is behaving throughout an entire day. For example you may observe that during the night your system experiences much more I/O wait than during the day.

To examine resource consumption over time

1. In the Instance list, choose the item you want to analyze.
2. In the Time Frame list, choose the period of time you want to analyze.
3. On the View controls in the Main area, click **In Sybase**.
4. Move the cursor over the bars in the graph to view the instance's resource consumption over time.

Resolving blocking situations

While many think that locking is bad, the truth is that locking is good - it maintains data integrity. The situations to avoid are contentions. Contentions are situations in which one application waits for a resource that is being held by another application.

Locking can become problematic if the portion of time the application is blocked lasts too long. If you identify a lock wait time that seems too long, further investigation is required. The Lock wait time is reported as part of the In Sybase breakdown.

Analyzing the blocking situation is much easier when the lock is currently active, but this is usually not the case. For this reason it is important to understand why the locks occurred in the past.

There are several areas you can focus on when resolving blocking situation:

- Look at the Overview view of your instance (database, program) to understand the percentage of Lock Wait your application suffers.
- Examine the Lock Wait over a selected time period. By examining resource consumption over time you can identify if Lock Wait occurred all the time, or only on specific dates.
- Try to gain an understanding of the Lock types your applications were blocked on. An overtime graph of the different Lock types can be displayed in the Locking view.
- Examine the locked objects. You can display a list of all locked objects and compare them according to lock time.
- Once you have identified all of the above, you also need to assess who is blocking who. For example, you will need to determine which session is blocking which session? The locks tree displays all the contentions that occurred and shows which sessions were involved in contentions.

To examine locking over time of a database entity

1. In the Instance list, choose the item you want to analyze.
2. In the Time Frame list, choose the period of time you want to analyze.
3. In the Association area drill down to the database entity you want to analyze.
4. On the View controls in the Main area, click **Locking**.
5. Examine the locking of the selected entity in the graph, for the period of time you selected. Move the cursor over the bars in the graph to view the Lock Wait breakdown.

Examine the locks tree to observe all the contentions that occurred and determine which sessions were involved in contentions.
6. Assess who is blocking who (For example, determine which session is blocking which session. By drilling from the application component with Lock Wait to the Blocker Session entity, you can find the sessions that caused the contention).

Identifying offending SQL statements

Having identified a major resource-consuming entity, you can further drill down to the statements and batches executed by that entity. These may include long running SQL statements and SQL statements that use few resources but were executed frequently. In many cases analyzing the queries executed by the entity and tuning them will result in the most performance gain. For each statement, you can do the following:

- View which entities executed the statement.
- View resources consumed by the statement, which enables you to determine if the statement is CPU or I/O bound, waiting for a lock, etc.
- View resource consumption patterns, which allow you to observe if the statement execution is affected by excessive use of the same resource by another statement.
- Tune the statement in the SQL tab.

To identify offending SQL Statements

1. In the Instance list, choose the item you want to analyze.
2. In the Time Frame list, choose the period of time you want to analyze.
3. In the Association area, drill down to the top resource-consuming database.
4. On the Association controls in the Association area, click **Statements**. On the Performance tab view which entities executed the statement and view the resources consumed by the statement.
5. In the Plan tab, in the Association area, view a summary of the statement's access plan.

Examining the scalability of your application or Sybase servers

Applications are called upon to support the addition of more and more users, and data, over the years. One of today's tasks is ensure that our application and servers are scalable for tomorrow's tasks.

Precise for Sybase lets you examine if the application is scalable, or if the Sybase servers are scalable.

To analyze scalability of an application you need to compare the number of sessions and the impact of the average session duration. In a well scalable instance, growth in the number of sessions will not impact the average session duration. This can be observed in the Scalability View in the Session Duration (Avg) vs. Sessions graph. The same applies to I/O and CPU, as you can see in the Physical I/O Operations vs. I/O Wait graph, which shows how more and more I/O requests affect the I/O waits.

When examining if the server is scalable, check if Sybase server CPU consumption growth affects the server CPU consumption and I/O wait times. You will need to answer if growth in the I/O requests increases the time I/Os are waiting in the Queue to be served.

The first example below shows that the application is not scalable. The Session Duration (Avg) decreases when the number of sessions increase.

To examine the scalability of your application or Sybase servers

1. In the Instance list, choose the item you want to analyze.
2. In the Time Frame list, choose the period of time you want to analyze.
3. On the View controls in the Main area, click **Scalability**.
4. In the Session Duration (Avg) vs. Sessions graph, observe the number of sessions and the impact of the average session duration, over the selected time period. For example, if you observe that average session duration decreases when the number of sessions increase, this may indicate that the application has a scalability problem.

Examining the relationships between database objects

This section includes the following topics:

- [About the Objects tab](#)
- [How the Objects tab is structured](#)
- [About the entities you can examine in the Objects tab](#)
- [How the Objects tab can help you identify performance problems](#)

About the Objects tab

The Objects tab is a comprehensive browser, which allows you to understand the relationships and associations between database schema objects. It enables a more productive tuning process by reducing the need of another tool for exploring the Sybase catalog tables.

The Objects tab also allows you to observe storage utilization, over time, at varying levels, such as, database, segment, database device, and so on.

How the Objects tab is structured

The Objects tab displays information on a selected entity and its associated entities. When you open the Objects tab the selected entity is by default Instance, meaning that information is displayed on the instance level. When you launch the Objects tab from the Activity tab with the database entity - the tab is launched in-context with the database you selected in the Activity tab.

If you open the Objects tab from another tab, the historical settings (meaning those settings which were selected when you left the tab, such as the last entity you drilled down to) are taken into account and the information displayed the last time you viewed this tab is displayed (similar to clicking the History button and returning to a previous tab).

The selected entity is always reflected in the Tab heading, which serves as a point of orientation. The highest-level entity you can view information for in the Objects tab is the Instance level. You can select an instance from the Instance list.

The entities displayed in the Association area are associated with the selected entity displayed in the Main area. At times, the relationship between the entity displayed in the Main area and those displayed in the Association area is that of parent to child and sometimes it merely represents that there is a relationship between the selected entity and the entities displayed in the Association area.

About the Main area in the Objects tab

The Main area shows comprehensive information on the selected entity. You can choose from several views to examine the entity from different angles. You can, for example, focus exclusively on Index Properties, Index Statistics, or Index Histograms.

About the Association area in the Objects tab

The Association area provides corresponding information on the child entities. You can view information on one type of child entity at a time, such as only instances or only databases, by selecting a control from the View Controls list. The selection you make is only reflected in the Association area; the Main area remains unchanged.

From the Association area, you can also drill down to another entity by clicking a table row. A drilldown affects the whole tab: When you drill down to another entity, the Tab heading reflects the new selection, the Main area displays information on the newly selected entity, and the Association area shows data on its associated entities.

For example, when you want to view information on a specific database, choose Database from the Association controls. The Association area changes to display all the databases associated with the current instance. Note that the Tab heading and the Main area remain unchanged.

In the Association area, click the row of the database you want to view detailed information for. The Tab heading indicates the newly selected entity; the Main area displays an overview of the database you want to investigate, and the Association area shows general details on the tables that are associated with the database you selected.

See [About Precise for Sybase tabs](#) on page 15.

About the entities you can examine in the Objects tab

The Objects tab displays information on different entities. This section provides an overview of all entities, their meaning, and their views.

- About the Instance entity
- About the Database entity
- About the Table entity and Tables Contain Same Column entity
- About the Index entity
- About viewing data storage and fragmentation information for the selected index
- About the Column entity
- About the View entity and Views Contain Same Column entity
- About the User entity
- About the Database Device entity
- About the Segment entity
- About the Data Cache entity
- About the Unused Indexes entity
- About the Statement entity
- About the Batch entity

About the Instance entity

The Instance entity displays configuration information, and performance information, as collected by the Precise for Sybase Collector agent for the entire instance. The following views are available:

- Overview
- Instance parameters

About getting an overview of instance configuration and performance information

The Overview displays configuration and performance information for the entire instance. The Overview is divided into the following areas:

Instance details	Details server configuration details, such as version and number of engines.
In Sybase	Displays resource consumption during the selected time frame, of the selected instance, as collected by the Precise for Sybase Collector agent
Space Allocation	Displays the total free space and total used space (in MB) during the selected time frame, of all the files in the selected instance, as collected by the Collect Space Utilization process.

About viewing instance configuration parameters

The Instance Parameter view displays the selected instance configuration parameters. This includes the possible range values, configuration (pending) values and run values, for each parameter.

About the Database entity

The Database entity displays configuration information, and performance information collected by the Precise for Sybase collector of the selected database. The following views are available:

- Overview
- Database options
- Replication Agent configuration parameter

About getting an overview configuration, performance, and Space Allocation information for a selected database

The Overview displays configuration, performance, and Space Allocation information for the selected database, as follows:

Properties	Displays general database details, such as status, owner, and log segment free space.
In Sybase	Displays resource consumption of the selected database during selected time frame, as collected by the Precise for Sybase Collector agent.
Space Allocation	Displays total free space and total used space (in MB) during the selected time frame, of all the segments in the selected database, as collected by the Collect Space Utilization process.
Data vs. log	Displays total free space and total used space (in MB) during the selected time frame, of all the segments in the selected database, as collected by the Collect Space Utilization process, and displayed according to segment usage as data only, log only and data and log segments.

About viewing Database options

The Database Options view displays database options, such as abort tran on log full and trunc. log on chkpt, as follows:

Database options	Displays the current value of the database options.
In Sybase	Displays resource consumption information of the selected database during the selected time frame.

About the Replication Agent configuration parameter view

This table describes the Replication Agent configuration parameters for the specified database. This view is only enabled for published databases. The Replication Agent Configuration Parameters in the Main area can be accessed as follows:

1. In the Objects tab in Precise for Sybase, select the desired database in the Association area.
2. Click **Overview** in the Main area, and then click **Replication Agent Configuration Parameters**. The Replication Agent Configuration Parameters appear in the Main area.

The headings in the Replication Agent Configuration Parameters main area are described below.

Table 7-1 Replication Agent Configuration Parameters headers

Column	Description
Parameter Name	Displays the parameter name.
Default Value	Displays the default value for each parameter.
Configuration Value	Displays the configuration value for each parameter.
Run Value	Displays the run value for each parameter.

About the Database entity in the Association area

You can view the entities associated with the selected Database entity in the Association area. It is possible to associate to the Database entity from the Instance entity.

About the Table entity and Tables Contain Same Column entity

The Table entity displays general details regarding space utilization for the selected table.

The Tables Contain same Column Entity displays tables that have the same column as the column displayed in the Main area. The information displayed is the same as displayed in the Table entity and can only be associated to from the Column entity.

The following views are available:

- Overview
- Optdiag
- Table partitions

About getting an overview of table entities

The Overview displays general table properties, along with space utilization, and number of rows in the selected table during the selected time frame. It is divided in the following areas:

Properties	Displays general table details, such as, number of columns, data cache, and whether or not the table uses a clustered index.
Space Allocation	<p>Displays the amount of free space and used space (in MB) in the selected table during the selected time frame, and specifies whether the space is populated by an index or data.</p> <p>The values may be incorrect due to out-of-date space usage information. To update the values, use the <code>DBCC FLUSHSTATS</code> command to flush the most up-to-date information regarding the objects size from memory and run <code>UPDATE STATISTICS</code> to recalculate the statistics. You can also use the <code>DBCC CHECKSTORAGE</code> command to verify and recalculate the space usage information.</p>
Rows	Displays the number of rows in the selected table during the selected time frame, as collected by the Collect Space Utilization process.

About viewing data storage and fragmentation information

The Optdiag view displays data storage and fragmentation information for the selected table. The following information is displayed:

Optdiag

Optdiag displays the following information:

- Pages - displays the number of data pages.
- Empty Pages - the number of empty data pages. A non-zero count, indicates fragmentation, which can be resolved with `reorg`.
- Rows - the total number of data rows.
- Number of forwarded records - when there is not enough free space to write an updated record in its page, Sybase allocates the updated record to another page and instructs the old record to point to the new location. The number of forwarded records, is the number of records allocated to another page, since there was no free space in the original page. Accessing forwarded records adversely affects performance, because additional I/O is performed. Check the number of forwarded records in your table or index in relation to the total number of records. If the ratio is high, consider reorganizing the table or index to achieve a better performance. To reorganize data-only-locked tables, use the `reorg` command. This number is always zero for allpages-locked (APL) tables.
- Deleted Rows - the number of committed deletes, whose space has not yet been reclaimed. An increase in this number means increase in fragmentation as well. This can be resolved using `reorg`.
- Extents - displays the number of extents used to store the data pages.
- OAM + Allocation Page Counts - the sum of Object Allocation Map (OAM) pages plus allocation pages for the table. An Object Allocation Map (OAM) shows how the Adaptive Server allocates space for objects. Each allocation page for an object has an entry in the OAM pages for that table or index. There is at least one OAM page per table.
- First Extent Data Pages - the number of pages that appear on the first extent in an allocation unit, and consequently share the extent with an allocation page. This number is useful for I/O estimation.
- Min Record Size - indicates the minimum record size in the data pages.
- Max Record Size - indicates the maximum record size in the data pages.
- Avg Record Size - indicates the average record size in the data pages.

Derived Statistics

Derived Statistics displays the following information:

- Average Extents Ratio - ratio of the size of the data to the size that can be stored in all pages allocated to the table. Gives an estimation of how full each page really is. Should be as close to 1.0 as possible. This statistic is only displayed in versions prior to version 12.5.0.3.
- Data Page Cluster Ratio - this is a measure of how well data pages are clustered within extents. The closer this number is to 1.0, the better. When the cluster ratio is low, then performance may be affected because it would take more I/O calls to read the table pages. This derived statistic is only displayed in version 12.5.0.3 and above.

Table Definition Related Values

Table Definition Related Values displays the following information:

- Max Rows Per Page - value of `max_rows_per_page` table parameter. Setting a maximum number of rows per page can reduce contention for allpages-locked tables and indexes, as fewer rows are placed on each page.
- Expected Row Size - value of the `exp_row_size` table definition parameter. Relevant for data-only-locked tables. Useful when an application allows rows that contain null values or short variable-length character fields to be inserted, and these rows grow in length with subsequent updates. The major purpose of setting an expected row size is to reduce row forwarding.
- Reserved Page Gap - value of `reservepagegap` table definition parameter. The `reservepagegap` space management property is used to reserve empty pages for expansion, when additional pages need to be allocated. Setting this parameter can reduce the frequency of maintenance activities, such as record rebuild and re-creating indexes.

Extent Fragmentation (%) The extent fragmentation indicates how well the extents, which contain the data pages, are physically contiguous on the disk.

The page number of its first page identifies an extent. If the extents 8, 16, 24, or 32 make up an index, then we can see that the extents are contiguous on the disks, which means there are no gaps. On the other hand, if the extents 8, 24, 32, and 64 make up an index, then we can see that we have a gap. The extent fragmentation is computed by dividing the number of gaps by the number of extents, so in this example, the extent fragmentation is 1/4, or 25 percent.

