Getting Started
Managing Objects
Managing Data Transfer
Generating SQL Scripts
Performing Queries

Getting Started

Click the button to go to the Help topic.

- Logging In to a SQL Server
- Changing Server Connections
- Changing Databases
- Disconnecting from a SQL Server
- Exiting from SQL Object Manager
- Setting SQL Object Manager Preferences
- Setting SQL Object Manager Configuration Information
- Editing, Printing, and Saving SQL Object Manager Windows
- Changing Your Identity

Managing Objects

Click the button to go to the Help topic.

- Managing Tables
- Managing Indexes
- Managing Triggers
- Managing Views
- Managing Keys
- Managing Rules
- Managing Defaults
- Managing User-defined Datatypes
- Managing Stored Procedures
- Managing Object Permissions
- Managing Objects

Managing Tables

Click the button to go to the Help topic.

- Creating a Table
- Altering a Table
- Copying a Table
- Moving a Table to a Segment
- Dropping a Table
- Viewing Table Properties
- Using DBCC CheckTable
- Selecting Data from a Table

Managing Indexes

Click the button to go to the Help topic.

- Creating Indexes
- Modifying an Index
- Rebuilding an Index
- Renaming an Index
- Updating Index Statistics
- Dropping an Index

Managing Triggers

Click the button to go to the Help topic.

- Creating a Trigger
- Editing a Trigger
- Dropping a Trigger
- Viewing Trigger Properties

Managing Views

Click the button to go to the Help topic.

- Creating a View
- Editing a View
- Dropping a View
- Viewing View Properties
- Selecting Data from a View

Managing Keys

Click the button to go to the Help topic.

- Creating a Key
- Modifying a Key
- Dropping a Key

Managing Rules

Click the button to go to the Help topic.

- Creating a Rule
- Binding a Rule
- Unbinding a Rule
- Dropping a Rule
- Displaying Rule Bindings Information
- Viewing Rule Properties

Managing Defaults

Click the button to go to the Help topic.

- Creating a Default
- Binding a Default
- Unbinding a Default
- Dropping a Default
- Displaying Default Bindings
- Viewing Default Properties

Managing User-defined Datatypes

Click the button to go to the Help topic.

- Creating a User-defined Datatype
- Renaming a User-defined Datatype
- Dropping a User-defined Datatype
- Viewing User-defined Datatype Properties

Managing Stored Procedures

Click the button to go to the Help topic.

- Creating a Stored Procedure
- Editing a Stored Procedure
- Dropping a Stored Procedure
- Viewing Stored Procedure Properties

Managing Object Permissions

Click the button to go to the Help topic.

- Granting Object Permissions
- Revoking Object Permissions
- Viewing Permissions for Objects
- Viewing a User's Permissions
- Viewing a User's Permissions on All Objects

Managing Objects

Click the button to go to the Help topic.

- Viewing Object Dependencies
- Viewing Object Properties
- Renaming an Object
- Dropping an Object

Managing Data Transfer

Click the button to go to the Help topic.

Transferring Data to a SQL Server

Transferring Data from a SQL Server

Generating SQL Scripts

Click the button to go to the Help topic.

Generating SQL Scripts

Performing Queries

Click the button to go to the Help topic.

- Querying a Database
- Saving a Query to a File
- Using a Saved Query
- Setting Query Options
- Editing Query Results
- Analyzing Queries with Graphical Showplan
- Analyzing Queries with Graphical Statistics I/O

Logging In to a SQL Server

Logging in identifies you to a SQL Server and allows you to use the databases on that server. Before you log in, you need to know your login ID and password. The SQL Server Login dialog box automatically appears when you start Microsoft SQL Object Manager. You can also log in to additional SQL Servers at any time while working with Microsoft SQL Object Manager.

To log in to a SQL Server

1. From the Microsoft SQL Object Manager window, choose the Connect button.
   The Connect Server dialog box appears.
2. Type or select a server name, and then type your login ID and password.
3. Choose Connect.

Related Topics

Disconnecting from a SQL Server

Exiting from SQL Object Manager

Changing Server Connections

You can log in to more than one SQL Server at a time. The Current Connections box lists the servers you are connected to. To perform tasks on a server that is not listed, you must first connect to that server in the SQL Server Login dialog box.

To change server connections

- In the Current Connections box, select a server.

Related Topic

Logging In to a SQL Server

Changing Databases

The Current Database box lists the databases you can access. When you are connected to a SQL Server, your default database is the database in focus. If you have permission to perform tasks on other databases, you can change the focus to those databases.

To change databases

- In the Current Databases box, select a database.

Disconnecting from a SQL Server

When you disconnect from a SQL Server, all windows associated with that server are closed. Any query report windows that contain information for that server remain open until you close them.

To disconnect from a SQL Server

1. From the Microsoft SQL Object Manager window, choose the Connect button.
   The Connect Server dialog box appears.
2. Select the server to disconnect from.
3. Choose Disconnect.

Related Topics

Logging In to a SQL Server

Exiting from SQL Object Manager

Exiting from SQL Object Manager

When you exit from SQL Object Manager, all your SQL Server connections are disconnected.

To exit from SQL Object Manager

- From the File menu, choose Exit.

Related Topics

Disconnecting from a SQL Server

Logging In to a SQL Server

Setting SQL Object Manager Preferences

You can set options that control whether to show the Transact-SQL statements used to query the server, to create a query archive file, and to be prompted before exiting SQL Object Manager.

To set SQL Object Manager preferences

1. From the Tools menu, choose Preferences.
   The Microsoft SQL Object Manager Preferences dialog box appears.
2. Set the appropriate preferences.
   To set all preferences to their default values, choose Defaults.
3. Choose OK.

Related Topic

Setting SQL Object Manager Configuration Information

Setting SQL Object Manager Configuration Information

You can specify how much time should elapse before login attempts are timed out, how much time should elapse while the server retrieves query results, and how much time the server should wait for major functions to complete before timing out.

To set SQL Object Manager configuration information

1. From the Tools menu, choose Configure.
   The Object Manager Configuration dialog box appears.
2. Set the appropriate server preferences.
   To set all options to their default values, choose Defaults.
3. Choose OK.

Related Topic

Setting SQL Object Manager Preferences

Editing, Printing, and Saving SQL Object Manager Windows

You can print or save the information that appears in a SQL Object Manager window. You can edit, print, or save the information that appears in a results window.

To edit information in a results window

1. Select the text to cut or to copy.
2. From the Edit menu, choose the appropriate editing task.

To print window information

From the File menu, choose Print.

To save window information

1. From the File menu, choose Save As.
   The Save As dialog box appears.
2. In the File Name box, type a filename.
3. Choose OK.

Changing Your Identity

If you are the system administrator or database owner, you can impersonate another user who owns objects in the database so that you can grant and revoke permissions on those objects.

To change your identity

1. From the Tools menu, choose Set User.
   The Set User dialog box appears.
2. Select the user to impersonate.
3. Choose OK.
   The menu item changes to Set User dbo.

To reset yourself to be the database owner

From the Tools menu, choose Set User dbo.

Creating a Table

When you Create a table you name its columns, supply a datatype for each column, and specify whether a particular column can have null values. You can also specify a segment on which to place the table. Each table can have as many as 250 columns and a database can have as many as 2 billion tables. Note that you must have Create Table permission to create a table.

To create a table

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. In the Tables box, select New Table.
3. Choose a segment for the table to reside on.
4. Select the Column Name cell. Then, in the Properties box, type the column name.
5. In the same row, select the Datatype cell and then, in the Properties box, select the datatype.
6. If you selected a char, varchar, binary, or varbinary datatype, select the Length cell in the same row. Then, in the Properties box, type the maximum number of characters that can be stored in the column.
7. To allow null values for the column you are defining, in the same row, select the Nulls cell.
8. To bind a default to the column, select the Default cell in the same row. Then, in the Properties box, select the default.
9. To bind a rule to the column, select the Rule cell in the same row. Then, in the Properties box, select the rule.
10. Select another row and repeat steps 4 through 9 for each table column.
11. Choose the Create button.
    The Specify New Table Name dialog box appears.
12. Type a name for the table.
13. Choose OK.

Related Topics

Altering a Table

Copying a Table

Moving a Table to a Segment

Dropping a Table

Viewing Table Properties

Using DBCC CheckTable

Altering a Table

After you have created a table, you can add new columns to it. You must be the table owner to alter the table.

To alter a table

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. In the Tables box, select the table to alter.
3. To rename a column, select the appropriate Column Name cell and then type a new name for the column.
4. To create a new column, in an empty Column Name cell, type a name for the new column.
   You can use the Edit menu to insert or delete row columns between existing columns.
5. In the same row, select the Datatype cell to specify a datatype for the new column. Then, in the Properties box, select a datatype for the new column.
6. If you selected a char, varchar, binary, or varbinary datatype, select the Length cell. Then, in the Properties box, type the maximum number of characters that can be stored in the column.
   Note that for new columns the Nulls cell is always checked when you add a new column to an existing table.
7. To bind a default to the column, select the Default cell. Then, in the Properties box, select the default to bind.
8. To bind a rule to the column, select the Rule cell. Then, in the Properties box, select the rule to bind.
9. Repeat Steps 4 through 8 for each new column.
   You can insert or delete the newly added column by choosing Insert or Delete from the Edit menu.
10. Choose the Alter button.

