

Examining the relationships between Precise for Sybase 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 drill-down 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. For more information about Precise for Sybase tabs, see [About Precise for Sybase tabs](#).

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 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.** 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.
The values may be incorrect due to out-of-date space usage information. To update the values, use the `DBCC FLUSHSTATS` command to flush the most up-to-date information regarding the objects size from memory and run `UPDATE STATISTICS` to recalculate the statistics. You can also use the `DBCC CHECKSTORAGE` command to verify and recalculate the space usage information.
- **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.



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 `DBCC FLUSHSTATS` command to flush the most up-to-date information regarding the objects size from memory and run `UPDATE STATISTICS` to recalculate the statistics. You can also use the `DBCC CHECKSTORAGE` 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 `reorg`.
- **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 `sp_chgattribute` 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 [Installation and Administration](#).
- **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 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 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.** Number of executions overtime.

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 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 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.
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:

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

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. For more information about table and view entities, see [About the Table entity and Tables Contain Same Column entity](#) and [About the View entity and Views Contain Same Column entity](#).

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, use the Statistics view of the Column entity to help you see the distribution of values in that column. For more information about viewing these statistics, see [About viewing density and selectivity statistics](#) and [About viewing column statistics](#).