The lower the number, the less fragmentation there is.

This number is not relevant to heaps because physical order has no importance in the data pages.

NOTE Derived statistics may show incorrect data, when table statistics are out-of-date. It is advised to run `update statistics`, before reviewing these figures.

About viewing data on table partitions

The Table Partitions view displays size and distribution of table storage over all of its partitions. The following information is displayed:

Avg. Partition Pages	Displays the average number of pages in each partition.
Min. Partition Pages	Displays the number of pages in the smallest partition table.
Max. Partition Pages	Displays the number of pages in the largest partition table.
Partition Pages Ratio (Max/Avg)	The ratio between the largest partitions to the average partition size. A number that is much bigger than 1.0 is an indication that the table rows are not well balanced between the partitions.
Partition size	Visual representation of the way table storage is spread among all partitions.

About the Index entity

The Index entity displays general details regarding space utilization for the selected index. The following views are available:

- Overview
- Optdiag
- Statistics

About getting an overview of index entities

The Overview displays the following information:

Properties	Displays general details on the index such as number of keys and whether the index is clustered or unique.
Space Allocation	Displays the amount of free space and used space (in MB) in the selected index during the selected time frame.
	The values may be incorrect due to out-of-date space usage information. To update the values, use the <code>DBCC FLUSHSTATS</code> command to flush the most up-to-date information regarding the objects size from memory and run <code>UPDATE STATISTICS</code> to recalculate the statistics. You can also use the <code>DBCC CHECKSTORAGE</code> command to verify and recalculate the space usage information.
Rows	Displays the number of rows in the underlying table during the selected time frame, as collected by the Collection Space Utilization process.

About viewing data storage and fragmentation information for the selected index

The Optdiag view displays data storage and fragmentation information for the selected index. The following information is displayed:

Pages	Displays the number of index pages at leaf level. For a clustered index on the APL table, this is the same as the value displayed for the table.
Empty Pages	The number of empty leaf pages in the index. Any empty pages are counted towards the total page count by the optimizer and are a waste both of space and page count. If the number increases, consider dropping and recreating the index or running <code>reorg</code> .
Data Page CR Count	This number is used by Sybase to derive the data page cluster ratio, when accessing the table with this index. All CR counts are a measure of the additional I/O required, over a perfect clustering scenario. The server uses CR counts to calculate various cluster ratios that can be obtained by running the optdiag utility.
Index Page CR Count	This number is used by Sybase to derive the index page cluster ratio. It is shown for dol-clustered and non-clustered indexes only.
Data Row CR Count	A measure of how well clustered data rows are on data pages, in relation to rows in the leaf level of this index. This value indicates how many I/O should be done to read row sequentially from the leaf of the index. This is important since it is not always true that for every read of a data row from the leaf of a non-clustered index, one I/O would have to be done from the data pages. This is because sequential rows may be on the same data page. This value now allows the optimizer to estimate how efficient a non-clustered access will be. It is shown for dol-clustered and non-clustered indexes only.
First Extent Leaf Pages	This is the number of leaf-pages that appear on the first extent in an allocation unit, and consequently share the extent with an allocation page. This number is useful for I/O estimations.
Number of Extents	Number of extents used to store the pages in the leaf level of the index.
Index Height	Number of levels in the index.
Min. Record Size (Bytes)	Indicates the minimum record size in the index page.
Max Record Size (Bytes)	Indicates the maximum record size in the index pages.
Avg. Record Size (Bytes)	Indicates the average record size in the index pages.
Max Non-Leaf Record Size (Bytes)	Indicates the maximum record size in the intermediate pages of the index.
Fill Factor (%)	The value for the fill factor of a table/index, set at table definition or using the <code>sp_chgattribute</code> command. It specifies how full Adaptive Server makes each page when it creates a new index on existing data. The fill factor percentage is relevant only when the index is created. As the data changes, the pages are not maintained at any particular level of fullness.

About viewing derived statistics

The items in the following table are only displayed for instances in version 12.5.0.3 and above.

Updated statistics	The last date where update statistics was run on the table.
Data page cluster ratio	This is a measure of the clustering of data pages in extents. The figure takes into account the type of table; the ratio for DOL tables measures clustering with respect to the order of index rows, whereas the ratio for APL tables measures clustering with respect to the order of the page chain. The data page cluster ratio is used to estimate the cost of large I/O. The higher this ratio, the more efficient large I/O will be
Index page cluster ratio	This ratio measures clustering of index leaf pages in extents with respect to the leaf page chain. It is used to estimate the cost of large I/O.
Data row cluster ratio	This is a measure of the clustering of data rows on pages with respect to the index key order. This ratio helps in estimating the cost of using a non-clustered index, or a clustered index on a DOL table.

About viewing clustered index details

This section applies to a clustered index on an APL table. The counters shown are the same as those shown at table level. This section is only shown if the index is a clustered index.

About viewing density and selectivity statistics

The Statistics view displays density and selectivity statistics for each prefix subset of the selected index keys. Histograms for the index columns can be seen at the relevant column statistics view.

The following information is displayed:

Columns	The left-based subset of key columns in the index (including the first column alone).
Statistics Updated	Date the index was created, date that update statistics was last run, or date that optdiag was last used to change statistics
Range Cell Density	A measure of the density of values in the column, after values in frequency count cells has been eliminated. Basically, this is the density of all values that are not highly duplicated. This value is used to estimate the cost of an equality search arguments when the value falls into a range cell. Indicates the number of rows sampled for the statistics calculation.
Total Density	A measure of the duplicate values in all rows of the column, no matter what cell type they fall into. This value is used to estimate the number of rows that will be returned for a join on this column. It is also used as the selectivity value when the value of an equality search arguments is not known at run time.
Range Selectivity	Prints the default value of .33, unless the value has been updated using the optdiag input mode. This is the value used for range queries if the open ended range search argument (<, <=, >, >=) is not known at optimize time.
In Between Selectivity	Prints the default value of .25, unless the value has been updated using the optdiag input mode. This is the value used for range queries if the closed range search argument (between, col >= X and col <= Y) is not known at optimize time.

Index entity Association area

You can view general details on the entities associated with the selected Index entity in the Association area. It is possible to associate to the Index entity from the following entities: Database, Table, Column, User, Segment, and Data Cache.

About the Column entity

The Column entity displays general details for the selected column. The following views are available:

- Overview
- Column Statistics

About getting an overview of the selected column

Displays general details on the selected column such as type, length, and default value.

About viewing column statistics

This view shows column-level statistics including:

- Statistics giving the density and selectivity of columns.
- A histogram, if the table contains one or more rows of data at the time the index is created or update statistics is run. There is a histogram for the leading column for:
 - Each index that currently exists (if there was at least one non-null value in the column when the index was created).
 - Any indexes that have been created and dropped (as long as delete statistics has not been run).
 - Any column list on which update statistics has been run. There is also a histogram for:
 - Every column in an index, if the update index statistics command was used.
 - Every column in the table, if the update all statistics command was used.

The following information is displayed:

Statistics Updated	The date and time that the last modification of the statistics occurred. This can be due to create index, update statistics, or by writing the statistics with optdiag.
Steps	The number of steps generated for the column.
Range Cell Density	A measure of the density of values in the column, after values in frequency count cells has been eliminated. Basically, this is the density of all values that are not highly duplicated. This value is used to estimate the cost of an equality search arguments when the value falls into a range cell.
Total Density	A measure of the duplicate values in all rows of the column, no matter what cell type they fall into. This value is used to estimate the number of rows that will be returned for a join on this column. It is also used as the selectivity value when the value of an equality search argument is not known at run time.
Range Selectivity	Prints the default value of .33, unless the value has been updated using the optdiag input mode. This is the value used for range queries if the open ended range search argument (<, <=, >, >=) is not known at optimize time.
In Between Selectivity	Prints the default value of .25, unless the value has been updated using the optdiag input mode. This is the value used for range queries if the closed range search argument (between, col >= X and col <= Y) is not known at optimize time.

Distribution

Shows the column histogram as the distribution of values within the column.

The histogram contains the step numbers, the weights, an operator used to help indicate the type of cell (Range / Frequency), and the value of the cell. The weight value indicates the percentage of rows in the column that are occupied by values in the cell.

For example, a weight of 0.050000 means that the value(s) in the cell occupy 5% of the rows in the column. The weight is very useful to determine when a value or values is becoming highly duplicated-if the weights are fairly uniform, the data is well distributed.

Column entity Association area

You can view general details on the entities that are associated with the selected Column entity in the Association area. It is possible to associate to the Column entity from the Table, Index, and View entities.

About the View entity and Views Contain Same Column entity

The View entity displays general details for the selected view.

The Views Contain Same Column entity displays views that have the same column as the column displayed in the Main area. The information displayed is the same as displayed in the View entity and can only be associated to from the Column entity.

About getting an overview of the View entity and Views Contain Same Column entity

The Overview displays general details on the View entity such as creation date, number of columns, and the view text.

View entity and Views Contain Same Column entity Association area

You can display general details on the entities that can be associated with the selected View entity in the Association area. It is possible to associate to the View entity from the Database, Table, Column and Index entities.

About the User entity

The User entity displays general details of the selected user, such as, Roles and Group Name.

User entity Association area

You can view general details on the entities that can be associated to the selected User entity in the Association area. It is possible to associate to the User entity from the Database entity.

About the Database Device entity

The Database Device entity displays general details, I/O Activity and storage information on the selected device. The following views are available:

- Overview
- Statistics

About getting an overview of the database device entity

The Overview displays the following information:

Properties	Displays general details on the device, such as full filename (including path), size, I/O Wait, and I/O Activity during the selected time frame.
Semaphores	Displays the number of times that a request for a device semaphore was NOT granted immediately (the semaphore was busy) and the task had to wait for the semaphore to be released. This data is meaningful for SMP applications only. When Adaptive Server needs to perform a disk I/O, it gives the task the semaphore for that device, to acquire a block I/O structure. On SMP systems, multiple engines can try to post I/Os to the same device simultaneously. This creates contention for that semaphore, especially if there are hot devices or if the data is not well distributed across devices. A large percentage of I/O requests that waited can indicate a semaphore contention issue. One solution might be to redistribute the data on the physical devices.
Pages requested	Displays the I/O Activity (pages read and written) during the selected time frame for the selected device, as collected by the Collect Statistics process.
Space Allocation	Displays free space and used space (in MB) during the selected time frame for the selected device, as collected by the Collect Space Utilization process.

About viewing I/O statistics on a database device

Displays I/O statistics on the device during the selected time frame. The following information is displayed:

Semaphores	Displays the number of times that a request for a device semaphore was NOT granted immediately (the semaphore was busy) and the task had to wait for the semaphore to be released. This data is meaningful for SMP applications only. When Adaptive Server needs to perform a disk I/O, it gives the task the semaphore for that device, to acquire a block I/O structure. On SMP systems, multiple engines can try to post I/Os to the same device simultaneously. This creates contention for that semaphore, especially if there are hot devices or if the data is not well distributed across devices. A large percentage of I/O requests that waited can indicate a semaphore contention issue. One solution might be to redistribute the data on the physical devices.
Pages requested	Displays the I/O Activity (pages read and written) during the selected time frame for the selected device, as collected by the Collect Statistics process.

Database Device entity Association area

You can view general details on the entities associated to the Database Device entity in the Association area. It is possible to associate to the Database Device entity from the following entities: Instance, Database, Segment, Table, and Index.

About the Segment entity

The Segment entity displays general details and space utilization during the selected time frame, for the selected Database Segment.

About getting an overview of segment details

The Overview displays general Segment details, such as number of device extents the Segment occupies and information regarding the segment's space utilization during the selected time frame. The following information is displayed:

Space Allocation	Displays the amount of free space and used space of the selected segment during the selected time frame, based on the data collected by the Collect Space Utilization process.
-------------------------	--

Segment entity Association area

You can view general details on the entities that are associated to the selected Segment in the Association area. It is possible to associate to the Segment from the following entities: Database, Table, Index, and Database Device.

About the Data Cache entity

The Data Cache entity displays general information about the selected Data Cache. The following views are available:

- Overview
- Data cache pools

About getting an overview of data cache pools

The Overview displays general Data Cache details, such as size, replacement policy and pool configuration.

About viewing data cache pools

Data Cache Pools can be configured in each cache to enable Sybase to perform large I/Os. The Data Cache Pools view displays information regarding the pools configuration of the selected Data Cache.

The following information is available:

I/O Size	Shows the size of the buffers in the pool. Specifies the amount of bytes that can be read at once. The default size of the pool is the size of the server's logical page.
Config Size	The amount of memory space configured for the pool.
Run Size	The amount of memory space currently assigned for the pool. The value may differ from the config size of the pool because one cannot explicitly configure its size. The value may differ for other pools if you have tried to move space between them, and some of the space could not be freed.
Wash Size	Size of the Wash area. A portion of each pool is configured as the wash area. After dirty pages (pages that have been changed in cache) pass the wash marker and enter the wash area, Adaptive Server starts an asynchronous I/O on the page. When the write completes, the page is marked clean and remains available in the cache. The space in the wash area must be large enough so that the I/O on the buffer can complete before the page needs to be replaced. For more information, see the <i>Precise Installation Guide</i> .
APF (%)	Displays the percentage of the pool that can hold unused buffers brought in by asynchronous prefetch.

Data Cache entity Association area

You can view general details on the selected Data Cache entity in the Association area. It is possible to associate to the Data Cache entity from the Instance, Database, Table and Index.

About the Unused Indexes entity

The Unused Indexes entity displays all indexes that are not used by any explained statement that is associated to a database, table, user or segment.

About the Statement entity

The Statement entity displays statements associated to tables and indexes in the Object tab. The following views are available:

- Overview
- Performance
- SQL text

About getting an overview of a statement

The Overview displays the following information:

Statement ID	Displays the statement ID.
Database	Displays the database name.
Parsing User	Displays the parsing user name.
Most Recent Plan	Displays the date of the most recent plan.
Last Show Plan	Displays the date of the last shown plan.
Access Plan	Displays a list of all Access operators in the statement's explain results.

About viewing the performance of a selected statement

The Performance view displays the following information:

In Sybase graph	Buffer Wait
	Network I/O Wait
	Log Full Wait
	Log Wait
	Internal Wait
	Remote Wait
	Lock Wait
	I/O Wait
	Using CPU
Statement Executions graph	Number of executions overtime.

About viewing the SQL text of a specific statement

The SQL Text view displays the SQL text of the statement.

Statement entity Association area

You can view general details on the statement associated with the selected Table, or Index entity in the Association area. It is possible to associate from the statement to the Table or Index accessed by this statement. There are Table, Index, and Object lists which can be associated to a statement.

The table below describes the information displayed in the Performance tab.

Table 7-2 Performance tab

Column	Description
Launch icon to SQL tab	Launches into SQL tab.
Statement ID	Displays the statement ID.
In Sybase	Displays the sum of each In Sybase parameter in the time period from the statistics table - bar with separated values or sum of all In Sybase values.
Text	Displays the statement's text.