Related Topics

Creating a Table

Copying a Table

Moving a Table to a Segment

Dropping a Table

Viewing Table Properties

Using DBCC CheckTable

Copying a Table

You can create a new table by copying an existing table's structure, including its columns, data, defaults, rules, and indexes. You can copy a table from the Manage Tables window or from the Copy Table dialog box.

To copy a table from the Manage Tables window

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. In the Tables box, select the table to copy.
3. Change information in the grid for the table, if appropriate.
4. Choose the Save As button.
   The Specify New Table Name dialog box appears.
5. Type a name for the new table.
6. Choose the OK button.
   The new table is created. Note that indexes, keys, permissions, triggers, and data from the existing table are not copied to the new table.

To copy a table from the Copy Table dialog box

1. From the Tools menu, choose Copy Table.
   The Copy Table dialog box appears.
2. Select the table to copy the information from.
3. Type the name of the table to copy the information to.
4. To select data based on a certain criteria, type the appropriate WHERE clause in the Data Selection Criteria box.
5. Choose Copy.
   The new table is created. Note that permissions, triggers, and keys from the existing table are not copied to the new table.

Related Topics

Creating a Table

Altering a Table

Moving a Table to a Segment

Dropping a Table

Viewing Table Properties

Moving a Table to a Segment

After you have created a table, you can move the table and its clustered index, the table's nonclustered indexes, or a selected table index to a different segment.

To place a table and/or its indexes on a segment

1. From the Tools menu, choose Place on Segment.
   The Place Table on Segment dialog box appears.
2. Select the table to move, or select the table whose indexes you want to move.
3. Select the segment to place the information on.
4. Select the items to place on that segment.
5. Choose OK.

Related Topics

Creating a Table

Altering a Table

Copying a Table

Dropping a Table

Viewing Table Properties

Dropping a Table

Dropping a table removes the table, its contents, and all indexes and permissions associated with it. You must be the table or database owner to drop it. You cannot drop a table while it is in use (being read or written to by a user or a front-end program).

To drop a table

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. Select the table to drop.
3. From the Object menu, choose Drop Object.

Related Topics

Creating a Table

Altering a Table

Copying a Table

Moving a Table to a Segment

Viewing Table Properties

Viewing Table Properties

You can view information about the creation date, owner, size, number of columns and rows of a table, as well as information about the columns, indexes, keys, and triggers associated with the table. You can also print or save the information.

To view table properties

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. Select the table to view properties for.
3. From the Object menu, choose Object Properties.
   The Table Properties dialog box appears.
4. Choose Close.

Related Topics

Creating a Table

Altering a Table

Copying a Table

Moving a Table to a Segment

Dropping a Table

Using DBCC CheckTable

DBCC CheckTable allows you to check a table to see that index and data pages are correctly linked, that indexes are in properly sorted order, that all pointers are consistent, that the data information on each page is reasonable, and that page offsets are reasonable. If the log segment is on its own device, running DBCC CheckTable on the *syslogs* table reports the log(s) used and the amount of free space available.

To use DBCC CheckTable

1. From the Tools menu, choose DBCC CheckTable.
   The Select Table for DBCC CheckTable dialog box appears.
2. Select the table to check.
3. Choose OK.
   The results are displayed in the DBCC CheckTable window.

Selecting Data from a Table

You can construct basic SELECT queries using the Select Data from Table dialog box. This dialog box allows you to specify the WHERE, GROUP BY, HAVING, ORDER BY, and COMPUTE clauses of the SELECT statement. You can also save the information as a view.

To select data from a table

1. From the Tools menu, choose Select Data.
   The Select Data dialog box appears.
2. Select the table to query.
3. Select the columns to query and construct the query statement.
4. Save the query as a view, or execute the query.
   If you execute the query, the Select Results window appears and displays the results of the query.

Related Topic

Selecting Data from a View

Creating Indexes

Indexes are created on columns to speed retrieval of data and to impose uniqueness constraints. There are two types of indexes, clustered and nonclustered. You must be the table owner to create an index. You can create an index at any time, even if there is no data in the table.

To create an index

1. From the Manage menu, choose Indexes.
   The Manage Indexes dialog box appears.
2. Select the table to create an index for.
3. Choose New.
   The Specify New Index Name dialog box appears.
4. Type an index name.
5. Choose OK.
6. Select an index type and the segment where the index will reside.
7. Select the table columns to create the index on.
8. Select any applicable row options.
9. Choose Create.
   The Create New Index dialog box appears.
10. Specify any applicable options for the index.
11. Choose Yes.

Related Topics

Modifying an Index

Rebuilding an Index

Renaming an Index

Updating Index Statistics

Dropping an Index

Modifying an Index

You can change the columns that are indexed or change the index options. You must be the table owner to modify an index.

To modify an index

1. From the Manage menu, choose Indexes.
   The Manage Indexes dialog box appears.
2. Select the table to modify the index for.
3. Edit any information as needed.
4. Choose Modify.
   The Save Changes to Existing Index dialog box appears.
5. Specify any applicable options for the index.
6. Choose Yes.

Related Topics

Creating Indexes

Rebuilding an Index

Renaming an Index

Updating Index Statistics

Dropping an Index

Rebuilding an Index

You can rebuild an index at any time. When an index is built, SQL Server drops the existing index and creates a new one.

To rebuild an index

1. From the Manage menu, choose Indexes.
   The Manage Indexes dialog box appears.
2. Select the table to rebuild the index for.
3. Select the index to rebuild.
4. Choose Rebuild.
   The Save Changes to Existing Index dialog box appears.
5. Specify any applicable options for the index.
6. Choose Yes.

Related Topics

Creating Indexes

Modifying an Index

Renaming an Index

Updating Index Statistics

Dropping an Index

Renaming an Index

If you are the table owner, you can rename an index.

To rename an index

1. From the Manage menu, choose Indexes.
   The Manage Indexes dialog box appears.
2. Select the table that contains the index to rename.
3. Select the index to rename.
4. Choose Rename.
   The Rename Index dialog box appears.
5. Type a new name for the index.
6. Choose OK.

Related Topics

Creating Indexes

Modifying an Index

Rebuilding an Index

Updating Index Statistics

Dropping an Index

Updating Index Statistics

You can update information about the distribution of key values in specified indexes.

To update index statistics

1. From the Tools menu, choose Update Statistics.
   The Update Table Statistics dialog box appears.
2. Select the table to update index statistics for.
3. Select the index to update statistics on.
4. Choose OK.

Related Topics

Creating Indexes

Modifying an Index

Rebuilding an Index

Renaming an Index

Dropping an Index

Dropping an Index

When you drop an index, SQL Server removes the specified index(es) from the database and reclaims the associated storage space.

To drop an index

1. From the Manage menu, choose Indexes.
   The Manage Indexes dialog box appears.
2. Select the table that contains the index to drop.
3. Select the name of the index to drop.
4. Choose Drop.

Related Topics

Creating Indexes

Modifying an Index

Rebuilding an Index

Renaming an Index

Updating Index Statistics

Creating a Trigger

A trigger is a special type of stored procedure that is automatically invoked when data is changed in the database. Triggers are frequently used to enforce the integrity of the database. A trigger executes automatically when a user attempts to modify data on a table specified in the trigger. There be only one Insert, Update, and Delete trigger per table. If a trigger exists for the table and you create another trigger, the existing trigger is replaced.

To create a trigger

1. From the Manage menu, choose Triggers.
   The Manage Triggers window appears.
2. In the Tables box, select the table to create the trigger for.
3. In the Triggers box, select the type of trigger to create. Or, from the File menu, choose New.
   The Create Trigger on Table dialog box appears.
4. In the Table Name box, select the table to create a trigger for.
   To view information about the table, choose the Table Info button.
5. In the Trigger Name box, type a trigger name.
6. Select one or more TransactSQL statements to associate the trigger with. The name of any existing trigger of that type is displayed to the right of that type.

- To associate the trigger with Insert statements, select For Insert.
- To associate the trigger with Update statements, select For Update.
- To associate the trigger with Delete statements, select For Delete.

For update and delete triggers, you can select either restrictive or cascading referential integrity. If you select the Cascade option button, changes to the primary key tables are also made to the foreign key tables. If you select the Restrict option button, changes to the primary key tables are disallowed if a foreign key exists with the same value as the primary key. The default for Update is Cascade. The default for Delete is Restrict.

7. If you are creating an Insert or Update trigger and want to specify the update status of specific table columns, in the Column Name box, select the columns to check and then select the And or Or button to state how the column will be related to other columns being checked for update status. If there are no other columns, the And or Or buttons are not used.
8. Choose the OK button.
   The TransactSQL script for the trigger appears in the Manage Triggers window. If no primary or foreign keys exist between the current table and any other table or view, a template of the CREATE TRIGGER statement appears.
9. Modify the information as necessary.
10. Save the trigger.

- To create the trigger on the server, choose the Execute button.


To save the trigger to a data file, from the File menu choose Save As. If you save the trigger to a file, it is not created on the server. You must execute the trigger in order for it to take effect.

Related Topics

Editing a Trigger

Dropping a Trigger

Viewing Trigger Properties

Editing a Trigger

After you have created a trigger, you can change its contents.

To edit a trigger

1. From the Manage menu, choose Triggers.
   The Manage Triggers window appears.
2. In the Tables box, select the table that contains the trigger to edit.
3. In the Triggers box, select the existing trigger to edit.
   The contents of the trigger appear in the window.
4. Edit the trigger.
5. Save the trigger.
   - To save the trigger to the server, choose Execute.
   - To save the trigger to a data file, from the File menu choose Save As. If you save the trigger to a file, the changes do not take effect on the server. You must execute the trigger in order for the changes to the trigger to take effect.

Related Topics

Creating a Trigger

Dropping a Trigger

Viewing Trigger Properties

Dropping a Trigger

If a trigger is no longer applicable, you can remove it.

To drop a trigger

1. From the Manage menu, choose Triggers.
   The Manage Triggers window appears.
2. Select the table that contains the trigger to drop.
3. Select the trigger to drop.
4. From the Object menu, choose Drop Object.

Related Topics

Creating a Trigger

Editing a Trigger

Viewing Trigger Properties

Viewing Trigger Properties

You can view information about a trigger -- its type, creation date, the table it is on, the table owner, and the objects referenced by the trigger's SQL statements. You can also print and save the information.

To view trigger properties

1. From the Manage menu, choose Triggers.
   The Manage Triggers window appears.
2. Select the table that contains the trigger to view properties for.
3. Select the trigger whose properties you want to view.
4. From the Object menu, choose Object Properties.
   The Trigger Properties dialog box appears.
5. Choose Close.

Related Topics

Creating a Trigger

Editing a Trigger

Dropping a Trigger

Creating a View

A view is an alternative way of looking at the data in one or more tables. It is usually a subset of columns from one or more tables in the database.

To create a view

1. From the Manage menu, choose Views.
   The Manage Views window appears.
2. From the File menu, choose New or Select (New View) from the Views box.
   The Create View dialog box appears.
3. In the Tables box, select the table to create the view from.
4. Select the columns to build the view with and select the function to apply to the column, if applicable.
5. Select the search condition to use with the query, if applicable.
6. Choose Add To Clause.
   The table column is added to the SELECT statement.
7. Add appropriate conditions to the query.
8. Choose Save View.
   The Save Select As View dialog box appears.
9. Type a name for the view and then choose OK.
   The generated text of the view appears in the Manage Views window.
10. Choose Execute to execute the view. Or, to save the view to a file, from the File menu, choose Save As.
    If you save the view to a file, it is not available to the server. You must execute the view for it to take effect.

Related Topics

Editing a View

Dropping a View

Viewing View Properties

Selecting Data from a View

Editing a View

After you have created a view, you can add or remove columns or change the SELECT statement the view is based on.

To edit a view

1. From the Manage menu, choose Views.
   The Manage Views window appears.
2. In the Views box, select the view to edit.
   The text of the view appears in the Views window.
3. Edit the appropriate information.
4. Choose Execute to execute the view. Or, to save the view to a file, from the File menu choose Save As.
   You must execute the view in order for the view to be updated on the server.

Related Topics

Creating a View

Dropping a View

Viewing View Properties

Selecting Data from a View

Dropping a View

If you no longer need a view, you can drop it. All permissions on that view are dropped as well.

To drop a view

1. From the Manage menu, choose Views.
   The Manage Views window appears.
2. Select the view to drop.
3. From the Object menu, choose Drop Object.

Related Topics

Creating a View

Editing a View

Viewing View Properties

Selecting Data from a View

Viewing View Properties

You can view information about a view's owner, its creation date, the base tables and views it is created from, and the columns of the tables and views it includes. You can print or save this information.

To view view properties

1. From the Manage menu, choose Views.
   The Manage Views window appears.
2. Select the view whose properties you want to view.
3. From the Object menu, choose Object Properties.
4. Choose Close.

Related Topics

Creating a View

Editing a View

Dropping a View

Selecting Data from a View

Selecting Data from a View

You can construct basic SELECT queries using the Select Data from Table dialog box. This dialog box allows you to specify the WHERE, GROUP BY, HAVING, ORDER BY, and COMPUTE clause of the SELECT statement. You can save the information as a different view.

To select data from a view

1. From the Tools menu, choose Select Data.
   The Select Data dialog box appears.
2. Select the view to query.
3. Select the columns to query and construct the query statement.
4. Save the query as a view, or execute the query.
   If you execute the query, the Select Results window appears and displays the results of the query.

Related Topics

Creating a View

Editing a View

Dropping a View

Viewing View Properties

Selecting Data from a Table

Creating a Key

You can create three types of keys for tables and views: primary, foreign, and common.

To create a key

1. From the Manage menu, choose Keys.
   The Manage Keys dialog box appears.
2. Select a key type.
3. Create the key.

To create a primary key, select the table or view and the column to associate the key with.

To create a foreign key, select the working table/view to associate the key with and then select the related table/view that foreign key is the primary key for. Then select the column to associate the key with.

To create a common key, select the working table/view and columns for the common key. Then, in the selected table box, select the common table and then select the column to associate the common key with.

4. Choose Create.
5. Choose Close.

Related Topics

Modifying a Key

Dropping a Key

Modifying a Key

After you have created a key on a table or a view, you can modify it.

To modify a key

1. From the Manage menu, choose Keys.
   The Manage Keys dialog box appears.
2. Select the Working Table key type and its related table (if it's a foreign or common key).
3. Modify the appropriate information.
4. Choose Modify.
5. Choose Close.

Related Topics

Creating a Key

Dropping a Key

Dropping a Key

If you no longer need a key, you can drop it.

To drop a key

1. From the Manage menu, choose Keys.
   The Manage Keys dialog box appears.
2. Select the working table/view.
3. Select the type of key to drop.
4. Choose Drop.
5. Choose Close.

Related Topics

Creating a Key

Modifying a Key

Creating a Rule

A rule determines what data can be entered in a particular column of a table. A rule can also be bound to a user-defined datatype so that any columns that use the datatype automatically have the rule associated with them.

To create a rule

1. From the Manage menu, choose Rules.
   The Manage Rules dialog box appears.
2. Type the contents of the rule. To enter a line break, press CTRL+ENTER.
3. Choose Create.
4. Choose Close.

Related Topics

Binding a Rule

Unbinding a Rule

Dropping a Rule

Displaying Rule Bindings Information

Viewing Rule Properties

Binding a Rule

You can bind a rule to a table column or to a user-defined datatype. You can bind a rule to a table column from the Manage Tables window or from the Manage Rules dialog box. You can bind a rule to a user-defined datatype from the Manage Rules or Manage Datatypes dialog boxes. Note that if you intend to bind the rule only to future columns the user-defined datatype is associated with, you must bind the rule in the Manage Rules dialog box to specify future only.

To bind a rule from the Manage Tables window

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. Select the table to bind the rule on.
3. Select the Rule cell of the column row to bind the rule to.
4. In the Properties box, select the rule to bind. The full name appears in the Rule cell.
5. Choose Alter.

To bind a rule from the Manage Rules dialog box

1. From the Manage menu, choose Rules.
   The Manage Rules dialog box appears.
2. Select the rule to bind.
3. Select the table and column or the user-defined datatype to bind the rule to.
   To bind the rule only on future columns that the user-defined datatype is associated with, select Futureonly.
4. Choose Bind.
5. Choose Close.

To bind a rule from the Manage Datatypes dialog box

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the Rule cell of the datatype row to bind the rule to.
3. In the Properties box, select the rule to bind.
4. Choose Modify.

Related Topics

Creating a Rule

Unbinding a Rule

Dropping a Rule

Displaying Rule Bindings Information

Viewing Rule Properties

Unbinding a Rule

If you no longer want a table column or user-defined datatype to have a rule, you can unbind it. You can unbind a rule bound to a table column from the Manage Tables window or from the Manage Rules dialog box. You can unbind a rule bound to a user-defined datatype from the Manage Rules or Manage Datatypes dialog boxes. You can also unbind a rule from all its bindings from the Current Bindings dialog box.

To unbind a rule from the Manage Tables window

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. Select the table to unbind the rule from.
3. Select the Rule cell of the column row to unbind the rule from.
4. In the Properties box, select (none).
5. Choose Alter.