The table below describes the information displayed in the Plan tab.

Table 7-3 Plan tab

Column	Description
Launch icon to SQL tab	Launches into SQL tab.
Statement ID	Displays the statement ID.
Database	Displays the database name.
Parsing User	Displays the parsing user name.
Most Recent Plan	Displays the date of the most recent plan.
Last Show Plan	Displays the date of the last shown plan.
Total Estimate Cost	Displays a graph/number.
Table Scan	Displays Yes or No.
Clustered Index Scan	Displays Yes or No.
Clustered Index Seek	Displays Yes or No.
Index Scan	Displays Yes or No.
Index Seek	Displays Yes or No.
Table Lookup	Displays Yes or No.
Sort	Displays Yes or No.
Merge Join	Displays Yes or No.
Join	Displays Yes or No.
Nested Loop Join	Displays Yes or No.
Parallel Access	Displays Yes or No.

About the Batch entity

The Batch entity displays batches associated to tables and indexes in the Object tab. The following views are available:

- Overview
- Performance
- SQL text

About getting an overview of batch details

The Overview displays general batch details such as most recent plan and last show plan, as follows:

Batch ID	Displays the batch ID.
Database	Displays the database name.
Parsing User	Displays the parsing user name.
Most Recent Plan	Displays the date of the most recent plan.
Last Show Plan	Displays the date of the last shown plan.
Access Plan	Displays a list of all Access operators in the batch's explain results.

About viewing the performance of a batch

The Performance view displays resource consumption and number of statement executions for a batch, as follows:

In Sybase graph	Buffer Wait
	Network I/O Wait
	Log Full Wait
	Log Wait
	Internal Wait
	Remote Wait
	Lock Wait
	I/O Wait
	Using CPU
	Statement Executions graph

About viewing the SQL text of a batch

The SQL text view displays the SQL Text of the batch.

Batch entity Association area

You can view general details on the batch associated with the selected Table, or Index entity in the Association area. It is possible to associate from the batch to the Table or Index entity accessed by this batch. There are Table, Index, and Object lists that can be associated to a batch.

The table below describes the information displayed in the Performance tab.

Table 7-4 Batch entity Performance tab

Column	Description
Launch icon to SQL tab	Launches into SQL tab.
Batch ID	Displays the batch ID.
In Sybase	Displays the sum of each In Sybase parameter in the time period from the statistics table - bar with separated values or sum of all In Sybase values.
Text	Displays the text of the batch.

The table below describes the information displayed in the Plan tab.

Table 7-5 Batch entity Plan tab

Column	Description
Launch icon to SQL tab	Launches into SQL tab.
Batch ID	Displays the batch ID.
Database	Displays the database name.
Parsing User	Displays the parsing user name.
Most Recent Plan	Displays the date of the most recent plan.
Last Show Plan	Displays the date of the last shown plan.
Total Estimate Cost	Displays a graph/number.
Table Scan	Displays Yes or No.
Clustered Index Scan	Displays Yes or No.
Clustered Index Seek	Displays Yes or No.
Index Scan	Displays Yes or No.
Index Seek	Displays Yes or No.

Column	Description
Table Lookup	Displays Yes or No.
Sort	Displays Yes or No.
Merge Join	Displays Yes or No.
Join	Displays Yes or No.
Nested Loop Join	Displays Yes or No.
Parallel Access	Displays Yes or No.

How the Objects tab can help you identify performance problems

It is possible to use the Objects tab to identify performance problems by:

- See [Examining an object's space usage over time](#) on page 80.
- See [Examining fragmentation of table data and indexes](#) on page 80.
- See [Examining distribution statistics of indexes and columns](#) on page 80.

Examining an object's space usage over time

Space usage information, collected by the Collect Space Utilization process, lets you examine the space allocated versus the space used, in a table and index. This enables you to easily keep track of object growth over time and decide how to spread database segments over various devices and how to connect objects to segments in the specified database.

Examining fragmentation of table data and indexes

Fragmentation of table data and indexes may affect the performance of your system. The fragmentation occurs as a result of data modification (`INSERT`, `UPDATE`, and `DELETE` statements) performed on the table. Fragmentation increases the amount of page reads performed by the query, increasing their resource consumption.

To examine fragmentation of table data and indexes

1. Use the Optdiag view in the Table or Index entity to observe fragmentation information.
2. Determine whether the table or index is heavily fragmented and needs to be reorganized. See [About the Table entity and Tables Contain Same Column entity](#) on page 6974.

See [About the View entity and Views Contain Same Column entity](#) on page 74.

Examining distribution statistics of indexes and columns

Density and selectivity statistics of an index can help you determine whether or not the index is useful to the query optimizer. The information displayed in the Statistics view of the Index entity indicates the selectivity of an index; the more selective an index is, the more useful it is, because this means more rows can be eliminated from consideration. You can also determine what columns in the index contribute to the selectivity and thus decide whether all should be included in the index.

To examine distribution statistics of indexes and columns

1. Use the Statistics view of the Column entity to help you see the distribution of values in that column. See [About viewing density and selectivity statistics](#) on page 73.

See [About viewing column statistics](#) on page 73.

Examining Sybase instance statistics

This section includes the following topics:

- [About the Statistics tab](#)
- [How the Statistics tab is structured](#)
- [About the entities you can examine in the Statistics tab](#)
- [How the Statistics tab can help you identify performance problems](#)

About the Statistics tab

The Statistics tab displays statistical information on all Sybase instances. The tab can be used to monitor your system's historical statistical information.

The Statistics tab lets you provide answers to the following types of questions: "Does Sybase have too many or too few Adaptive Server engines?" or "There are a lot of task context switches in the system that affect performance. What is the main reason for those context switches?"

The Statistics tab provides hundreds of raw performance counters. It is possible to view status information in 15 minute intervals or historical information roll-up at a higher summary level.

Examining the ready-to-use set of graphs that display several related counters is more important than examining a single counter. This is now possible using the Statistics tab.

Use the Statistics tab to periodically monitor the health of your system - for example, view the statistics on CPU Utilization or Hit Ratios. Alternatively, you can use the Statistics tab to fully analyze a performance problem reported by the Collector agent.

NOTE Statistics tab data is based on the results of the statistics counters that are gathered by ASE. These counters are cleared when using `sp_sysmon` to view statistics counters. If `sp_sysmon` is run, statistics data gathered by Precise for Sybase will not be complete. This problem can be avoided by not running `sp_sysmon` or by using the `noclear` option of `sp_sysmon` (available in Sybase version 12.5.3 and higher).

How the Statistics tab is structured

The Statistics tab displays information on a selected entity and its associated entities. When you open the Statistics tab from the Dashboard tab or the Activity tab, the tab is launched, in-context, and depending upon your selection in the All Sybase Instances table (All or Instance), the information is displayed on the Tier or Instance level. Navigating between these tabs and the Statistics tab will change the chosen time frame. If you open the Statistics tab from another tab, the historical settings (meaning those settings which were selected when you left the tab, such as the last entity you drilled down to) are taken into account and the information displayed the last time you viewed this tab is displayed (similar to clicking the History control and returning to a previous tab).

The selected entity is always reflected in the Tab heading, which serves as a point of orientation. The highest-level entity you can view information for in the Statistics tab is the Tier level. You can select an instance from the Instance list.

The entities displayed in the Association area are related to the selected entity displayed in the Main area.

About the Main area in the Statistics tab

The Main area shows comprehensive information on the selected entity. You can choose from several views to examine the entity from different angles. For example, the Instance entity reports on Disk I/O, Memory and Lock Management.

About the Association area in the Statistics tab

The Association area displays relevant information on the entities associated with the selected entity (displayed in the Tab heading in the Main area) in a table format. For example, it is possible to associate to all Engines or Devices that are related to a specific Instance, by selecting an entity from the Association controls list. Notice that some entities, show additional information on the associated entities in different tabs. The tabs are located above the Association area table. Clicking on a tab displays different table columns showing different information for the associated entities.

If you want to view additional information on an associated entity, drill down to it, by selecting the entity's row. A drilldown affects the entire tab. When you drill down to another entity, the Tab heading reflects the new selection, the Main area displays information on the selected entity, and the Association area displays the entities that are now associated with the newly selected entity.

For example, when you want to view information on a specific device, choose Devices from the Association controls. The Association area changes to display information on the devices associated with the selected Instance. Note that the Tab heading and the Main area remain unchanged.

In the Association area, click the row of the device you want to view detailed information for. The Tab heading indicates the newly selected entity; the Main area displays statistical information on the device entity you drilled down to, and the Association area shows information on the associated performance counters.

About the entities you can examine in the Statistics tab

The Statistics tab displays information on the following different entities in the Association area, as described in the following sections:

- About the Tier entity
- About the Instance entity
- About the Engine entity
- About the Device entity
- About the Data Cache entity
- About the Performance Counter or Counter Instance entity
- About the Wait Events entity
- About the Replication Server entity
- About the Replicate(d) Databases entity

About the Tier entity

Choose All Instances from the Instances List to display information on the Tier entity. The Tier entity provides a view of all Sybase instances in the application.

The following views are available:

- Overview
- Instance Grouping

About viewing a summary of statistical information on top instance consumers

The Overview provides statistical information on the top six instance consumers of various resources and enables you to compare them.

The Tier entity shows the top six consumers as follows:

Top 6 instances (% Engine Utilization)	Displays the top six instances for which the engine utilization was highest within a certain time frame.
Top 6 instances (Committed Transactions)	Displays the top six instances with the highest number of committed transactions within a certain time frame.
Top 6 instances (Connection Opened – Sum)	Displays the top six instances with the most connections opened within a certain time frame.

About viewing performance counters according to instance groups

The Instance Grouping view displays performance counters broken down into instance groups (the instance groups are defined by the user).

For an explanation on how to define instance groups, see the *Precise Installation Guide*. The table below describes the information displayed in the Instance Grouping table.

Table 8-1 Instance Grouping table

Item	Description
Group	Displays the name of the group.
Instances	Displays the number of instances linked to the group.
Engine Utilization	Displays the average Engine Utilization for the instances in the specified group.
Committed Transactions	Displays the SUM number of Committed Transactions for the instances in the specified group.
Connections Opened (Sum)	Displays the sum of opened connections for the instances in the specified group.

About viewing instances associated with an Tier in the Association area

Displays information on the Instances associated with the selected Tier in the Association area. The table below describes the information displayed in the tabs in the Association area.

Table 8-2 Viewing information on associated instances

Tab	Description
Overview	Displays counters reporting on CPU Usage, Memory Usage and amount of Locked Data.
Network	Displays counters reporting on the number of times an engine polls incoming and outgoing packets and bytes.
Disk I/O	Displays various counters reporting on the amount of disk I/O.
Transactions	Displays counters reporting on the number of rows changed in the instance, grouped by type (insert, update, and delete).

About the Instance entity

The instance entity displays statistical information for the entire instance. Precise for Sybase collects hundreds of counters that are relevant to the entire instance. The Instance entity displays predefined graphs that display several counters in each. Each graph contains counters that have some relation to each other and can indicate where performance problems lie.

The following views are available:

- Overview
- Disk I/O
- Data cache
- Memory
- Network
- Lock management
- Transaction log
- Transaction profile
- HouseKeeper management
- Wait events
- Replication Inbound
- Replication Outbound

About getting an overview of an instance based on statistical data

The Overview displays predefined overtime graphs, which can indicate where performance problems lie. The predefined graphs contain counters reported by Sybase and collected by the Precise for Sybase Collector agent.

The following information is displayed:

In Sybase	Displays resource consumption counters collected by the Precise for Sybase Collector agent, such as Using CPU, I/O Wait, and Lock Wait.
CPU Yields	Displays information on how often the engines yielded to the Operating System.
Engine Utilization	Displays the percentage of time that the Adaptive Server Kernel is busy executing tasks.
Connections Opened	Displays the number of opened connections (max and sum).
Task Context Switches - Reasons	Displays the number of times each Adaptive Server engine switched context from one user task to another. Possible reasons are: a Network Packet sent and received, modify conflicts, last log page, writes group, committed sleeps, and so on.

About viewing Disk I/O activity

The Disk I/O view displays statistics that report on disk I/O activity, such as Percentage of Checks Returning I/O, number of Pages read and written, and so on.

The following information is displayed:

Checks Returning I/O Percentage	Displays the percentage of checks that returned an I/O out of the total number of performed checks.
Page Requests	Displays the number of pages read synchronous or asynchronous and the number of pages written.
Avg Disk I/O Returned	Displays the average number of engine checks that returned an I/O.
Disk I/O Delays	Displays the disk I/O delays due to the following reasons: the engine limit, the OS limit or the ASE limit.

About viewing Data Cache activity

The Data Cache view displays statistics that report on data cache activity, such as Hit Ratio, Utilization, and so on.

The following information is displayed:

Hit Ratio	Displays the ratio between the number of times a page was found in the cache and the number of times caches were searched for a specific page.
Utilization	Displays the number of times a requested page was found and not found in all caches (hits) and needed to be read from disk.
Wash Behavior	Displays the number of buffers that were clean when they passed the wash marker, the number of times that the I/O was already active on a buffer when it entered the wash area, and the number of times that a buffer entered the wash area dirty and was not already in I/O.
Dirty Pages Requests	Displays the average number of pages requested at isolation level 0 for a cache.

About viewing Memory activity

The Memory view displays statistics that report on memory activity, such as number of Memory Requests, Memory Pages, Worker Processes, and so on.

The following information is displayed:

Memory Requests	Displays the number of successful/failed requests for memory allocation for the worker processes.
Worker Process Requests	Displays the number of granted and denied worker process requests.
Memory Pages	Displays the number of times that a new page was allocated/released in the memory.
Avg Memory Ever Used by Worker Process	Displays the average memory used by the worker process.

About viewing Network activity

The Network view displays statistics that report on network activity, such as number of Packets, Bytes, Checks, and so on.

The following information is displayed:

Packets	Displays the total number of packets received and sent.
Checks	Displays the number of times ASE performed blocking and non-blocking network checks.
Bytes	Displays the number of bytes received from the network and sent to the network
I/O Delays	Displays the total number of times the network I/O was delayed.

About viewing Lock Management activity

The Lock Management view displays statistics that report on lock management activity, such as Lock Contention, Lock Waits, and so on.

The following information is displayed:

Lock Contention Percentage	Displays the number of times that there were different lock contentions as a percentage of the total number of lock requests.
Lock Waits	Displays the number of times that there was a lock contention, divided by types.

About viewing Transaction Log activity

The Transaction Log view displays statistics that report on transaction log activity, such as Log Writes, Log Allocations, and so on.

The following information is displayed:

ULC Flushes to Log	Displays the number of times that the ULC was flushed to the transaction log, because of the following reasons: the ULC was full, a transaction ended, the database was changed, or a system transaction occurred within the user transaction.
Log Writes	Displays the number of times the ASE server writes transaction log pages to disk.
Log Allocations	Displays the number of times additional pages were allocated to the transaction log.
Avg Writes per Log Page	Displays the average number of times each log page was written to disk.