To unbind a rule from the Manage Rules dialog box

1. From the Manage menu, choose Rules.
   The Manage Rules dialog box appears.
2. Select the rule to unbind.
3. Select the table and column or the user-defined datatype to unbind the rule from.
4. Choose Unbind.
5. Choose Close.

To unbind a rule from the Manage Datatypes dialog box

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the Rule cell of the datatype row to unbind the rule from.
3. In the Properties box, select (none).
4. Choose Modify.

To unbind a rule from all its bindings

1. From the Manage menu, choose Rules.
   The Manage Rules dialog box appears.
2. Select the rule to unbind all bindings from.
3. Choose Info.
   The Current Bindings dialog box appears.
4. Choose Unbind All.

Related Topics

Creating a Rule

Binding a Rule

Dropping a Rule

Displaying Rule Bindings Information

Viewing Rule Properties

Dropping a Rule

If you no longer need a rule, you can drop it. You can drop a rule from the Database Objects window or from the Manage Rules dialog box. Before you can drop it, a rule must be unbound from all columns and user-defined datatypes.

To drop a rule from the Database Objects window

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. Select the rule to drop.
3. From the Object menu, choose Drop Object.

To drop a rule from the Manage Rules dialog box

1. From the Manage menu, choose Rules.
   The Manage Rules dialog box appears.
2. Select the rule to drop.
3. Choose Drop.
4. Choose Close.

Related Topics

Creating a Rule

Binding a Rule

Unbinding a Rule

Displaying Rule Bindings Information

Viewing Rule Properties

Displaying Rule Bindings Information

You can display information about the table columns and/or user-defined datatypes a rule is bound to. You can also print or save the information.

To display rule bindings information

1. From the Manage menu, choose Rules.
   The Manage Rules dialog box appears.
2. Select the rule to display binding information for.
3. Choose Info.
   The Current Bindings dialog box appears.
   You can unbind the rule from all its bindings by choosing Unbind All.
4. Choose Close.

Related Topics

Creating a Rule

Binding a Rule

Unbinding a Rule

Dropping a Rule

Viewing Rule Properties

Viewing Rule Properties

You can display information about a rule, its owner, creation date, contents, and what user-defined datatypes or table columns are bound to it. You can also print or save the information.

To view rule properties

1. From the Microsoft SQL Object Manager window, choose the Object button.
   The Database Objects window appears.
2. Select the rule to view properties for.
3. From the Object menu, choose Object Properties.
   The Rule Properties dialog box appears.
4. Choose Close.

Related Topics

Creating a Rule

Binding a Rule

Unbinding a Rule

Dropping a Rule

Displaying Rule Bindings Information

Creating a Default

A default is created for a column or a user-defined datatype and causes SQL Server to automatically supply a specified value if the user does not explicitly provide a value for the column or the user-defined datatype. Note that if you intend to bind the default only to future columns the user-defined datatype is associated with, you must bind the default in the Manage Defaults dialog box to specify future only.

To create a default

1. From the Manage menu, choose Defaults.
   The Manage Defaults dialog box appears.
2. Type the contents of the default using Transact-SQL syntax. To enter a line break, press CTRL+ENTER.
3. Choose Create.
4. Choose Close.

Related Topics

Binding a Default

Unbinding a Default

Dropping a Default

Displaying Default Bindings

Viewing Default Properties

Binding a Default

You can bind a default to a table column or to a user-defined datatype. You can bind a default to a table column from the Manage Tables window or the Manage Defaults dialog box. You can bind a default to a user-defined datatype from the Manage Datatypes or Manage Defaults dialog box.

To bind a default from the Manage Tables window

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. Select the table whose columns you want to bind the default to.
3. Select the Default cell of the column row to bind the default to.
4. In the Properties box, select the default to bind.
   The default name appears in the Default cell.
5. Choose Alter.

To bind a default from the Manage Defaults dialog box

1. From the Manage menu, choose Defaults.
   The Manage Defaults dialog box appears.
2. Select the default to bind.
3. Select the table column or the user-defined datatype to bind the default to.
   To bind the default only on future columns associated with the user-defined datatype, select Futureonly.
4. Choose Bind.
5. Choose Close.

To bind a default from the Manage Datatypes dialog box

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the Default cell of the datatype row to bind the default to.
3. In the Properties box, select the default to bind.
   The default name appears in the Default cell.
4. Choose Modify.

Related Topics

Creating a Default

Unbinding a Default

Dropping a Default

Displaying Default Bindings

Viewing Default Properties

Unbinding a Default

If you no longer want a default to be bound to a table column or to a user-defined datatype, you can unbind it. You can unbind a default bound to a table column from the Manage Tables window or from the Manage Defaults dialog box. You can unbind a default bound to a user-defined datatype from the Manage Defaults or Manage Datatypes dialog boxes. You can also unbind a default from all its bindings from the Current Bindings dialog box.

To unbind a default from the Manage Tables window

1. From the Manage menu, choose Tables.
   The Manage Tables window appears.
2. Select the table whose columns you want to unbind the default from.
3. Select the Default cell of the column row to unbind the default from.
4. In the Properties box, select (none).
5. Choose Alter.

To unbind a default from the Manage Defaults dialog box

1. From the Manage menu, choose Defaults.
   The Manage Defaults dialog box appears.
2. Select the default to unbind.
3. Select the table column or the user-defined datatype to unbind the default from.
4. Choose Unbind.
5. Choose Close.

To unbind a default from the Manage Datatypes dialog box

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the Default cell of the datatype row to unbind the default from.
3. In the Properties box, select (none).
4. Choose Modify.

To unbind a default from all its bindings

1. From the Manage menu, choose Defaults.
   The Manage Defaults dialog box appears.
3. Select the default to unbind all bindings from.
4. Choose Info.
   The Current Bindings dialog box appears.
5. Choose Unbind All.

Related Topics

Creating a Default

Binding a Default

Dropping a Default

Displaying Default Bindings

Viewing Default Properties

Dropping a Default

If you no longer need a default, you can drop it. You can drop a default from the Database Objects window or from the Manage Defaults dialog box.

To drop a default from the Database Objects window

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. Select the default to drop.
3. From the Manage menu, choose Drop Object.

To drop a default from the Manage Defaults dialog box

1. From the Manage menu, choose Defaults.
   The Manage Defaults dialog box appears.
2. Select the default to drop.
3. Choose Drop.
4. Choose Close.

Related Topics

Creating a Default

Binding a Default

Unbinding a Default

Displaying Default Bindings

Viewing Default Properties

Displaying Default Bindings

You can display the table columns and/or the user-defined datatypes that a default is bound to. You can also print or save the information.

To display default bindings

1. From the Manage menu, choose Defaults.
   The Manage Defaults dialog box appears.
2. Choose Info.
   The Current Bindings dialog box appears. You can unbind the default from all its bindings by choosing the Unbind All button.
3. Choose Close.

Related Topics

Creating a Default

Binding a Default

Unbinding a Default

Dropping a Default

Viewing Default Properties

Viewing Default Properties

You can view information about a default -- its owner, its creation date, the contents of the default, and its current bindings. You can also print or save the information.

To view default properties

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. Select the default to view properties for.
3. From the Object menu, choose Object Properties.
   The Default Properties dialog box appears.
4. Choose Close.

Related Topics

Creating a Default

Binding a Default

Unbinding a Default

Dropping a Default

Displaying Default Bindings

Creating a User-defined Datatype

A user-defined datatype is created by a user to specify what type of data a column can contain. After you create a user-defined datatype, you can bind rules and defaults to the datatype. Note that if you intend to bind the rule or default only to future columns the user-defined datatype is associated with, you must bind the rule or default in the Manage Rules or Manage Defaults dialog boxes to specify future only.

To create a user-defined datatype

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Choose New.
   The Specify New Type Name dialog box appears.
3. Type the datatype name.
4. Choose OK.
5. Select the Type cell of the new datatype row.
6. In the Properties box, select a system-defined datatype to base the user-defined datatype on.
7. Specify the length, if necessary for the datatype, and whether the datatype will allow null values.
8. To bind a rule to the datatype, select the Rule cell. Then, in the Properties box, select the rule.
9. To bind a default to the datatype, select the Default cell. Then, in the Properties box, select the default.
10. Choose Create.
11. Choose Close.

Related Topics

Renaming a User-defined Datatype

Dropping a User-defined Datatype

Viewing User-defined Datatype Properties

Binding a Rule

Binding a Default

Renaming a User-defined Datatype

You can rename a datatype you have created.

To rename a user-defined datatype

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the datatype to rename by clicking on the Name cell.
3. Choose Rename.
   The Specify New Type Name dialog box appears.
4. Type a new name for the datatype.
5. Choose OK.
6. Choose Close.

Related Topics

Creating a User-defined Datatype

Dropping a User-defined Datatype

Viewing User-defined Datatype Properties

Dropping a User-defined Datatype

If you no longer need a user-defined datatype, you can drop it. However, you cannot drop a user-defined datatype that is already in use in a table.