About viewing a Transaction Profile

The Transaction Profile view displays statistics that report on the number of Inserts, Updates, and Deletes. The following information is displayed:

Inserts	Displays the number of rows inserted into heap tables, clustered tables or data-only -lock tables.
Deletes	Displays the number of rows deleted on the deferred mode, directly or from DOL tables.
Updates	Displays the number of rows updated in the deferred mode, directly in-place, directly not in-place, directly expensive or deferred mode in DOL tables.
DOL Updates	Displays the number of rows replaced in the DOL tables, rows shrunken while updated in the DOL tables, rows expanded while updated in the DOL tables, rows expanded and forwarded while updated in the DOL tables or rows that were forwarded already and now fit in original page and returned to the original page in the DOL tables.

About viewing Housekeeper activity

The Housekeeper Management view displays statistics that report on housekeeper activity, such as Buffer Cache Wash, Garbage Collection, and so on. The following information is displayed:

Buffer Cache Washes	Displays the number of times that the housekeeper activity performs buffer cache washes and find it clean or dirty.
Garbage Collections	Displays the number of times the housekeeper garbage collection activity checks for space that can be reclaimed.
Pages Processed in the Garbage Collection	Displays the number of pages reclaimed by the housekeeper garbage collection activity.
Statistics Updates	Displays the number of times the housekeeper chore activities checked to see if statistics needed to be written.

About viewing Wait events

The Wait Event view displays statistics on the top ten wait events for the selected instance.

About viewing Replication Inbound view graphs

The Replication Inbound view is a combination of a few predefined graphs related to the inbound processing phase (replication agent); the view helps you identify what may be causing the latency in the replication system. The information displayed in the Instance level is a summation of all the replicated databases of the instance.

Table 8-3 I/O vs Cmds Sent

Graph	Description
IO vs. Cmds Sent	<p>Highlights</p> <p>The IO vs. Cmds Sent graph shows the number of commands processed by the RepAgent during the given period and the amount of time the RepAgent had to wait while writing to the inbound queue. Each bar represents a time period. The speed of the RepAgent is directly proportional to the speed of ASE processing the network that sends requests coupled with the speed of the Replication Server processing.</p> <p>What to do next</p> <ul style="list-style-type: none"> • If the amount of I/O wait is significant (compared to the total time each bar represents), then the specific issue delaying the RepAgent must be isolated. The amount of time a RepAgent will have to wait while writing to the inbound queue depends on how much cache is available (exec_sqm_write_request_limit) as well as the values for init_sqm_write_delay/init_sqm_write_max_delay. For more information, check the online help available on the Sybase Web site. • If you see a shift in the trend of the Cmds Sent (fewer commands are sent even though the activity is the same), then perhaps something has changed in the ASE processing. Examine the Activity tab and see what was the ASE doing at that time and what, if anything, was holding it back.

Table 8-4 Truncation Point Movement

Graph	Description
Truncation Point Movement	<p>Highlights</p> <p>The Truncation Point Movement graph shows the number of times the RepAgent has asked the RS for a new secondary truncation point and then moved the secondary truncation point in the log. If 'Moved' is more than one less than 'Gotten', the likely cause is that a large or open transaction exists from the Replication Server perspective (either it is indeed still open in ASE or the RepAgent has not forwarded the commit record yet).</p> <p>What to do next</p> <ul style="list-style-type: none"> • Check for an open transaction in ASE. • To determine if the numbers are high - you can gauge the number by dividing the 'Replicated Log Records' (which can be seen in the Log Scan Summary graph) by the RepAgent 'scan_batch_size' configuration parameter and add one to count for the additional move when it reaches the end of the log. The number should be close to what is shown.

About the entities you can examine in the Statistics tab

Table 8-5 Network

Graph	Description
Network	<p>Highlights</p> <p>The Network graph shows information related to the network activity. If the Number of Full Packets is a considerable percentage out of the total package sent ($\geq 90\%$), then you may need to consider increasing the <code>send_buffer_size</code>.</p> <p>What to do next</p> <ul style="list-style-type: none"> Examine the <code>send_buffer_size</code> configuration parameters in the Objects tab.

Table 8-6 Log Scan Summary

Graph	Description
Log Scan Summary	<p>Highlights</p> <p>The Log Scan Summary graph is an indicator of how much work the RepAgent is doing and how much information is being sent to the Replication Server. Replicated log records are records converted into LTL and sent to the RS. Not replicated log records are all log records that were scanned but not sent.</p> <p>What to do next</p> <ul style="list-style-type: none"> If the replicated log records number is small compared to the total log records scanned, it means that most of the transactions in the log are not aimed for replication — so a large part of the log is scanned needlessly. Consider moving the few tables that need to be replicated to another database so that the total log records scanned will be smaller.

Table 8-7 Replicated Commands Breakdown

Graph	Description
Replicated Commands Breakdown	<p>Highlights</p> <p>The Replicated Commands Breakdown lists the number and type of commands that were replicated and which most affect the replication.</p> <p>DDL – refers to DDL statements that were replicated. Generally this number should be zero, with only minor increases in a Warm Standby when DDL changes are made.</p> <p>Writetext – shows how many Writetext operations are being replicated. DML with Text/Image – displays how many row images are processed. What to do next</p> <ul style="list-style-type: none"> If a large number of text rows are being replicated, you may want to investigate whether a text/image column was inappropriately marked or left at "always_replicate" instead of "replicate_if_changed".

About viewing Replication Outbound view graphs

The Replication Outbound view graphs view is a combination of a few predefined graphs related to the outbound processing phase. This enables you to understand what specific component(s) may be causing the latency in the replication system. The information displayed in the instance level is a summation of all the outbound processing performed for the specified instance by all the Replication servers attached to it.

Table 8-8 SQM Processing

Graph	Description
SQM Processing	<p data-bbox="532 443 634 468">Highlights</p> <p data-bbox="532 489 1398 577">The SQM Processing graph shows the space usage of SQM queue overtime versus the number of commands written to the queue. The SQM Commands Written indicates the amount of activity performed by the replication server.</p> <p data-bbox="532 598 691 623">What to do next</p> <ul data-bbox="581 644 1341 1157" style="list-style-type: none"> <li data-bbox="581 644 1341 732">• Check the percentage used of the SQM and see if it is sufficient, over-allocated, or near maximum space capacity to accommodate the written commands. <li data-bbox="581 743 1341 884">• Examine the trend of written commands, with special attention to trend changes. Is the activity evenly distributed, or does it peak at certain times? A deviation from normal trends can be acceptable, but it may also indicate a replication activity that is unauthorized or poorly defined. <li data-bbox="581 894 1341 982">• If the number of commands are lower, then maybe there is a problem in inbound queue processing or in the network. Check the replication inbound statistics. <li data-bbox="581 993 1341 1157">• If the amount of free space is lower and there is no change in the number of written commands, this may indicate that the subscriber was unable to apply the transaction. For example, an insert statement may have failed because of a duplicate key. If this example occurs, the truncation point will indicate this failed transaction and will not proceed until it is cleared manually.

Table 8-9 DSI SQT Processing

Graph	Description
DSI SQT Processing	<p data-bbox="532 247 634 275">Highlights</p> <p data-bbox="532 296 1349 323">The DSI SQT Processing graph shows space usage of the outbound queue cache.</p> <p data-bbox="532 344 1422 464">The Cache size lists the size of the SQT cache. Cached Used shows SQT thread memory usage over time. Each command structure (allocated by an SQT thread) is freed when its transaction context is removed. Consequently, if no transactions are active in SQT, then SQT cache usage is zero.</p> <p data-bbox="532 485 1422 575">If Cached Used is near the maximum cache size capacity and Trans Removed is constantly greater than zero, you may decide to increase the SQT cache size by increasing <code>dsi_sqt_max_cache_size</code>.</p> <p data-bbox="532 596 1406 751">Don't increase the cache size if Trans Removed is occasionally greater than zero; this situation usually means that a transaction was removed from the cache because it was greater than the <code>dsi_sqt_max_cache_size</code>. For occasional increases of Trans Removed, increasing the cache size will have the opposite effect because the latency in processing transactions ahead of them will likely result in their being removed.</p> <p data-bbox="532 772 1422 892">The SQM Read Cached indicates how many 16K blocks of cache were read by the SQM Reader thread. This number should be as high as possible. If it is constantly hitting zero, most likely there is a latency in executing the SQL at the replicate; resulting in the DSI SQT cache to reach its maximum.</p> <p data-bbox="532 913 691 940">What to do next</p> <ul data-bbox="581 961 1341 1850" style="list-style-type: none"> <li data-bbox="581 961 1317 1073">• If Cached Used is near the maximum cache size capacity and Trans Removed is constantly greater than zero, you may decide to increase the SQT cache size by increasing <code>dsi_sqt_max_cache_size</code>. <li data-bbox="581 1094 1341 1205">• For occasional increases of Trans Removed, increasing the cache size will have the opposite effect because the latency in processing transactions ahead of them will likely result in their being removed. <li data-bbox="581 1226 1341 1463">• Examine <code>dsi_sqt_max_cache_size</code> and <code>sqt_max_cache_size</code>. The <code>sqt_max_cache_size</code> should usually be between 4-8 MB. For non-parallel DSI configuration and no Warm StandBy, the <code>dsi_sqt_max_cache_size</code> should be smaller than the <code>sqt_max_cache_size</code>. If set too high (or left at zero as default), the DSI-S thread will continuously try to fill the available DSI SQT cache from the outbound queue - often at the expense of yielding the CPU to the DSI EXEC. Consequently, the optimum size for <code>dsi_sqt_max_cache_size</code> is between 1-2 MB. <li data-bbox="581 1484 1300 1596">• For Warm StandBy Configurations, the DSI SQT cache should normally be equal to the SQT cache or higher because it is the DSI SQT thread that is actually sorting the transactions into commit order. <li data-bbox="581 1617 1341 1850">• For parallel DSI configurations, the DSI SQT cache should contain double the <code>dsi_max_xacts_in_group</code> transactions for each DSI EXEC thread. The DSI EXEC thread can effectively process many row modifications because the load can be distributed among several available DSIs. This can result in a situation where the DSI transaction rate is higher than the amount of rows read from the outbound queue. In these situations, raising the DSI SQT cache allows the DSI to "read ahead" into the queue and begin preparing transactions before they are needed.

Table 8-10 DSIE Transactions Processing

Graph	Description
DSIE Transactions Processing	<p data-bbox="532 247 634 275">Highlights</p> <p data-bbox="532 296 1425 352">The DSIE Transactions Processing graph shows how long it took to process a transaction by a DSI/E thread breakdown to various phases.</p> <ul data-bbox="581 380 1344 1373" style="list-style-type: none"> <li data-bbox="581 380 1344 541">• FS Map Time: The amount of time used to translate the replicated row functions into SQL commands. If a large amount of time is spent in this area, this may indicate that there are relatively large customized function strings; if true, there may not be any options you can take. However, you may want to check that the STS cache is sized appropriately. <li data-bbox="581 562 1344 667">• Sent Time: This represents the amount of time spent sending the command batch to the replicate data server. A long lag time may indicate inefficient batching, or a slow response to client applications from the replicate server. <li data-bbox="581 688 1344 745">• SendRPCTimeAvg: This records the time spent sending RPCs to the RDB. <li data-bbox="581 766 1344 823">• SendDTTimeAvg: This records time spent sending LOB data to the RDB. <li data-bbox="581 844 1344 1005">• Result Time: This calculated value can be used to determine the amount of time spent processing results from the replicate server. This value also includes the execution time, because RS does very little result processing. Often, these metrics will be among the highest. You need to speed-up the replicate DBMS to improve the RS throughput. <li data-bbox="581 1026 1344 1157">• Commit Seq. Time: The amount of time spent waiting to commit. Here too, a high value may indicate a near-serial dsi-serialization_method, such as wait_for_commit. Alternatively, it may indicate a contention within the replicate server, possibly within the rs_threads_group. <li data-bbox="581 1178 1344 1308">• Batch Seq. Time: This is the time spent trying to coordinate the sending of the first batch in parallel DSIs. A long time may indicate that the dsi_serialization_method is wait_for_commit and a previous transaction is running a long time, or that the DSI thread is simply too busy to respond to the Batch Sequencing message. <li data-bbox="581 1329 1344 1373">• Other Time: The execution time by the replicate database outside of the replication server. <p data-bbox="532 1394 695 1421">What to do next</p> <ul data-bbox="581 1442 1344 1694" style="list-style-type: none"> <li data-bbox="581 1442 1344 1547">• If Other Time is the largest chunk of time out of the total transaction time, tuning RS will not help. You must do one of the following: tune the replicate database by using parallel DSIs, or use minimal columns/repdefs to reduce the SQL execution time. <li data-bbox="581 1568 1344 1625">• If FS Map Time is high, consider rewriting the queries and simplifying them, or work with a stored procedure. <li data-bbox="581 1646 1344 1694">• If Commit Seq. Time is high, consider lowering the number of parallel threads.

Table 8-11 DSIE Commands Applied

Graph	Description
DSIE Commands Applied	<p>Highlights</p> <p>The DSIE Commands Applied graph shows how many commands were successfully applied to the target database by DSI/E overtime (relevant for version 15 and higher). It can be used to measure the amount of work imposed by the replication server. You can also compare that number to the SQM Commands Written and check if there is a lag in processing commands inside the replication server.</p> <p>What to do next</p> <ul style="list-style-type: none"> • Study the trend of applied commands. If it has suddenly dropped, this may indicate a problem in the subscriber ASE. • Go to the Activity tab and examine the resulting batches from replication (rep server program).

About the Engine entity

The Engine entity displays predefined graphs that shows several counters in one. This enables you to immediately view relevant counters, grouped according to topic.

The following views are available:

- Overview
- Network
- Disk I/O

About getting an overview of CPU yields and engine utilization

The Overview displays predefined overtime graphs, which represent CPU Yields and Engine Utilization. The following information is displayed:

CPU Yields	Displays information about how often the engine yielded to the operating system.
Engine Utilization	Displays the percentage of time that the Adaptive Server Engine is busy executing tasks.

About viewing Network activity

The Network view displays predefined overtime graphs, which represent the number of Bytes sent and received, and Packet transportation.

The following information is displayed:

Packets	Displays the total number of packets received and sent by the engine.
Bytes	Displays the total number of bytes received and sent by the engine.

About getting an overview of Disk I/O

The Disk I/O view displays predefined overtime graphs, which represent Checks Returning I/O and Average Returned Disk I/O.

The following information is displayed:

Checks Returning I/O	Displays the percentage of I/O checks that found an I/O request completed.
Avg Disk I/Os Returned	Displays the average number of engine checks that returned an I/O.

About the Device entity

The Device entity displays predefined graphs that report on device activity and performance. The following views are available:

- Overview
- Semaphores

About getting an overview on a Device activity

The Overview displays an overtime graph, which represents Page Requests (read and write) of the device. The following information is displayed:

Page Requests Displays the number of pages read synchronous or asynchronous and the number of pages that are written.