To drop a user-defined datatype

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the row of the datatype to drop.
3. Choose Drop.
4. Choose Done.

Related Topics

Creating a User-defined Datatype

Renaming a User-defined Datatype

Viewing User-defined Datatype Properties

Viewing User-defined Datatype Properties

You can view information about a user-defined datatype, its owner, its physical type, its length, whether it allows null values, if it is bound to a rule or to a default, and the table columns that it is associated with. You can also print or save the information.

To view user-defined datatype properties

1. From the Manage menu, choose User Defined Datatypes.
   The Manage Datatypes dialog box appears.
2. Select the user-defined datatype to view properties for.
3. Choose the Info button.
   The User Datatype Properties dialog box appears.
4. Choose Close.

Related Topics

Creating a User-defined Datatype

Renaming a User-defined Datatype

Dropping a User-defined Datatype

Creating a Stored Procedure

A stored procedure is a collection of SQL statements and optional control-of-flow language. Stored procedures are compiled the first time they are executed and their execution plans are stored so that subsequent execution is very fast. Once defined, a stored procedure can be executed by name.

To create a stored procedure

1. From the Manage menu, choose Stored Procedures.
   The Manage Stored Procedures window appears and displays the CREATE PROCEDURE statement.
2. Add the appropriate information.
3. To execute the stored procedure, choose Execute. To save the stored procedure to a file, from the File menu choose Save As.
   You must execute the stored procedure in order for it to be available to the server.

Related Topics

Editing a Stored Procedure

Dropping a Stored Procedure

Viewing Stored Procedure Properties

Editing a Stored Procedure

You can edit the contents of a stored procedure.

To edit a stored procedure

1. From the Manage menu, choose Stored Procedures.
   The Manage Stored Procedures window appears.
2. Select the stored procedure to edit.
   The contents of the stored procedure appear in the Manage Stored Procedures window.
3. Edit the contents of the stored procedure.
4. To execute the stored procedure, choose Execute. To save the stored procedure to a file, from the File menu choose Save As.
   You must execute the stored procedure in order for it to be available to the server.

Related Topics

Creating a Stored Procedure

Dropping a Stored Procedure

Viewing Stored Procedure Properties

Dropping a Stored Procedure

If you no longer need a stored procedure, you can drop it. You must be in the database the stored procedure was created in to drop it.

To drop a stored procedure

1. From the Manage menu, choose Stored Procedures.
   The Manage Stored Procedures window appears.
2. Select the stored procedure to drop.
3. From the Object menu, choose Drop Object.

Related Topics

Creating a Stored Procedure

Editing a Stored Procedure

Viewing Stored Procedure Properties

Viewing Stored Procedure Properties

You can view information about a stored procedure -- its creator, creation date, and parameters. You can also print or save the information.

To view stored procedure properties

1. From the Manage menu, choose Stored Procedures.
   The Manage Stored Procedures window appears.
2. Select the stored procedure whose properties you want to view.
3. From the Object menu, choose Object Properties.
   The Stored Procedure Properties dialog box appears.
4. Choose Close.

Related Topics

Creating a Stored Procedure

Editing a Stored Procedure

Dropping a Stored Procedure

Granting Object Permissions

You can grant users permissions for tables, views, and stored procedures. If users already have permissions on those objects, any new permissions that you grant will replace of the existing permissions. You can also grant permissions to user groups you have created.

To grant object permissions

1. From the Object menu, choose Object Permissions.
   The Object Permissions dialog box appears.
2. Select the object to grant permissions for.
3. Select the type of permissions to grant and, if necessary, the columns the permissions are for.
4. Select the user(s) or group(s) to grant permission to.
5. Choose Add.
   Repeat steps 4 and 5 for each user or group.
6. Choose Grant.
   Repeat steps 2 through 6 for each set of permissions to grant.
7. Choose Close.

Related Topics

Revoking Object Permissions

Viewing Permissions for Objects

Viewing a User's Permissions

Viewing a User's Permissions on All Objects

Revoking Object Permissions

If you don't want a user to have object permissions, you can revoke them.

To revoke object permissions

1. From the Object menu, choose Object Permissions.
   The Object Permissions dialog box appears.
2. Select the user or group from which to revoke permissions.
3. Choose Add.
4. In the Users with Selected Permissions box, select the user.
5. Remove the appropriate permissions by selecting the appropriate check boxes.
6. Choose Revoke.
7. Choose Close.

Related Topics

Granting Object Permissions

Viewing Permissions for Objects

Viewing a User's Permissions

Viewing a User's Permissions on All Objects

Viewing Permissions for Objects

You can view users who have specific permissions on a table, view, or stored procedure.

To view object permissions

1. From the Objects menu, choose Object Permissions.
   The Object Permissions dialog box appears.
2. Select the object to view permissions for.
3. Select the type of permissions to view.
4. Choose Show.
   Users with those permissions appear in the Users with Permissions box.
5. Choose Close.

Related Topics

Granting Object Permissions

Revoking Object Permissions

Viewing a User's Permissions

Viewing a User's Permissions on All Objects

Viewing a User's Permissions

You can view a user's permissions for a table, view, or stored procedure.

To view a user's permissions

1. From the Object menu, choose Object Permissions.
   The Object Permissions dialog box appears.
2. Select the object to view the user's permissions for.
3. Select the user.
4. Choose User Zoom.
   The permissions the user has are displayed.
5. Choose Close.

Related Topics

Granting Object Permissions

Revoking Object Permissions

Viewing Permissions for Objects

Viewing a User's Permissions on All Objects

Viewing a User's Permissions on All Objects

You can view a user's permissions on all database objects. You can also print or save the information.

To view a user's permissions on all objects

1. From the Object menu, choose Object Permissions.
   The Object Permissions dialog box appears.
2. Select the user to view all permissions for.
3. Choose User Info.
   The User Permissions dialog box appears.
4. Choose Close.

Related Topics

Granting Object Permissions

Revoking Object Permissions

Viewing Permissions for Objects

Viewing a User's Permissions

Viewing Object Dependencies

You can view an object's dependencies on other objects as well as those objects that depend on the object. You can also print or save the information.

To view object dependencies

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. From the Object menu, choose Object Dependencies.
   The Object Dependencies dialog box appears.
3. Select the object to view dependencies for.
   The dependency information appears.
4. Choose Close.

Related Topics

Viewing Object Properties

Renaming an Object

Dropping an Object

Viewing Object Properties

You can view information about an object -- its owner, creation date, base tables and views, and columns. You can also print or save the information.

To view object properties

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. Select the object to view properties for.
3. From the Object menu, choose Object Properties.
   The Object Properties dialog box appears.
4. Choose Close.

Related Topics

Viewing Object Dependencies

Renaming an Object

Dropping an Object

Renaming an Object

You can change the name of an object. When you change the name of any object, you must rebuild any views, stored procedures, or triggers that reference the object.

To rename an object

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. Select the object to rename.
3. From the Object menu, choose Rename Object.
   The Rename Object dialog box appears.
4. Type a new name for the object.
5. Choose OK.

Related Topics

Viewing Object Dependencies

Viewing Object Properties

Dropping an Object

Dropping an Object

If you no longer need an object, you can remove it.

To drop an object

1. From the Microsoft SQL Object Manager window, choose the Objects button.
   The Database Objects window appears.
2. Select the object to drop.
3. From the Object menu, choose Drop Object.

Related Topics

Viewing Object Dependencies

Viewing Object Properties

Renaming an Object

Transferring Data to a SQL Server

From data files, you can transfer data to a SQL Server. You can transfer data into a new table in the database, or you can transfer data into an existing table. You must have Insert permission for the table into which you are copying data.

To transfer data into a new table

1. From the Microsoft SQL Object Manager window, choose the Transfer button.
   The Transfer Data dialog box appears.
2. Select Import into SQL Server.
3. Choose the Define New Table button.
   The Create New Table dialog box appears.
4. Type a name for the new table.
5. Create the columns for the table.
6. Choose Create.
   The Transfer Data dialog box reappears.
7. Select the type of file to transfer the data from.
   If you choose BCP defined, choose Column Mapping to define the columns.
8. Select the directory and filename to transfer data from.
   If you choose BCP Format File, choose Format File to select the format file to read.
   You can also create a format file for transfer types other than BCP Format File by choosing Format File.
9. To set data transfer options, choose Options.
10. To begin the data transfer, choose Import.
    During the transfer, the status line in the Microsoft SQL Object Manager window periodically changes to inform you of how many rows are successfully transferred to the server. A summary count is displayed when the transfer is complete. If the data transfer encounters any errors, the Data Transfer Error Report window appears and displays the error.

To transfer data into a selected table

1. From the Microsoft SQL Object Manager window, choose the Transfer button.
   The Transfer Data dialog box appears.
2. Choose Import into SQL Server.
3. Select the table to transfer data into.
4. Select the type of file to transfer the data from.
   If you choose BCP defined, choose Column Mapping to define the columns.
5. Select the directory and filename to transfer data from.
   If you choose BCP Format File, choose Format File to select the format file to read.
   You can also create a format file for transfer types other than BCP Format File by choosing Format File.
6. To set data transfer options, choose Options.
7. To begin the data transfer, choose Import.
   During the transfer, the status line in the Microsoft SQL Object Manager window periodically changes to inform you of how many rows are successfully transferred to the server. A summary count is displayed when the transfer is complete. If the data transfer encounters any errors, the Data Transfer Error Report window appears and displays the errors.

Related Topic

Transferring Data from a SQL Server

Transferring Data from a SQL Server

You can transfer data from a SQL Server to a data file.

To transfer data from a SQL Server

1. From the Microsoft SQL Object Manager window, choose the Transfer button.
   The Transfer Data dialog box appears.
2. Choose Export from SQL Server.
3. Select the table to transfer data from.
4. Select the type of file to transfer.
   If you choose BCP defined, choose the Column Mapping button to define the columns.
5. Select or type the name of the file and directory to transfer the data to.
   If you choose BCP format as the file type, choose Format File to select the format file to read for the data transfer.
   You can also create a format file for transfer types other than BCP format file by choosing Format File.
6. To set the data transfer options, choose Options.
7. To begin the data transfer, choose Export.
   During the transfer, the status line in the Microsoft SQL Object Manager window periodically changes to inform you of how many rows are successfully transferred to the server. A summary count is displayed when the transfer is complete. If the data transfer encounters any errors, the Data Transfer Error Report window appears and displays the errors.

Related Topic

Transferring Data to a SQL Server

Generating SQL Scripts

SQL Scripts are descriptions of statements used to create the objects of the database. You can generate scripts for an entire database or for particular database objects.

To generate database scripts

1. From the Microsoft SQL Object Manager window, choose the Scripts button.
   The Generate SQL Scripts dialog box appears.
2. Set the appropriate options.
3. Type the name of the person generating the scripts.
4. Choose Generate.
   The Save As dialog box appears.
5. Type a name for the file in which to put the Transact-SQL code.
6. Choose OK.
   The script is generated and stored in the selected file.

Querying a Database

Using Transact-SQL statements, you can query a database. While SQL Server processes a query, you can perform other administration tasks. When the query results start to return to the workstation, the results window is displayed and becomes the active window. You can edit, print, or save the query results.

To query a database

1. From the Microsoft SQL Object Manager window, choose the Query button.
   The Query window appears.
2. Using Transact-SQL syntax, type your query in the Query window.
3. Choose Execute, or press ALT+X or CTRL+E. To execute only a portion of a query, select that portion and choose Execute.

Related Topics

Using a Saved Query

Saving a Query to a File

Setting Query Options

Editing Query Results

Saving a Query to a File

If you plan to use a query more than once, you can save the query to a file.

To save a query to a file

1. From the Microsoft SQL Object Manager window, choose the Query button.
   The Query window appears.
2. Using Transact-SQL syntax, type a query in the Query window.
3. From the File menu, choose Save As.
   The Save As dialog box appears.
4. Type a filename.
5. Choose OK.

Related Topics

Querying a Database

Using a Saved Query

Setting Query Options

Editing Query Results

Using a Saved Query

You can execute queries that have been saved to a file.

To use a saved query

1. From the Microsoft SQL Object Manager window, choose the Query button.
   The Query window appears.
2. From the File menu, choose Open.
   The Open dialog box appears.
3. Select a file.
4. Choose OK.
   The query is appears in the Query window.
5. Choose the Execute button, or press ALT+X or CTRL+E.

Related Topics

Querying a Database

Saving a Query to a File

Setting Query Options

Editing Query Results

Setting Query Options

You can set the options for how a query is processed each time you execute the query.

To set query options

1. From the Microsoft SQL Object Manager window, choose the Query button.
   The Query window appears.
2. From the Query menu, choose Set Options.
   The Query Options dialog box appears.
3. Set the appropriate options.
4. Choose OK.

Related Topics

Querying a Database

Using a Saved Query

Saving a Query to a File

Editing Query Results

Editing Query Results

After you have executed a query and the results window appears, you can edit the contents of that window. You can also print or save all or a portion of the information.

To edit query results

1. From the Microsoft SQL Object Manager window, choose the Query button.
   The Query window appears.
2. Using Transact-SQL syntax, type a query in the Query window.
3. Choose the Execute button.
4. When the results appear, edit the information.

Related Topics

Querying a Database

Using a Saved Query

Saving a Query to a File

Setting Query Options

Analyzing Queries with Graphical Showplan

Graphical Showplan shows the query execution plan for a query. You can generate the query plan without executing the query by choosing the NO EXEC option. You can also generate a textual execution plan of the query by setting Show Query Plan in the Query Options dialog box.

To use graphical showplan

1. Type a query in the Query window.
2. Choose the Showplan button.

   To generate Showplan results without executing the query, choose the NO EXEC button.
3. Choose the Execute button, or press ALT+X or CTRL+E.

   The Query Plan window appears. You can change the font of the query plan text, zoom out to view an overview of the query execution plan, zoom in to view the detailed query execution plan, and display the steps of the query plan.

Details About the Query Chart

Details About the Query Plan Steps

Related Topic

Analyzing Queries with Graphical Statistics I/O

Details About the Query Chart

The Query Plan chart shows the name of the table(s) used, the method used to access the table(s), and the join path between tables. SQL Server uses one of three methods to access the table(s):

- Using the clustered index (indicated by a green border)
- Using a specified index (indicated by a yellow border)
- Using a table scan, in which no indexes are used and the table is scanned sequentially (indicated by a red border)

Details About the Query Plan Steps

The query plan steps contain three parts: the result number, the steps that indicate subprocesses that occur during the query execution, and the action. The steps indicate the path that the optimizer chose to most efficiently process the query. Multistep processes usually occur when data must first be gathered and then sorted in a worktable. The action indicates the processing occurring on the tables listed in the chart.

Action	Result
SELECT	The query will retrieve data from the tables listed.
SELECT (into worktable)	The query will retrieve data from the tables listed and put the data into a worktable for further processing.
INSERT	The query is inserting data into the listed tables.
DELETE	The query is deleting data from the listed tables.
UPDATE	The query is updating data in the listed tables.
/Sorting	The query will sort the data in the listed tables.

Analyzing Queries with Graphical Statistics I/O

Graphical Statistics I/O depicts the amount of disk I/O required to access data for a query. You must execute the query in order to get Statistics I/O information, so the NO EXEC option must be turned off. You can also generate a textual description of the Statistics I/O by choosing Show Stats I/O in the Query Options dialog box. There are three I/O statistics: Scan Count, Logical Reads, and Physical Reads.

To use graphical Statistics I/O

1. Type a query in the Query window.
2. Choose the Statistics I/O button.
3. Choose the Execute button, or press ALT+X or CTRL+E.

   The Statistics I/O window appears. You can change the colors of the bar chart, change the text font, and display a legend for the chart.

Related Topics

Analyzing Queries with Graphical Showplan

Scan Count

Indicates the number of times the query engine "visited" the table during the process of retrieving the data. Generally, for single table queries, the table is scanned once. For "lookup tables" (tables that are joined during processing in order to get more information), the scan count indicates the number of times the table was accessed.

Logical Reads

Indicates the number of index and data pages that were required to get the data needed for the query. This does not equal to the number of rows retrieved for a table, because a page of data can contain many rows.

Physical Reads

Indicates the number of logical reads that actually resulted in data access on the disk. The difference between logical and physical reads indicates the amount of data that was cached. Tables and worktables can appear multiple times in order to process all the data required for the query.

To log in to a SQL Server

1. In the Server box, select or type the name of a server.
   The Server box contains a list of the last 5 servers logged on to. To get a list of servers on the network, choose the List Servers button.
2. In the Login ID box, type your login ID.
   If you selected one of the servers listed in the Server box, the last login ID you used to log in to that server is automatically supplied.
3. In the Password box, type your password.
4. If the server is not started, choose the Start Server before Connect box to start the server.
5. Choose the Connect button.

To disconnect from a SQL Server

1. In the Server box, select the server.
2. Choose the Disconnect button.

To set preferences

- To view queries sent from SQL Object Manager to SQL Server from a dialog box, select the Show Generated T-SQL check box. Or clear the check box to not view the queries sent.

  This option allows you to view the queries generated by SQL Object Manager based on your selections in a dialog box.
- To specify the default directory in which to save files, select the directory in the File Directory box.
- To write the Transact-SQL statements used to create objects to a file, select the T-SQL Archive Filename check box and specify the name of the file.

  This option allows you to keep a record of what you did to create the objects in your database.
- To be prompted before closing SQL Object Manager, select the Prompt before Exiting check box.

  If this box is not selected, SQL Object Manager automatically closes all active connections when you exit SQL Object Manager.

To set all options to their default values

- Choose the Defaults button.