About viewing Semaphores

The Semaphores view displays an overtime graph, which represents Semaphore Hits and Misses. The following information is displayed:

Semaphores Displays the number of times that a request for a device semaphore was granted immediately and the number of times the semaphore was busy and the task had to wait for the semaphore to be released.

About the Data Cache entity

The Data Cache entity displays predefined graphs that report on the activity of the Data Cache. The following views are available:

- Overview
- Cache details

About getting an overview on Data Cache activity

The Overview displays statistics that report on the Data Cache Hit Ratio and Utilization. The following information is displayed:

Hit Ratio Displays the percentage of time a requested page was found in a cache and was found in the wash area.

Utilization Displays the percentage of searches using this cache, as a percentage of searches across all caches.

About viewing Cache Details

The Cache Details view displays statistics that report on the Data Cache Buffer Wash (Pages) and large I/Os. The following information is displayed:

Buffer Wash (Pages) Displays the number of buffers that were clean when they passed the wash marker, the number of times that I/O was already active on a buffer when entered the wash area, and the number of times that a buffer entered the wash area dirty and was not already in I/O.

Large I/O Displays the number of times a large I/O was requested for a buffer and the number of times it was denied.

About the Performance Counter or Counter Instance entity

The Performance counter or Counter Instance entity displays statistical information on the selected counter. Precise for Sybase collects counters reported by Sybase. These counters determine which aspects of the objects to monitor. For example, the Sybase Memory Object contains the Mempages_allocated counter.

The Overview displays general details of the selected counter and overtime graph, which represent the counter value. In some of the selected counters, two overtime graphs are displayed: a graph that displays the number of transactions over time and another graph that displays the counter value per committed transaction.

About the Wait Events entity

Available for Sybase versions 12.5.0.3 and higher, the Wait Events entity displays the name of the wait event group, the number of waits experienced by this event group and the summed amount of time accumulated by wait events and summed occurrences.

The following view is displayed:

- Overview

About getting an overview of Wait Events

The Overview displays general details of the selected counter. The following information is available:

Event group	Displays the name of the wait type.
Event name	Provides a description of what is causing the event.
Time waited (avg)	Average time waited for wait events, over the selected time period.
Times waited graph	Displays the time waited for wait events vs. the time waited for group events.
Wait graphs	Displays the number of wait events vs. the number of wait events for group events.

About the Replication Server entity

The Replication Server entity displays predefined graphs which represent statistics related to the outbound processing phase and inbound (repagent statistics) processing phase for the selected Replication Server in the Selected time frame.

NOTE The Precise for Sybase Replication server collector agent collects statistics only for the Sybase ASEs which are monitored by Precise and use the specified replication server for replicating data. If the same RS also serves other ASEs that are not monitored by Precise or another RDBMS (such as: Oracle, SQL Server, etc.), the statistics of those instances will not be collected and displayed.

The following views are available:

- Replication Inbound
- Replication Outbound

See [About the Instance entity](#) on page 83.

About the Replicate(d) Databases entity

The Replicate(d) databases entity displays predefined graphs which represent statistics related to the outbound processing phase and inbound (repagent statistics) processing phase for the selected database in the Selected time frame.

NOTE The Precise for Sybase Replication server collector agent collects statistics only for the Sybase ASEs that are monitored by Precise and use the specified replication server for replicating data. If the same RS also serves other ASEs that are not monitored by Precise or another RDBMS (such as: Oracle, SQL Server, etc.), the statistics of those instances will not be collected and displayed.

The following views are available:

- Replication Inbound
- Replication Outbound

See [About the Instance entity](#) on page 83.

How the Statistics tab can help you identify performance problems

To determine whether Sybase is performing optimally there is a need to analyze performance measurements over time. The Statistics tab provides performance counters grouped into several predefined graphs that let you locate performance problems in your system (For example CPU, Paging, I/O, and Network Statistics). In addition, In Sybase states let you correlate between the performance counters and In Sybase states, to determine how an activity in the Sybase instance affects the system.

You can identify a performance problem by doing one or more of the following:

- See [Examining the kernel utilization of Sybase over time](#) on page 94.
- See [Examining I/O statistics](#) on page 94.
- See [Examining network activity](#) on page 95.
- See [Examining memory bottlenecks](#) on page 95.
- See [Examining a specific counter over time](#) on page 96.

Examining the kernel utilization of Sybase over time

You can examine a Sybase instance over time to confirm that the amount of Sybase engines configured in your system is optimal.

The Engine busy utilization graph in the Overview of the Instance entity reports the percentage of time the Sybase Kernel is busy executing tasks (rather than time spent idle). In other words, it measures how busy Sybase engines were during the time slice.

Viewing this graph in correlation with the CPU Yields graph can help you determine whether there are too many engines or if it is likely that response time and throughput can benefit from an additional engine.

The CPU Yields graph reports the number of times the engines yielded back to the Operating System.

If the Engine Busy Utilization graph data indicates a low engine utilization, use CPU Yields to determine whether the Engine Busy Utilization data reflects a truly inactive engine or one that is frequently starved out of CPU time by the Operating System.

When an engine is not busy, it yields to the CPU after a period of time. A high value for CPU Yields indicates that the engine yielded voluntarily.

The table below describes what the utilization status of an engine indicates.

Table 8-12 Utilization status

Engine Busy	CPU Yields	Status
Low	Low	Engine is CPU starved
Low	High	Engine is inactive
High	Low	Engine is busy
High	High	Engine is busy

To improve kernel utilization, it is recommended to check the Disk I/O and Network view. This is done to determine if the number of checks for I/O or Network that the engine performs, is ideal or are overhead.

For example, in the Disk I/O view, check the Checks Returning I/O Percentage graph. It reports the percentage that a Requested I/O had completed when an engine checked for Disk I/O. For example, if an engine checks for Expected I/O 100,000 times, this average indicates the percentage of time that there actually was I/O pending. If, of those 100,000 checks, I/O was pending 10,000 times, then 10% of the checks were effective, and the other 90% were overhead. However, you should also check the Average Disk I/O Returned graph and how busy the engines were during the sample interval.

If the sample includes idle time, or the I/O traffic is "bursty", it is possible that a high percentage of the checks were returning I/O during the busy period. If the results in this category seem low or high, you can configure the frequency of the checks.

Increasing the amount of time that Adaptive Server engines wait between checks may result in better throughput, because Adaptive Server engines can spend more time processing if they spend less time checking for I/O.

Examining I/O statistics

To examine I/O statistics

1. Open the Overview view of the specific instance entity.
2. Examine the I/O Wait value in the In Sybase graph to detect bottlenecks within the device subsystem.
3. Check for I/O Pacing and I/O Device Contention values in the Task Context Switches - Reasons graph.
4. If increased I/O Wait is the result of an activity in the instance, open the Activity tab to continue your analysis of the application's performance.

I/O Pacing reports how many times an I/O-intensive task switched off an engine, due to exceeding an I/O batch limit. This limit can be configured.

I/O Device Contention reports the number of times a task was put to sleep while waiting for a semaphore for a specific device. If there is significant contention for I/O Device Semaphores, try reducing it by redistributing the tables across devices or by adding devices and moving tables and indexes to them.

To determine which devices are having possible performance problems, drill down to the Device entity located in the Association area. In the Association area table you can see the Page Requests and the Sum of Semaphores for the different devices. Drilling down to one of the devices will enable you to see these counters over time.

How the Statistics tab can help you identify performance problems

In the Page Requests graph you can see the amount of I/O reads and writes.

High read requests may indicate the need for tuning a Select query. High write requests may indicate the need for tuning an Update query.

You can also use this information to check I/O distribution patterns over the disks and to make object placement decisions that can help balance Disk I/O across devices. For example, does the data show that some disks are more heavily used than others? If you see that a large percentage of all I/O went to a specifically named device, you can investigate the tables residing on the device and then determine how to remedy the problem.

Examining network activity

You can examine network traffic over a specified time period to detect bottlenecks. Network performance can be improved by one or more of the following:

- Enhancing the application
- Reducing the amount of data returned
- Configuring the network packet size parameter with a reasonable number that avoids many packets from being sent or received (causing overhead)
- Reconfiguring the network configuration

To examine network activity do one of the following

1. Examine the In Sybase graph and check the Network I/O Wait value.
2. Check the Task Context Switches - Reasons graph to see if Network Packet Sent is a major cause of task switching.
3. If so, check the applications on the problematic instance to see if it's possible to reduce the amount of data returned.

OR

1. Switch to the Network view.
2. Check the ratio between the Packets graph and the Bytes graph to see if the network packet size is used to its maximum capacity.
3. Consider configuring its size.

Examining memory bottlenecks

Lack of memory or a high memory consumption rate can cause paging. A page fault occurs when a process requires code or data that is not in its working set (its space in physical memory), thereby causing the page to be retrieved from disk and written to the physical memory. If there is not enough space in the physical memory, a chosen page needs to be written to disk to free space.

To examine memory bottlenecks

1. Examine the Memory view that reports the number of pages allocated and deallocated in each time slice.
2. If you spot a high value in one of the time slices, open the Activity tab to continue your analysis of the application's performance.
3. Switch to the Memory view to see information about the success and failure of Memory Requests for Worker Processed in the Memory Requests graph.
4. If some of the requests have failed you may need to increase the value of the `Memory per Worker Process` parameter.
5. Switch to Data Cache view to see information on how effective cache design is. A high Hit Ratio from the Hit Ratio graph and small amount of Page Request Misses from the Utilization graph may indicate effective cache design.
6. If you see a low Hit Ratio and a lot of Page Request Misses, investigate statistics for each of the Data Caches by selecting the Data Cache entity in the Association area.

Examining a specific counter over time

A wide range of performance counters generated by Sybase are collected by the Precise for Sybase collect instance statistics process. These performance counters are sampled every 15 minutes and displayed in this tab over time.

A performance group is a Sybase Performance Group resource, such as a Sybase Lock. Each performance group contains a few counters that determine which performance aspects to monitor. For example, the Sybase Memory Performance group contains counters called Pages Allocated (Sybase counter: `mempages_allocated`) and Pages Released (Sybase counter: `mempages_freed`).

Some performance counters have separate instances. An instance refers to a particular occurrence of a counter. For example, the Device-Counters Performance group has counters such as Disk Cache Misses (Sybase counter: `p_misses`) and Disk Cache Hits (Sybase counter: `p_hits`), but because several devices exist, each counter is divided into instances.

In relevant counters, the graphs appear with the counter value per transaction in addition to the counter value alone. In those counters, the number of transactions overtime graph appears as well.

To examine a specific counter

1. Select the associated required performance group from the Associations list.
2. Select the required counter.

Evaluating explain plans

This section includes the following topics:

- [About the SQL tab](#)
- [How the SQL tab is structured](#)
- [How the SQL tab can help you identify performance problems](#)

About the SQL tab

The SQL tab enables you to evaluate SQL statements and manage a database warehouse of your application's SQL statements together with their respective Explain Plans. The process of explaining statements is a prerequisite for tuning. The Explain process is designed to clarify the access path chosen for a statement and translate it into a visual medium.

After a statement is explained, Explain results are stored in the PMDB. This information includes the objects referenced by the statement and the operations performed on these objects. The statements are automatically explained every day. In addition, the top 10 resource-consuming statements during the last 4 hours, are automatically explained every 15 minutes.

Understanding the Execution Plan chosen by the Adaptive Server Query Optimizer is extremely important when tuning your application. You can ensure optimal system performance by ensuring that the best plans are used for your queries.

Precise for Sybase provides you with a special tab just for this purpose - the SQL tab. It provides several views of the same plan, making the process of analyzing long and complicated queries much easier to do.

To analyze the Plan, Precise for Sybase provides you with a full picture of the objects (table, indexes) participating in the Plan. You can change a plan by modifying your query or changing the schema by, for example, the addition of an index (outside of Precise for Sybase).

View the historical information of statements showing performance degradation to identify the source of the problem. A change in Execution Plan may explain the impact on performance.

NOTE The information described herein for this tab applies equally to statements and batches.

The table below shows from which tabs and entities it is possible to launch to the SQL tab, in context.

Table 9-1 Launching in context

Tab	Entities
Current	Session, Statement, Batch
Activity	Statement, Batch
SQL	Statement entered manually by the user through a New action.

NOTE If the SQL tab is opened with no statement in context, an error message is displayed requesting you open a statement.

See [About the Current tab](#) on page 40. See [About the Activity tab](#) on page 49. See [About the Objects tab](#) on page 67.

How the SQL tab is structured

The SQL tab lets you analyze execution plans and explain results so that you can tune statements and achieve optimal results.

Two sets of tabs control the information displayed in the SQL tab:

- View tabs
- Actions tab

About the View tabs

The SQL tab displays different information regarding the selected statement or batch, in different views. A statement is selected when you drill down to it in another tab (from Current or Activity tabs) or when you open a new or existing statement in the SQL tab. Click a View tab to display additional information regarding the statement or batch. Each view has a different layout.

The following View tabs are available in the SQL tab:

- Plan
- History
- Statements

About viewing the execution plan of a statement

The Plan view tab lets you display the execution plan of a statement and various related information such as statistics, referenced objects and operations performed in the Execution Plan.

The Plan view is divided into two panes. The Execution Plan tree is displayed in the left pane. The Details area is displayed in the right pane. The information displayed in the Details area is controlled by the information tabs (Highlights, Objects, Statistics, Operations, Properties) located above the Details area.

About the Execution plan tree

The Plan view displays the Execution Plan tree of the specified statement. When you analyze the access plan of a selected statement, you can examine the access path that was chosen by the Adaptive Server Query Optimizer. The Explain result is displayed as an Explain tree. Each branch in the Explain tree represents one operation. Each operation is displayed in the following format:

[<execution order>] <operation type> <accessed object>

It is also possible to display the Execution Plan's text, at the bottom of the Execution Plan tree. This allows you to view the text of the statement, the Execution Plan and additional information all at once, such as the objects referenced by the statement.

Actions that can be performed on the tree

It is possible to perform the following actions on the Execution Plan tree:

- The Playback controls, located at the top of the tree enable you to freely move within the Execution Plan of an explained statement.
- By moving the pointer over the Execution Plan steps, you can view a ToolTip that contains statistical information, such as Estimated Cost and Estimated Rows of the specified step.
- Selecting a specific step will highlight its sons and affect the information displayed in the information tabs.
- Clicking the plus sign (+) located at the bottom of the tree, displays the statement's text.

Available operation options

Precise for Sybase enables you to obtain a reasonable understanding of the Execution Plan selected by the Adaptive Server Query Optimizer, by displaying the textual output returned from the `SET SHOWPLAN ON` command in the Plan tree format. In addition to the Execution Plan, the Estimated I/O Costs returned from the `DBCC TRACEON (310, 311)` command are displayed in all the table-based operations.

The following options are available to different operations:

- Clustered index - the leaf pages of the index hold the actual data, instead of RIDs (Records ID). A table without a clustered index is called a Heap.
- Non-clustered index - the index consists only of the columns that combine the key and a pointer to the records in the table data. If the table has a clustered index, then the pointer is the key of the clustered index. If the table does not have clustered index, then the pointer is an RID.

There are several operations for accessing indexes or tables:

- Seeking - the index tree can be used for quickly locating the matching records. Seeking can only be done on indexes.
- Scanning - The leaf pages of the index or the data pages of the table are consecutively scanned.

NOTE Seeking is the recommended way to use indexes.

The table below describes the available operation options.

Table 9-2 Operation options

Operator	Description
Sort	The Sort operator presents the sorting of the rows returned from the previously executed operation (appears just below this operator) in a specific order, optionally eliminating duplicate entries. Sorting is required when no index that satisfies the requested ordering exists, or when an Index Scan is more expensive than sorting. It is usually used as the final step to retrieve the fetched data or a prior step for the Join or Stream Aggregate (Group By) operators. In some cases, the Sort operation requires more than the available memory, in which case a temporary table is used. This involves Disk I/O, which has an impact on performance. Sorting a large number of rows makes heavy use of the CPU, which can also affect overall server performance. The amount of memory needed for sorting can be estimated by multiplying the number of rows by the average row size. Therefore, minimizing the row size by selecting only the necessary fields will decrease the memory required for the Sort operation.
Nested joins	Nested Joins (also called Nested Loops) use the Nested algorithm to join two tables. The Nested algorithm works as follows: For each row in the outer table (the first table, appearing directly below the Nested Loop operator in the Explain tree), Sybase executes the Access Plan for the inner table (the second one). A large number of iterations may indicate a problem in choosing the Access Plan due to out-of-date statistics.
Merge joins	A Merge Join algorithm is used to join two tables by merging them. Merge Join algorithms read both tables only once; therefore it is required that the input be sorted according to the Join predicates. Merge Joins are available in Precise for Sybase from version 12.0 and above. This option is grayed-out by default. To enable it on session level, execute the <code>set sort_merge</code> on command. See the Sybase Help for more details on how to enable Merge Joins on server level.

About the operators that are available in the plan tree

The table below describes the operators that are available in the Plan tree.

Table 9-3 Operation options

Operator	Description
Step	This operator indicates the start of the specified step Subtree.
Table Scan	This operator indicates that a table scan was performed on the specified table to receive the required data. The More column displayed in the ToolTip and in the Statistics tab specify the scan direction (forward/backward), the I/O size and the buffer replacement strategy. When there is no useful index, the optimizer performs a table scan to receive the required data.
Index Scan	This operator indicates that an index scan was performed on the specified index to receive the required data. The More column displayed in the ToolTip and in the Statistics tab specify the scan direction (forward/backward), the I/O size and the buffer replacement strategy.
Index Seek	This operator indicates that an index seek was performed on the specified index to receive the required data. The More column displayed in the ToolTip and in the Statistics tab specify the scan direction (forward/backward), the I/O size and the buffer replacement strategy.
Clustered Index Scan	This operator indicates that a clustered index scan was performed on the specified index to receive the required data. The More column displayed in the ToolTip and in the Statistics tab specify the scan direction (forward or backward), the I/O size and the buffer replacement strategy.
Clustered Index Seek	This operator indicates that a clustered index seek was performed on the specified index to receive the required data. The More column displayed in the ToolTip and in the Statistics tab specify the scan direction (forward or backward), the I/O size and the buffer.
Table Lookup	This operator indicates that in addition to an access to the index pages, there is an access to the data pages of the table to receive the required data which is not included in the index itself. The More column displayed in the ToolTip and in the Statistics tab specify the I/O size and the buffer replacement strategy.
Nested Loops	This operator indicates that the join between the two result sets under this operator is of type Nested Loops. The first operator under the Nested Loops operator is the outer table in the join and the second operator under the Nested Loops operator is the inner table in the join.
Merge Join	This operator indicates that the join between the two result sets under this operator is of type Merge Join.
Sort	This operator indicates that a sort operation was performed on the rows returned from the previously executed operation (appears just below this operator) in a specific order, optionally eliminating duplicate entries.
Aggregate	This operator indicates that an Aggregate operation was performed on the rows returned from the previously executed operation (appears just below this operator). This operator appears in steps, which perform an evaluation of grouped or ungrouped aggregate functions such as <code>SUM</code> , <code>AVERAGE</code> and <code>COUNT</code> .
Subquery	This operator indicates that a subquery was executed. The Subquery tree will appear under this operator. The More column displayed in the ToolTip and in the Statistics tab specify the nesting level of the subquery, the predicate type (<code>IN</code> , <code>ANY</code> , <code>EXISTS</code> , or <code>EXPRESSION</code>), and the type of query (correlated/ non-correlated).
Spooling	This operator indicates that a temporary worktable was created to hold intermediate results. Temporary worktables are created in cases when the statement cannot be executed in one step. For example, if the statement contains an Order By clause and the optimizer cannot use an index to order the resulting rows, it creates a worktable to sort the resulting rows before returning them. The More column displayed in the ToolTip and in the Statistics tab specify why the temporary worktable was created (Order By, Distinct, Reformatting, and so on).
Join	This operator indicates that a join operation was performed between a subquery output and another result set.
Parallelism Merge	This operator indicates that a parallelism merge was done to the result sets returned by all the worker processes involved in the parallel execution. The More column displayed in the ToolTip and in the Statistics tab specify the way in which the results of parallel scans were merged.
Log Scan	This operator indicates that a log scan operation was performed.

About the costs that are available in the plan tree

Precise for Sybase uses the `DBCC TRACEON (310,311)` command to receive the Estimated I/O needed to execute the statement. Since Sybase supplies Estimated I/O only to table-based operations, not all operators in the Plan tree contain I/O Costs. In cases of access both to the index pages and data pages, meaning a Table Lookup is being performed, Sybase provides the Total I/O operations needed while accessing the table. Therefore, the I/O Costs are displayed in the Table Lookup operator and it applies both to the index access and table access. The following operators contain the Total Estimated I/O Costs of their Subtree operators: The entire statement, STEP N and Subquery N.

The table below shows the I/O Cost information operators available in the Plan tree.

Table 9-4 I/O cost information operators

Counter	Description
Estimated Cost	Logical Reads * 2 milliseconds (which is the assumption of how long a logical read takes) + Physical Reads * 18 milliseconds (which is the assumption of how long a physical read takes).
Estimated Subtree Cost	Total estimated cost of all the operator's Subtree operators.
Logical Reads	Number of logical reads performed on the specified table. The number of logical reads estimated by the optimizer is based on the access method and the size of cache. In this case, the optimizer is estimating how many pages are likely to be in cache.
Physical Reads	Physical reads are estimated based on the number of logical reads and the cache size. The optimizer is estimating how many pages will not be available in cache and will need to be read from disk.
Scan Count	A "scan" can represent any of these access methods: <ul style="list-style-type: none"> • A table scan. • Access using a clustered index. Each time the query starts at the root page of the index and follows pointers to the data pages, it is counted as a scan. • Access using a non-clustered index. Each time the query starts at the root page of the index and follows pointers to the leaf level of the index (for a covered query) or to the data pages, it is counted. • If queries run in parallel, each worker process access to the table is counted as a scan.
Estimated Rows	Number of rows from the current table that is estimated to satisfy the join.

About viewing a statement's text

The Highlights tab displays the statement's text.

About viewing which objects are referenced by the execution plan



The Objects tab displays three tables (Tables used in the plan, Indexes of table, and Columns of table) that list all referenced objects in the Execution Plan, including their indexes and columns. Statistical details and general details are displayed for each object and its sub-entities.

About tables used in the plan

Displays a list of all the tables referenced in the tree.

The table below describes the information displayed in the tables used in the plan.



Table 9-5 Tables used in the plan

Column	Description
	Click to launch the Objects tab and focus on the specified table.
	Click to locate and highlight all the operators in the Execution Plan that access the specified table.
Used	Indicates whether the specified table is used in the selected operator in the Execution Plan tree.
Table	Displays the full table name ([Database Name].[Owner].[Table Name]).
Cost	Total cost of the operators that accessed the specified table.
Cost (%)	Total cost in percentage of the operator that accessed the specified table.
Rows	Number of rows in the specified table.
Pages Allocated	The number of pages allocated in the disk for the table records. The value might be incorrect due to out-of-date space usage information. If you suspect the value to be incorrect, use the <code>sp_flushstats</code> command to flush current space usage information from memory or run the <code>update statistics</code> command to recalculate the space usage fields.
Pages Used	The number of pages used in the disk for the table records. The value might be incorrect due to out-of-date space usage information. If you suspect the value to be incorrect, use the <code>sp_flushstats</code> command to flush current space usage information from memory or run the <code>update statistics</code> command to recalculate the space usage fields.

About indexes of table

Displays a list of all the indexes of the selected table displayed in Tables Used in the Plan. The table below describes the information displayed in indexes of table.

Table 9-6 Indexes of table

Column	Description
	Click to launch the Objects tab and focus on the specified index.
	Click to locate and highlight all the operators in the Execution Plan that access the specified index.
Used	Indicates whether the specified index is used in the selected operator in the Execution Plan tree.
Index	Displays the index name.
Pages Allocated	The number of pages allocated in the disk for the index records. The value may be incorrect due to out-of-date space usage information. If you suspect the value to be incorrect, use the <code>sp_flushstats</code> command to flush current space usage information from memory or run the <code>update statistics</code> command to recalculate the space usage fields.
Pages Used	The number of pages used in the disk for the index records. The value may be incorrect due to out-of-date space usage information. If you suspect the value to be incorrect, use the <code>sp_flushstats</code> command to flush current space usage information from memory or run the <code>update statistics</code> command to recalculate the space usage fields.
Cost	Total cost of the operators that accessed the specified index.
Cost (%)	Total cost in percentage of the operators that accessed the specified index.
Index Depth	Number of levels in the index.
Unique	Indicates whether the index is unique.
Clustered	Indicates whether the index is clustered.

About Columns of Table

Displays a list of all columns in the selected table displayed in Tables Used in Plan. When an index is selected in Indexes of Table, the first column constitutes the index column sorted by the position of the column in the index, and marked with an Ascending or Descending icon.

The table below describes the information displayed in the columns of the table.

Table 9-7 Columns of table


Column	Description
A↑ D↓	Indicates whether the column is part of the selected index and whether it is sorted in ascending or descending order.
Column	Displays the column name.
Key Number	Displays the key number of the column in the index.
Type	Displays the physical storage type of the column.

Displaying statistical information on operators in the execution plan tree

The Statistics tab displays statistical information of all operators in the Execution Plan tree. When you sort a table, you can locate operators by their estimated cost or by any other column in the table.

The table below shows the statistical information displayed for Operator ID and Type.

Table 9-8 About viewing statistics on Operator ID and Type


Column	Description
	Click to locate and highlight the operator in the Execution Plan tree that matches the selected operator in the grid.
Operator ID and Type	Displays the operator ID and type.
Logical Reads	Number of logical reads performed on the specified table. The number of logical reads estimated by the optimizer is based on the access method and the size of cache. In this case, the optimizer is estimating how many pages are likely to be in cache.
Physical Reads	Physical reads are estimated based on the number of logical reads and the cache size. The optimizer is estimating how many pages will not be available in cache and will need to be read from disk.
Subtree Cost	Displays the total estimated cost of all the current operation's sons.
Estimated Cost	Displays the estimated cost of the current operation. A high cost value may indicate a problem in the current implementation of the operation. Estimated cost is calculated as followed: Logical Reads * 2 milliseconds (which is the assumption of how long a logical read takes) + Physical Reads * 18 milliseconds (which is the assumption of how long a physical read takes).
Estimated Cost (%)	Displays the estimated cost (in percentage) of the current operator.
Estimated Rows	Number of rows from the current table that is estimated to satisfy the join.
Scan Count	A "scan" can represent any of these access methods: <ul style="list-style-type: none"> • A table scan. • Access using a clustered index. Each time the query starts at the root page of the index and follows pointers to the data pages, it is counted as a scan. • Access using a non-clustered index. Each time the query starts at the root page of the index and follows pointers to the leaf level of the index (for a covered query) or to the data pages, it is counted. • If queries run in parallel, each worker process access to the table is counted as a scan.
More	Displays additional information according to the specified operation. For example, in a table scan operation, the scan type (forward/backward), I/O size and buffer replacement strategy will be displayed.
Parallel	Indicates whether the operator is running in parallel.

About viewing statistical information on operation types

The Operations item under the Menu tab displays two tables that show statistical information broken down into operation types of all operators in the Execution Plan tree.


The table below describes the information displayed for major operation types .

Table 9-9 Major operations

Column	Description
	Click to locate and highlight the operators in the Execution Plan tree that match the selected major operation type.
Major Operation Type	Displays the major types of the operators. <ul style="list-style-type: none"> • The operators in the Execution Plan tree are separated into the following major types: • Table (any Read access to table) • Index (any Read access to index) • Sort (any sort operation) • Join (any join operation) • Spooling • Other (all other operations)
Estimated Cost	Displays the total estimated cost breakdown according to major operation types.
Operations	Indicates the number of operators of a specific major operation type.
Objects	Indicates the number of referenced tables/indexes of the specified operation type. This column is only relevant for table and index operation types.

The table below describes the information displayed for regular operation types.

Table 9-10 Operations

Column	Description
	Click to locate and highlight the operators in the Execution Plan tree that match the selected operation type.
Operation Type	Indicates type of operator, such as, Clustered Index Scan, Index Seek, Nested Loop, and so on.
Major Operation Type	Displays the major types of the operators. The operators in the Execution Plan tree are separated into the following major types: <ul style="list-style-type: none"> • Table (any access to table) • Index (any access to index) • Sort (any sort operation) • Join (any join operation) • Spooling • Other (all other operations)
Estimated Cost	Total estimated cost breakdown according to operation types.
Operations	Indicates the number of operators of the specified operation type.
Objects	Indicates the number of referenced tables/indexes of the specified operation type.

About viewing general properties of a selected statement

The Properties item under the More tab displays general details of the selected statement. The table below describes the properties displayed for a selected statement.

Table 9-11 Selected statement properties

Column	Description
Statement ID	Displays the name of the statement or batch that the SQL tab was launched with. Can also be the name of the statement opened manually.
Database	Indicates the database that was in use when the statement was run.
Parsing User	Indicates the user that was in use when the statement was run.
Cabinet	Indicates the cabinet in which the statement was saved.
Folder	Indicates the folder in which the statement was saved.
Estimate Cost	Indicates the total estimate cost of the statement as calculated for the last access plan.
Most Recent Show Plan	Indicates the last time the statement or batch was explained.
Plan Time	Indicates when the Access Plan was saved. An Access Plan is saved in two cases: 1) there is no Access Plan for the statement and this is the first time it is saved; 2) changes were made to the statement's Access Plan.

Observing a history of resource consumption

The History view tab displays Resource Consumption Over Time vs. Access Plan Changes, to determine how the changes affected the performance of the statements or batches.

The History view is divided into the following areas:

- Main area
- Change history area

About the Main area

The Main area (upper area) displays Resource Consumption Over Time vs. Cost and Changes Over Time.