To set configuration information

- In the Login box, type the amount of time to elapse before the server terminates a login attempt.
- In the Query box, type the amount of time for the server to retrieve query results.
- In the System Functions box, type the amount of time the server should wait for major functions to complete before timing out.
- Select the ANSI to OEM check box to activate the AutoANSItoOEM conversion.

To set all options to their default values

- Choose the Defaults button.

To copy a table

1. In the Table Name box, select the table to copy.
   To view information about the table, choose the Table Info button.
2. In the Destination Table box, type the name of the table to copy the information to.
3. Under Copy Options, select the options you want.
   - To copy all columns to the new table, select All Columns.
   - To copy the table structure with its data, select Copy Data.
   - To copy the defaults associated with the table, select Defaults.
   - To copy the rules associated with the table, select Rules.
   - To copy the indexes associated with the table, select Indexes.
4. If you did not select All Columns, in the Table Columns box, select the columns to add.
5. Choose the Add button or double-click the column to add.
6. Repeat steps 4 and 5 for each column you want to copy to the new table.
   If you add a column by mistake, double-click that column to remove it, or select the column and choose the Remove button.
7. To set the criteria the data must meet to be copied to the new table, type the appropriate Transact-SQL WHERE clause in the Data Selection Criteria box.
8. Choose the Copy button.

To move a table and/or its indexes to a segment

1. In the Table Name box, select the table to move to a different segment or select the table whose indexes you want to move to a different segment.
   To view information about the table, choose the Table Info button.
2. In the Segment box, select the segment to place the table or its indexes on.
3. Under Table Indexes, select the items to place on the segment.
   - To place the table on the segment, select Table and Clustered Index.
   - To place the nonclustered indexes on the segment, select All Nonclustered Indexes.
   - To choose the nonclustered index to place on the segment, choose Select Index, and then select the index.
4. Choose the OK button.

To select data from a table

You can construct basic SELECT queries. You can specify the WHERE, GROUP BY, HAVING, ORDER BY, and COMPUTE clauses of the SELECT statement. You can also save the information as a view.

1. In the Tables box, select the table to query.
2. In the Table Columns box, select the columns to query.
3. In the Expression Type box, select the function to apply to the column, if applicable. To add the columns without an expression, select Column.
4. In the Operator Type box, select the search condition to use with the query, if applicable. Or select (none).
5. Choose the Add To Clause button.
   The information is added to the SELECT statement.
6. Add the appropriate conditions to the query.
   - Under Query, type the conditions in the WHERE, GROUP BY, HAVING, ORDER BY, and COMPUTE boxes. Use Transact-SQL syntax. Or select the conditions from the Expression Type and Operator Type boxes.
   - To include only unique rows in the results, select the Distinct check box.
   - To remove information added to the query by mistake, choose the Undo button.
7. Save or execute the query.
   - To save the query as a view, choose the Save View button. The Save Select As View dialog box appears. Type a name for the view. Choose the OK button.
   - To execute the query, choose the Execute button. The Select Results window appears and displays the results of the query.

To create a trigger

1. In the Table Name box, select the table to create a trigger for.
   To view information about the table, choose the Table Info button.
2. In the Trigger Name box, type a trigger name.
   Trigger names must conform to the rules for identifiers and must be unique in the database.
3. Select one or more Transact-SQL statements to associate the trigger with. The name of any existing trigger of that type is displayed to the right of that type.
   - To associate the trigger with INSERT statements, select For Insert.
   - To associate a trigger with UPDATE statements, select For Update.
   - To associate a trigger with DELETE statements, select For Delete.

   For Update and Delete triggers, you can select either restrictive or cascading referential integrity. If you select the Cascade Option button, changes to the primary key tables are also made to the foreign key tables. If you select the Restrict option button, changes to the primary key tables are disallowed if a foreign key exists with the same value as the primary key. For an Update trigger, the default is Cascade. For a Delete trigger, the default is Restrict.
4. If you are creating an Insert or Update trigger and want to check the update status of specific table columns, in the Column Name box, select the columns to check, and then select the And or Or button to state how the column will be related to other columns being checked for update status.
5. Choose the OK button.
   The Transact-SQL script for the trigger appears in the Manage Triggers window. If no primary or foreign keys exist between the current table and any other table or view, a template of the CREATE TRIGGER statement appears.
6. Modify the information as necessary.
7. Save the trigger.
   - To create the trigger on the server, choose Execute.
   - To save the trigger to a data file, from the File menu, choose Save As. If you save the trigger to file, it is not created on the server. You must execute the trigger in order for it to take effect.

To create an index

1. In the Tables box, select the table to create an index for.
   To view information about the table, choose the Table Info button.
2. Choose the New button.
   The Specify New Index Name dialog box appears.
3. Type a name for the index.
   Index names must be unique within a table, but do not need to be unique within a database.
4. Choose the OK button.
   The Manage Indexes dialog box reappears.
5. Select the Index Type.
   To create a unique index, select the Unique cell.
   To create a clustered index, select the Clustered cell. If you do not explicitly select Clustered, a nonclustered index is created.
6. To place the index on a separate segment, select the Segment cell, and then, in the Properties box, select a segment to place the index on.
   A nonclustered index can be created on a different segment from the table. If you have created a table on a segment and do not specify that segment when you create a clustered index, the clustered index is created on the default segment and the table moves to the default segment as well.
   Before creating an index on a segment, see your System Administrator or the database owner for a list of segments that you can use. Certain segments can be allocated for specific tables or indexes for performance reasons or for other considerations.
7. Select the cell of the column to create the index for. You must specify at least one column.
   You can specify two or more columns to create a composite index on the combined values in the specified columns.
8. If you are creating a nonunique clustered index, set row options.
   To ignore duplicates in batches of data, select the Ignore Duplicate Row cell. If you set this option, when a duplicate row is entered in the table with the INSERT statement, that INSERT statement is canceled.
   To allow duplicate data in the table, select the Allow Duplicate Row cell. When this option is set, the table is allowed to have duplicate rows and a user can insert duplicate rows as well.
9. Choose the Create button.
   To disallow the addition of rows that create duplicate keys, select the Ignore Duplicate Key check box.

To modify an index

1. In the Tables box, select the table to modify the index for.
2. Edit the index information.
3. Choose the Modify button.

To rebuild an index

1. In the Tables box, select the table to rebuild the index for.
2. Select the name of the index to rebuild.
3. Choose the Rebuild button.

To rename an index

1. In the Tables box, select the table that contains the index to rename.
2. Select the name of the index to rename.
3. Choose the Rename button.
   The Rename Index dialog box appears.
4. In the New Name box, type the new index name.
5. Choose the OK button.

To drop an index

1. In the Tables box, select the table that contains the index to drop.

2. Select the name of the index to drop.
3. Choose the Drop button.

To update index statistics

1. In the Table Name box, select the table to update statistics on.
   To view information about the table, choose the Table Info button.
2. In the Index Name box, select the index to update statistics on.
   To select more than one index at a time, hold down the CTRL or SHIFT key as you select index names.
3. Choose the OK button.

To create a primary key

1. Under Key Type, select Primary.
2. In the Working Table/View box, select the table to associate the primary key with.
   To view information about the table, choose the Table Info button.
3. In the Columns box, select the first column cell and then select the column to associate the key with.
4. Choose the Create button.
5. Choose the Close button.

To create a foreign key

1. Under Key Type, select Foreign.
2. In the Working Table/View box, select the table to associate the key with. (The primary key is listed in the working list and is dimmed.)
   To view information about the table, choose the Table Info button.
3. In the Related Table/View box, select the columns to associate the foreign key with.
   These columns match the primary key columns in the Working Table Columns box. The datatypes of the columns must match.
4. Choose the Create button.
5. Choose the Close button.

To create a common key

1. Under Key Type, select Common.
2. In the Working Tables box, select the table to associate the common key with.
   To view information about the table, choose the Table Info button.
3. In the Working Table Columns box, select the column for the common key.
4. In the Related Table box, select the common table.
5. In The Related Table Columns box, select the column to associate the common key with.
   These columns match those in the Working Table Columns box.
6. Choose the Create button.
7. Choose the Close button.

To modify a key

1. Select the Working Table key type and its related table (if it is a foreign or common key).
2. Modify the appropriate information.
3. Choose the Modify button.
4. Choose the Close button.

To drop a key

1. Select the working table.
2. Under Key Type, select the type of key to drop.
   The columns the key is associated with are listed in either the Working Table Columns or the Related Table Columns box. Select the related table if the key type is foreign or common.
3. Choose the Drop button.
4. Choose the Close button.

To create a rule

1. In the Rule Name box, select <New Rule>.
2. In the Rule Contents box, type the contents of the rule using Transact-SQL syntax. To enter a line break in the Rule Contents box, press CTRL+ENTER.
   To use an existing rule as the base for a new rule, select the rule in the Rule Name box and then alter the contents of the rule as necessary. Be sure to change the name of the rule in the CREATE RULE statement.
3. Choose the Create button.
4. Choose the Close button.