About viewing a history of resource consumption over time

The History view displays Resource Consumption Over Time vs. Cost and Access Plan Changes Over Time in a number of overtime graphs.

In Sybase	Displays the Resource Consumption of the selected statement over a specified time period.
Cost	Displays the Estimated Cost of the selected statement over a specified time period.
Changes	Displays Access Plan Change indications for the statement vs. Average Duration over a specified time period.
Executions	Displays the number of executions of the selected statement over a specified time period.

About viewing execution plan text

The Text view displays the Execution Plan's text.


About the Change History area

The Change History area (lower area) displays all the Access Plan changes made over a specified time period.

About viewing Access Plan Changes

The Access Plan Changes view displays a list of all the Access Plans of the statement as detected by the Explain Statement process.

Table 9-12 Access Plan Changes table

Column	Description
	Click to switch to Plan view and display the selected Access Plan.
Date	Indicates the date the specified Access Plan was detected by the Explain Statements process.
Estimate Cost	Indicates the estimated cost of the specified Access Plan.
Table Scans	Displays the number of table scans performed in the specified Access Plan.
Index Scans	Displays the number of index scans performed in the specified Access Plan.
Parallel	Displays the number of parallel operations performed in the specified Access Plan.
Nested Loops	Displays the number of nested loops performed in the specified Access Plan.
Merge Joins	Displays the number of merge joins performed in the specified Access Plan.
Sorts	Displays the number of sort operations performed in the specified Access Plan.

Breaking down a batch's access plan into statements

In many cases a statement is part of a long batch (such as a Stored Procedure or improvised batch). Sometimes you want to examine a particular statement, and sometimes you want to view the statement in the context of the entire batch.

The Statements view breaks down the Access Plan of the entire batch into statements and correlates the statements in the Access Plan with the statements that were captured by the Precise for Sybase Collector agent. The statement IDs of these statements that were captured are displayed in Precise. It is possible to view the Access Plan of a different statement by clicking its respective Plan icon.

NOTE Only statements that were captured by the Precise for Sybase Collector agent can be chosen.

When you switch to this view in the context of a batch, all the batch's statements are listed. In this case, all the information displayed is related to the selected batch.

The table below describes the information displayed in the Statement's View table.

Table 9-13 Statement's View table


Column	Description
	Click to switch to Plan view, in-context with the selected statement. NOTE: This icon is only enabled at batch level and in the statement that was captured by the Collector agent.
ID	Indicates the ID of the statement in the batch.
Collector Statement	Displays the ID of the statement and batch as assigned by the Precise for Sybase Collector agent.
Physical Operator	Describes how a query or update was executed. The physical operator describes the physical implementation used to process a statement, for example: Batch level, Select, Update, Cond, and so on.
Operations	Displays the number of operations in the specified statement.
Table Scans	Displays the number of Table Scans performed in the specified statement.

Table 9-13 Statement's View table

Column	Description
Index Scans	Displays the number of Index Scans performed in the specified statement.
Joins	Displays the number of join operations performed in the specified statement.
Sorts	Displays the number of sort operations performed in the specified statement.
Estimated Cost	Displays the estimated cost for the specified statement.

About the Actions tab

The Actions tab, on the Precise bar, lets you perform additional actions. You can perform the following actions:

- Create a new statement
- Open an existing statement
- Edit an existing statement
- Re-Explain an existing statement
- Display the Sybase plan using Sybase SQL Advantage and Interactive SQL

Creating a new statement

You can create a new Statement and save it in the PMDB in a logical cabinet and folder hierarchy. You can also rewrite statements and view their Access Plans.

To create a new statement

1. On the Precise bar, click **Actions>New**.
2. In the New Statement dialog box, choose the relevant properties for the new statement from the drop-down lists and enter a statement name, as follows:

Instance	Indicates the name of the instance that the statement belongs to.
Cabinet	Indicates the name of the cabinet that the statement is saved in (creates a cabinet if it does not already exist).
Folder	Indicates the name of the folder that the statement is saved in (creates a folder if it does not already exist).
Name	Indicates the name of the statement (generated randomly).
Database	Indicates the name of the database where the statement is running.
User	Indicates the statement's parsing user.

3. Enter the SQL text for the statement in the text frame.
4. Click **OK**.

The new statement is saved in the PMDB in a logical cabinet and folder hierarchy.

Opening an existing statement

You can view a statement that was saved in the PMDB. To open an existing statement

1. On the Precise bar, click **Actions>Open**.
2. In the Open Statement dialog box, choose the relevant properties for the statement you want to view from the Instance, Cabinet, Folder and Name drop-down lists, as follows:

Instance	Indicates the name of the instance that the statement belongs to.
Cabinet	Indicates the name of the cabinet that the statement is saved in (creates a cabinet if it does not already exist).
Folder	Indicates the name of the folder that the statement is saved in (creates a folder if it does not already exist).
Name	Indicates the name of the statement (generated randomly).
Database	Indicates the name of the database where the statement is running.
User	Indicates the statement's parsing user.

The statement's text is displayed in the text frame.

3. Click **OK**.

Editing an existing statement

You can edit a statement that was saved in the PMDB. The statement that was collected by the Collector agent is saved as a New Statement in a logical cabinet and folder hierarchy.

To edit an existing statement

1. On the Precise bar, click **Actions>Edit**.
2. In the Edit Statement dialog box, define the statement you want to edit by choosing a relevant property from the drop-down lists, as follows:

Instance	Indicates the name of the instance that the statement belongs to.
Cabinet	Indicates the name of the cabinet that the statement is saved in (creates a cabinet if it does not already exist).
Folder	Indicates the name of the folder that the statement is saved in (creates a folder if it does not already exist).
Database	Indicates the name of the database where the statement is running.
User	Indicates the statement's parsing user.

3. Enter a New Statement name in the Name field.
4. Edit the SQL text for the statement in the text frame.
5. Click **OK**.

The edited statement is saved in a logical cabinet and folder hierarchy.

Re-Explaining an existing statement

You can Re-Explain an existing statement. If the new access plan is different from the most recent access plan, it is saved in the PMDB and displayed.

To re-explain an existing statement

- Click the Re-Explain tab

Display the Sybase plan using Sybase SQL Advantage and Interactive SQL

You can display the Sybase plan using Sybase SQL Advantage and Interactive SQL. To display the Sybase plan using Sybase SQL Advantage and Interactive SQL

- Click the Sybase Plan tab

The Sybase SQL Advantage window is opened with the statement in context and commands to display its access path are executed.

How the SQL tab can help you identify performance problems

After identifying a problematic statement that is slowing down the response time of a specific application, the first step in tuning the statement is to understand the access path that Sybase chose for the statement. The Explain procedure is designed to clarify the access path chosen for a statement and to translate it into a visual medium. Therefore you can easily see whether the optimizer chose the proper Execution Plan. For example, you can see whether the optimizer performed an Index Seek as expected. In addition, you can see if the Access Plan of the statement was changed and how the changes affected the performance of the statement.

You can identify a performance problem by doing one or more of the following:

- See “Identifying problematic operators in the latest access plan” on page 111.
- See “Locating referenced tables that pose potential problems” on page 111.
- See “Examining how access path changes affect statement performance” on page 112.
- See “Locating the most resource-consuming statements in a batch” on page 112.

Identifying problematic operators in the latest access plan

The first step in tuning a statement is to identify problematic operators. Problematic operators are operators whose estimated cost is high.

To identify problematic operators

1. Launch to the SQL tab with a statement in-context or open the statement you want to analyze in the SQL tab.
2. Open the Statistics tab in Plan view
3. Sort the grid by the Estimated Cost value (desc.)

NOTE The lower Estimated Cost the better.

Following are examples of how to improve problematic operators:

- If the operator is a Table Scan, try creating an index.
- If the operator is of a Join type (such as Nested Loop or Merge), examine a different join method.
- If the operator is a Sort operation, check if the sort is required by the application or can be removed. For example, remove `distinct` or change the `union` clause to `union all`.

See “About the Execution plan tree” on page 100.

Locating referenced tables that pose potential problems

When it is not possible to access the statement, such as in third company products, the only way to improve the statement is by improving access to the data. This can be achieved by performing changes in the referenced objects by examining the operators that access the tables and identifying potential problematic operators.

To locate reference tables that pose potential problems

1. Launch to the SQL tab with a statement in-context or open the statement you want to analyze in the SQL tab.
2. On the Plan tab, click **Objects**.
3. Click the Locator icon (to locate and highlight all the operators in the execution plan that access the specified table).
4. Examine the cost of these operators (to locate the problematic operators).
5. Identify the objects that are referenced by the problematic operators and try to tune those objects.

See “About viewing which objects are referenced by the execution plan” on page 103.

Examining how access path changes affect statement performance

Precise for Sybase allows you to track the Access Plan changes and compare them with the Resource Consumption of your statement to understand how the changes affected the statement's performance.

To examine how access path changes affect statement performance

1. Launch to the SQL tab with a statement in-context or open the statement you want to analyze in the SQL tab.
2. On the History tab, on the View controls, click **History**.
3. Analyze the statement's resource consumption over time vs. execution cost, access path changes, and number of executions over time. This can help you to observe whether the change in the statement performance occurred due to the change in the access path or not.
4. In the Changes History table, you can view additional details on the change made. You can select the launch icon for a particular access path and launch to the Plan tab in-context to continue your analysis of the change.

Locating the most resource-consuming statements in a batch

The main difference between examining a statement's Access Plan and a batch's Access Plan is in the number of operators in the Access Plan. Understanding an Access Plan with hundreds of operators is a difficult task. The Statements view can assist you with this task since it shows the Access Plan statements and their Estimated Cost. The ID of the statements sampled by the Precise for Sybase Collector agent are also displayed.

To locate the most resource-consuming statements in a batch

1. Launch to the SQL tab with a statement in-context or open the statement you want to analyze in the SQL tab.
2. On the Statement tab, observe the access plan statements and their estimated cost.
3. Determine which of the statements has the highest estimated cost and focus your analysis on that statement.

abandonment rate	In Web, a counter that keeps track of the percentage of users that abandon the loading of the Web page before it completes downloading.
action	An operation that Alerts FocalPoint automatically performs when detecting a warning or critical status for a specific metric. According to the defined action, Alerts FocalPoint opens a message box, sends an email or SNMP trap, or executes a program.
action item (manual)	Some installation steps cannot be executed automatically by Precise Installer. The administrator is required to execute them manually. Action Items are presented (when necessary) at end of installation or product update.
AdminPoint	The central administration console of Precise that facilitates the maintenance, configuration, and management of all installed Precise components, such as monitoring the status of all Precise agents and PMDB processes, getting license information, starting and stopping the agents, getting log data on agents and events, changing PMDB settings, and installing patches. See also PMDB . See also agent .
advice	In Oracle, an algorithm that is designed to recommend on gathering statistics, creating new indexes, change existing indexes, and add or delete hints to make Oracle's Optimizer choose a better access path and make the statement perform better. Can be activated from any DML (Data Manipulation Language) statement.
agent	A program that runs on a server machine to collect, process, or communicate performance information. The Precise installation consists of multiple agents.
alert	The state of an Alerts metric that has reached a near-critical or critical status. An alert is issued by Alerts, triggering an action and informing of a problem that has occurred or is likely to occur within the area sampled by the specific metric. See also action .
Alerts	These product provide alerts to problematic conditions before they turn into performance problems, based on predefined metrics and thresholds. Alerts can automatically perform an action, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program.
Alerts FocalPoint	An agent that receives data from the InformPoint agents, stores it, and performs any action that has been user-defined for that specific alert, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program. See also InformPoint .
alert type	In Alerts, the status of all metrics belonging to a metric group or a monitored instance, indicating the current performance level through colors. See also alert . See also metric .
application server metric	In J2EE, a metrics that is provided by the application server or by customer code. This can include metrics published by the Java Management Extension (JMX) APIs or vendor-specific APIs, such as IBM's Performance Management Interface (PMI).
AppTier	The abbreviation for an application tier in a Precise environment. An AppTier contains one or more instances of the same technology and purpose. Application tiers do not necessarily correspond to distinct physical servers: in many cases, the tiers are logical, with one server running multiple AppTiers or one AppTier spans a cluster of servers. A Precise environment may contain multiple AppTiers on the same technology. For example, you may group J2EE instances (JVMs) into a J2EE Presentation AppTier and a J2EE Business Logic AppTier. Segmenting application service time into the contribution of individual application tiers is helpful in identifying the source of performance problems. Analyzing the performance and behavior of each tier separately is crucial for isolating the root causes of performance problems.
cabinet	In Oracle, SQL Server, DB2, and Sybase, the highest logical level in the SQL workspace hierarchy. A cabinet contains folders and, within folders, statements.

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

s a remote EJB. J2EE makes a distinction between service request invocations and other invocations. The first HTTP or EJB invocation within a call path is designated as an entry point.

environment	<p>The highest Precise logical group. It may contain multiple AppTiers of various technologies that serve an application together. For example, a Payroll Production environment may contain all Web servers, application servers, transaction managers, databases, and servers that serve this application. Alternatively, it may contain any set of instances that form together an administrative group. Since Precise version 8.5, an instance can belong to more than one environment. Let's say a single Oracle database serves two different applications plus the DBA wants to associate this database with a group of other databases under his responsibility regardless of served applications. In this case the Oracle database will be associated to three different Precise environments.</p> <p>See also AppTier.</p> <p>See also instance.</p>
ERP Extension	<p>In Oracle, DB2, and SQL Server, an extension to the Collector that provides detailed information on the activities and resource consumption of packaged application components. It correlates the database information, and the packaged application information and lets you see users, transactions, reports, and other elements of ERP applications, such as Oracle Applications, SAP, PeopleSoft, and Siebel.</p>
event	<p>In Alerts, the occurrence of a sampling or progress. A sampling occurrence occurs every time a metric samples. A progress occurrence occurs when a metric's progress status is changed or when the investigated status reaches the end of the given investigation time. In AdminPoint, all occurrences reported by Precise agents, including informational events, warnings, and errors related to one of the agents. All events are shown in the Events view.</p> <p>See also metric.</p> <p>See also progress.</p>
executions	<p>In Oracle, SQL Server, DB2, and Sybase, the number of times a SQL statement was executed during the selected time frame. In SmartLink and SAP, the number of times a transaction was executed during the selected time frame.</p>
explained statement	<p>In Oracle, SQL Server, DB2, and Sybase, a statement whose access path (chosen by the RDBMS Optimizer) is clarified and translated into a visual display. Explained results include information on the objects referenced by the statement and the operations performed on these objects.</p>
extended collection	<p>In Oracle, a function that proactively specifies a future period during which Oracle activity data is collected and organized for subsequent analysis. Extended collections are an easy means to view collected information, assess application resource consumption, and identify bottlenecks that are inhibiting application performance and end-user productivity.</p>
Federated Precise	<p>Federated Precise is a version of Precise (starting with version 8.5) that can manage multiple Precise installations within unified StartPoint and AdminPoint screens, displaying and managing all environments, instances, and installations.</p>
findings	<p>A ranked list of top performance problems in the selected context. They are presented in Oracle, SQL Server, .NET, and J2EE. Each finding provides: a short explanation of the problem, detailed background information, and most important: one or more links to further investigate the problem and possibly get more granular findings on the selected context.</p>
first byte time	<p>In Web, a counter that keeps track of the time that it takes from the moment a new Web page is called until the first byte arrives back from the Web server.</p>
FocalPoint	<p>An agent that communicates with the Listeners installed on the monitored servers, receives data from the Collectors, periodically processes and stores this data in the PMDB, and serves UI requests.</p>
folder	<p>In SQL Server, Oracle, DB2, and Sybase, the intermediate logical level in the SQL workspace hierarchy. Folders are grouped into cabinets and contain SQL statements.</p> <p>See also cabinet.</p>
Framework Installer	<p>The application that facilitates the installation of Precise framework components. It can be invoked from the installation DVD to install a new Precise deployment. It can also be invoked from an existing Precise deployment to install an additional framework node and attach it to the originating Precise deployment.</p> <p>See also Framework node.</p>
Framework node	<p>A set of FocalPoint agents that are installed together and manages a set of monitored instances. The performance data of these instances will be loaded into a dedicated PMDB. A single Precise deployment may contain multiple framework nodes (using a separate PMDB for each node). An environment cannot span over multiple framework nodes.</p>

garbage collection	An automatic process in the Java runtime environment that periodically reclaims memory used by objects that are no longer referenced. The process can impact an application's performance while memory is being reclaimed. Java programmers may initiate garbage collection explicitly.
group	In Web, the identifier that is used to group other identifiers, such as sites or URLs.
hint	In Oracle, an instruction directed at the Oracle Optimizer that includes considerations for an execution plan. The Oracle Optimizer will build an execution plan based on the hint, ignoring its own set of considerations.
hour group	A unit that reflects the type and level of activity within the system at different times. By defining the times of the day that are peak and off-peak, or day and night, the performance analysis can be focused on those particular times of the day. If, for example, most performance problems occur within nighttime and weekend batches, it can be useful to focus only on them.
InformPoint	In Alerts, an agent that retrieves performance data from all installed Precise products, analyzes it, and sends an alert to the Alerts FocalPoint when the predefined thresholds are exceeded. See also agent .
Insight	The Precise product family that facilitates the process of monitoring and correlating system performance. It consists of Insight. See also Report Manager . See also FocalPoint .
Insight FocalPoint	An agent that receives performance information from Insight Savvies, which monitor the environment. Insight FocalPoint then correlates, processes, and stores this information in a centralized location. The Insight performance history is stored in the PMDB. See also Savvy . See also PMDB .
instance	A monitored object of a specific technology. The following list specifies what constitutes an instance for the various supported technologies: J2EE - a Java Virtual Machine (JVM - a logical name set by the user), Microsoft .NET - a Common Language Runtime (CLR - a logical name set by the user), Oracle Applications - an Oracle Applications Form server, Oracle - an Oracle instance, SAP - a SAP system, SQL Server - an SQL instance, Sybase - a Sybase instance, Tuxedo - a Tuxedo domain, Web - a Web server, WebSphere MQ - an IBM WebSphere Queue Manager, DB2 - a DB2 non-partitioned database or a DB2 database partition. During installation, the instance is associated with one AppTier and environment. An Instance can be moved to a different environment or associated with multiple environments without re-installation nor losing historical data. See also environment .
instance statistics	See performance counter .
instance/database changes	In SQL Server and Oracle, the part of the Collect Schema Changes process that captures instance definition changes and database option changes and saves them in the PMDB.
instrumentation	The process of inserting fault-tolerant recording hooks in Java byte code, .NET MSIL, HTML pages, or other monitored components, resulting in the capture of performance metrics. In J2EE, a mechanism that enables collecting performance information when an application is executed. The process involves inserting special fault-tolerant recording hooks into application class objects. In WEB, the insertion of recording hooks into HTML pages can either be in memory (Dynamic Instrumentation) or file-based (Static Instrumentation).
instrumentation context	See invocation context .
internal invocation	The process of invoking a request from an HTTP request (Servlet or JSP) or an EJB. J2EE displays an internal invocation like any other invocation.
invocation	An execution of a J2EE entity (a Servlet, EJB, SQL Statement, method, and so on). When mentioned in plural (Invocations), means the amount of times that the entity was executed.
invocation context	The context within which a method is invoked. For example, if Method A is invoked both from Method B and from Method C, there will be two different invocation contexts for the performance metrics collected for Method A, one for when it is invoked from Method B and one for when it is invoked from Method C. See also call path and instrumentation context .
JRE (Java Runtime Environment)	As the runtime part of the Java software, the combination of the components that enable the execution of a Java program: a Java virtual machine, the core class libraries, and the files that form the Java platform.

JSP (Java Server Page)	An HTML page with special tags for Java scripting. An application server processes the tags and generates a Servlet.
JVM (Java Virtual Machine)	An instance of a JRE that executes Java programs. A server-side Java application server is itself a Java program that runs inside a JVM. Servlets, JSPs and EJBs are Java programs (applications) that run within the application server's JVM. J2EE monitors the JVM running the application server and the server-side Java applications within the application server.
JVM Heap Memory key metric	The amount of real computer memory that is allocated to the JVM for executing Java programs. In Alerts, a parameter that monitors a very important performance aspect. The status of the Oracle instance (up/down), for example, is crucial for system performance. If the instance is down, this is the first problem that needs to be solved. Marking a metric as key metric ensures that a critical alert raised for this metric receives top priority by the person that is responsible to handle alerts at any time. It can also help determine which alert to handle first in case of multiple alerts. In Alerts, key metrics are always displayed at the top of a metric table when it is sorted by alerts so that they get immediate attention.
Listener	The agent that facilitates the communication between the various Precise agents across different servers must be installed on every server where Collectors or FocalPoint agents are installed. The Listener allows communication with all other agents installed on the monitored server, while only the Precise Listener port is known by other servers.
login name	In SQL Server and Sybase, the session identifier that represents the credential used to connect to the database. When an ERP extension is installed, the user name of the packaged application's client overrides the login name. For example, when SAP extension is installed, the SAP user name overrides the login name.
machine	A session identifier. A machine as sampled, for example, by the SQL Server Collector is the identification of the machine where the client process executes. Machine is also sampled by Oracle, DB2 and Sybase. In Insight terminology Machine is called Client machine.
Main framework node	A main framework node is the single point for login and also serves as the Precise FocalPoint for the entire deployment.
major collection count	Number of estimated major garbage collection events that occurred during the last J2EE collector slice. A major garbage collection can stop the application while JVM heap memory is being reclaimed. See also garbage collection .
major garbage collection time	Percentage of time spent by the JVM executing major garbage collection events during the displayed interval. See also garbage collection .
MBean	A Java object that represents a manageable resource. In J2EE, MBeans, or Managed Beans, are used for application server metrics.
memory logger interval	The interval at which the J2EE Collector gathers JVM heap memory data. All snapshots of the memory logger's data collected according to the memory logger interval are summarized with counters in the current aggregation interval. The memory logger interval's time span is typically a small fraction of the aggregation interval.
metric	In Alerts, a query that helps measure performance in the environment. Three types of metrics are available: System metrics relate to the internal resources, operations, and objects of the monitored infrastructure; application metrics reflect the way the applications perform; user-defined metrics can be customized to specifically relate to a site. When a metric's value exceeds one of the defined thresholds, its status changes to near-critical (yellow bullet) or critical (red bullet). See also key metric .
metric set	In Alerts, a unit that groups metrics that measure related performance aspects. The following metric sets exist: Status: includes metrics that alert to functional problems related to the instance. Performance: includes metrics that alert to performance problems related to the instance. Load: includes metrics that alert to instance-related load problems that may later cause errors or crashes. Service: includes metrics that alert to instance-related SLA breaches. The metrics in this set are sampled by Insight. Performance Trending: includes metrics that alert to potential future performance problems. The metrics in this set are sampled by Report Manager.

	<p>Load Trending: includes metrics that alert to potential future load problems. The metrics in this set are sampled by Report Manager.</p> <p>Customized: includes user-defined metrics.</p> <p>Precise Status: includes metrics that alert to the near-critical or critical status of the installed Precise environment.</p>
module	A session identifier. In Oracle it contains the value of MODULE column in V\$SESSION table. In DB2 it contains the command name of non-SQL statements or the package name for SQL statements.
near-critical status	In Alerts, a status indicating that the value of the sampled metric has exceeded the defined near-critical threshold. A near-critical status is indicated by a yellow bullet.
network time	In Web, a counter that keeps track of the time spent on network activity from the server side perspective. This includes the time to read the request from the network and the network time to send the response back to the client.
node	<p>See Framework node.</p> <p>Note: when creating a new installation from DVD, it would be called "Framework." In the UI screens (columns in tables, choosing a system for a new environment etc), we would call them "Nodes." For example: choose a node for the new environment.</p> <p>When adding a new system within AdminPoint, it would be called a "Framework node".</p>
packaged application	<p>An application that is created and/or maintained by a third party and is not custom-built to one's specific needs. The following packaged applications have special treatment by Precise: SAP, Oracle Applications, PeopleSoft, and Siebel. They are harder to modify because the application code is either not available or hard to understand.</p> <p>See also ERP Extension.</p>
page size	In Web, a counter that keeps track of the amount of data loaded from the server to display the page.
page views	In Web, a counter that keeps track of the number of Web pages viewed at a specific Web site during a selected time period.
parent metric	<p>In Alerts, a joining of several child metrics. Each time a parent metric samples, it gathers data from a set of child metrics and presents it as a single metric query. The individual child metric values are displayed on the Thresholds tab of the Properties dialog box in Alerts.</p> <p>See also submetric.</p>
performance counter	In SQL Server, a Windows performance counter as reported by the operating system. In Oracle, an instance level statistic as gathered from some Oracle V\$ views. In Sybase, a sysmonitors counter collected by Sybase. In Web, an operating system or a Web server performance counter.
PMDB	The Precise data warehouse of performance and availability data. It can be hosted on an Oracle or SQL Server database.
portal server	An application server for Web-based applications that commonly provide personalization, single sign-on, and content aggregation from different sources and that host the presentation layer of Information Systems. Aggregation is the action of integrating content from different sources within a Web page. A portal may have sophisticated personalization features to provide customized content to users. Portal pages may have a different set of portlets creating content for different users.
portlet	A Java-based Web component, managed by a portlet container, that processes requests and generates dynamic content. Portlets are used by portals as pluggable user interface components that provide a presentation layer to Information Systems.
Precise deployment	An independent Precise system. It contains and manages various agent types and provides centralized monitoring and administration. A Precise deployment may contain multiple framework nodes. One of them is defined as the main framework node and it manages all other nodes of the Precise deployment.
Precise FocalPoint	See Precise FocalPoint agent .
Precise FocalPoint agent	An agent that manages all agents in a single Precise deployment. Additional product FocalPoint agents manage specific technologies and resources.
Precise Proxy agent	When multiple Framework Nodes are managed by a single Precise deployment, the Precise FocalPoint agent manages all of the agents of the main Framework node, while every other framework node has a Precise Proxy agent to manage all its agents.
Precise_root	This is the term used in a path for the Precise installation directory. The terms <code><i3_root></code> or <code>i3 root</code> can appear in text too.

program	<p>A session identifier in Oracle, DB2, SQL Server, Sybase, Insight, and Alerts. A program as sampled, for example, by the SQL Server Collector, is the name of the executable that connects to the database. Applications that do not set the application name have N/A as program. When a SAP extension is installed, the SAP transaction overrides the program, and it may change during the application's lifetime. When a Siebel extension is installed, the Siebel views override the program, and it may change during the application's lifetime. In Insight, program is the name of an executable as recognized by the operating system. If an executable is invoked from a script (a batch or a shell), the script is displayed as the command entity. In Alerts, program is part of the customized metrics definition and is the name of the executable or stored procedure executed in the database that runs when the metric is sampled. In action definitions, program is the name of the executable that will run if the metric exceeds its threshold.</p>
progress	<p>In Alerts, the management state of a metric for which an alert has been raised. The following statuses exist:</p> <p>Open: An alert is raised.</p> <p>Investigated: The alert is taken care of.</p> <p>Closed: The problem has been solved.</p>
Proxy FocalPoint	<p>See Precise Proxy agent</p>
recommended index	<p>In SQL Server, a function that uses the Microsoft Index Tuning Wizard to recommend on adding indexes or statistics for a selected statement, batch, or table. Based on the results of this function, the Optimizer will choose a better access plan and make the respective statement or batch perform better. For statements and batches, recommendations are based on the content of the statement or batch. For tables, recommendations are based on all the statements that are stored in the PMDB, are executed during the selected time period, and have an average duration time greater than the value defined in the registry (where 0 is the default).</p> <p>See also advice.</p>
related SQL	<p>In Oracle, a generated statement that uses alternative syntax to access the database in different ways and returns the same output as the original statement.</p>
relative frequency	<p>Number of contributor invocations per entry point. For example, if a service request calls three methods each time it is invoked, the Relative Frequency for the method is three invocations per service request execution. Similarly, if a service request calls one method every other time it is invoked (half of the time), the Relative Frequency for the method is 0.5 invocations per service request execution.</p>
rendering time	<p>In Web, a counter that keeps track of the time that it takes for a Web page to be loaded from the moment the first byte arrives until the Web page is fully loaded or the user interrupts or abandons the loading process.</p>
report	<p>A collection of queries, programming code, and layout settings that Report Manager executes to generate graphical results like tables and charts.</p>
Report Manager	<p>Uses historical information to identify problematic conditions, track long-term performance, volume trends and patterns, view availability problems over long periods of time and on different levels, compare the performance of similar systems, correlate between performance metrics of different products, assist in capacity planning, and generate demand-driven, user-defined reports.</p>
Report Manager FocalPoint	<p>An agent that examines the PMDB tables to produce scheduled performance reports.</p> <p>See also PMDB.</p>
report parameter	<p>In Report Manager, a keyword used in a report. Its value is set during the report execution. A parameter's value can be updated either permanently or for the current execution only.</p>
report property	<p>In Report Manager, the attributes that define a specific report, consisting of report parameters and scheduling information.</p> <p>See also report parameter.</p>
request	<p>In Web, a counter that keeps track of the number of HTTP requests sent for a viewed entity. request</p>
error	<p>In Web, a counter that keeps track of the percentage of requests completed with an HTTP error.</p>
sampling	<p>In Alerts, the process during which a metric queries a Precise product for a specific instance, retrieves values, and calculates the metric's alert level.</p>
sampling base	<p>In Alerts, the start time of a scheduled sampling process (by default Sunday, 00:00 AM).</p>
sampling period	<p>In specific Alerts metrics, the time frame during which statistical data is returned from other Precise products. Such a metric is for example Oracle's General Behavior metric, which returns database behavior for a certain period of time.</p>

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

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