To bind a rule

1. In the Rule Name box, select the rule to bind.
2. Bind the rule.
   - To bind the rule to a table column, under Column Bindings, select the table and column to bind the rule to. Then choose the Bind button.
   - To bind the rule to a user-defined datatype, under User Defined Datatype Bindings, select the datatype to bind the rule to. Then choose the Bind button.

   To bind the rule only on future columns that the user-defined datatype is associated with, select the Futureonly check box.
3. Choose the Close button.

To unbind a rule

1. In the Rule Name box, select the rule to unbind.
2. Under Column Bindings or User Defined Datatype Bindings, select either the table and column or the user-defined datatype to unbind the rule from.
3. Choose the Unbind button.
4. Choose the Close button.

To get information about rule bindings

- Choose the Info button.

To drop a rule

1. In the Rule Name box, select the rule to drop.
2. Choose the Drop button.
3. Choose the Close button.

To create a default

1. In the Default Name box, select <New Default>.
2. In the Default Contents box, type the contents of the default using Transact-SQL syntax. To enter a line break in the Default Contents box, press CTRL+ENTER.
   To use an existing default as the base for a new default, select the default in the Default Name box and then alter the contents of the default as necessary. Be sure to change the name of the default in the CREATE DEFAULT statement.
3. Choose the Create button.
4. Choose the Close button.

To bind a default

1. In the Default Name box, select the default to bind.
2. Bind the default.
   To bind the default to a table column, under Column Bindings, select the table and column to bind the default to. Choose the Bind button.
   To bind the default to a user-defined datatype, under User Defined Datatype Bindings, select the datatype to bind the default to. Choose the Bind button.
   To bind the default only on future columns that the user-defined datatype is associated with, select the Futureonly check box.
3. Choose the Bind button.
4. Choose the Close button.

To unbind a default

1. In the Default Name box, select the default to unbind.
2. Under Column Bindings or User Defined Datatype Bindings, select either the table and column or the user-defined datatype to unbind the default from.
3. Choose the Unbind button.
4. Choose the Close button.

To get information about default bindings

Choose the Info button.

To drop a default

1. In the Default Name box, select the default to drop.
2. Choose the Drop button.
3. Choose the Close button.

To create a user-defined datatype

1. Choose the New button.
   The Specify New Type Name dialog box appears.
2. In the New Type Name box, type a name for the datatype.
   The datatype name must follow the rules for identifiers and must be unique in each database.
3. Choose the OK button.
   The Manage Datatypes dialog box reappears.
4. Select the Type cell of the new datatype row.
5. In the Properties box, select a system-defined datatype to base the user-defined datatype on.
6. If you selected a *char*, *varchar*, *binary*, or *varbinary* datatype, select the length cell of the new datatype row. Then, in the Properties box, type the maximum number of characters that can be stored in the column.
7. To specify if the datatype will allow null values, select the Nulls cell.
   If null values are not allowed, SQL Server will require an explicit entry for the column. You can override the datatype assignment of NULL/NOT NULL when you create a table if the base type allows it.
8. Choose the Create button.
9. Choose the Close button.

To bind a rule to a user-defined datatype

1. Select the Rule cell. Then, in the Properties box, select the rule.
2. Choose the Modify button.

To bind a default to a user-defined datatype

1. Select the Default cell. Then, in the Properties box, select the default.
2. Choose the Modify button.

To rename a user-defined datatype

1. Select the datatype row to rename by clicking the Name cell.
2. Choose the Rename button.
   The Specify New Type Name dialog box appears.
3. In the New Name box, type a new name for the datatype.
4. Choose the OK button.
5. Choose the Close button.

To drop a user-defined datatype

1. Select the datatype row to drop.
2. Choose the Drop button.
3. Choose the Close button.

To view user-defined datatype properties

1. Select the datatype row of the datatype to view properties for.
2. Choose the Info button.

To grant object permissions

1. In the Object box, select the object to grant permissions for.
   You can view information about the object by choosing the Properties button.
2. Under Permissions, select the type of permissions to grant. The None check box is selected by default.
   - To grant INSERT permission, select the Insert check box.
   - To grant DELETE permission, select the Delete check box.
   - To grant EXECUTE permission for stored procedures, select the Execute check box.
   - To grant SELECT permission on all columns, select the Select check box. Or, to grant SELECT permission on specific columns, under Column-level Privileges select the columns to grant SELECT permission for.
   - To grant UPDATE permission on all columns, select the Update check box. Or, to grant UPDATE permission on specific columns, under Column-level Privileges, select the columns to grant UPDATE permission for.
   - To grant all permissions, select the All check box.
3. In the Existing Users/Groups box, select the user or group to receive the selected permissions.
4. Choose the Add button or double-click the username.
5. Repeat steps 3 and 4 for each user or group.
6. Choose the Grant button.
7. Repeat steps 2 through 6 for each set of permissions to grant.
8. Choose the Close button.

To view a user or group's permissions on an object

1. In the Object box, select the object to view the user or group's permissions for.
2. In the Existing Users/Groups box, select the user or group.
3. Choose the User Zoom button.
   The permissions that user has are displayed.
4. Choose the Close button.

To a view a user or group's permissions on all objects

1. In the Existing Users/Groups box, select the user to view all permissions for.
2. Choose the User Info button.
   The User Permissions dialog box appears.
3. Choose the Close button.

To view object permissions

1. In the Object box, select the object to view permissions for.
2. Under Permissions, select the permissions to view.
3. Choose the Show Users button.
   Users who have the specified permissions on the chosen object are listed in the Users With Selected Permissions box.
4. Choose the Close button.

To revoke object permissions

1. In the Existing Users/Groups box, select the user or group to remove permissions from.
2. Choose the Add button.
3. In the Users With Selected Permissions box, select the user or group.
4. Remove the appropriate permissions.
   - To remove one of the user or group's permissions, under Permissions, select the permission's check box.
   - To remove the user or group's permissions on certain columns, select the columns from the Select Columns and Update Columns cells.
   - To remove all the user or group's permissions, select the All check box.
6. Choose the Revoke button.

7. Choose the Close button.

To set query options

1. Under Query Options:

Select	To
Abort on Arithmetic Error	End a query when an overflow or divide-by-zero error occurs during a query. You cannot set both Abort on Arithmetic Error and Ignore Arithmetic Error.
Ignore Arithmetic Error	Return NULL when an overflow or divide-by-zero error occurs during a query. You cannot set both Abort on Arithmetic Error and Ignore Arithmetic Error.
No Count Display	Turn off the message returned at the end of each query result that tells how many rows were affected by the query statement.
No Execute	Compile a query but not execute it.
Parse Query Only	Check the syntax of each query and return any error messages without executing the query.
Show Query Plan	Generate a text description of the processing plan for the query as the query is processed. (To generate a graphical description of the processing plan, select the Showplan button on the Query window toolbar. If both the Showplan check box and button are selected, both a text and graphic description of the processing plan are generated.)
Show Statistics Time	Display the time it took to parse and compile each query statement and the time it takes to execute each step of the query statement.
Show Statistics I/O	Display the number of scans, logical reads, and physical reads for each table referenced in a query statement. (To generate a graphical description of the processing plan, select the Statistics I/O button on the Query window toolbar. If both the Statistics I/O button and check box are selected, both a text and graphical description are generated.)

2. In the Row Count box, type the number of rows the server should return before ending a query. The default (0) returns all rows.
3. In the Result Output Format box, select the type of format to display the results in. The choices are Column Aligned, Comma Separated (CSV), Tab Delimited, and Other Delimiter. If you select Other Delimiter, in the Delimiter box type the delimiter. The default is Column Aligned.
4. Select the Verbose Prints check box to display the PRINT statements output in the results window. This option is useful if you are executing the graphical Showplan and/or Statistics I/O.

5. Clear the Print Headers check box so the column headers will not appear in the query results.
6. Choose the OK button.

To generate database scripts

1. Under Scripts, select the appropriate objects.
   - To generate a script for all objects in the database, select the All Objects check box.
   - To generate a script for all objects of a specific type, select the object type's check box. A script for all objects of that type is generated. You can select more than one object type.
   - To generate a script for a specific object, clear the All Objects check box. Then select an object in the Available Objects box and choose the Add button. (You can add multiple objects one at at time, or you can select multiple objects and then choose the Add button.) If you add an object by mistake, you can cancel the addition by selecting the object in the Selected Objects box and choosing the Remove button. To view information about the selected object, choose the Info button.
   - To include all objects that depend on the specific object(s) or object type(s) you selected, select the Include Dependencies check box.
   - To include DROP statements before the CREATE statements, select the Include Drops check box.
   - To include the permissions for the selected objects, select the Include Permissions check box.
   - To include the statements used to add database logins, select the All Database Logins check box.
   - To include the statements used to add users and groups to a database, select the All Users/Groups check box.
2. The Generated By box defaults to your login ID. You can type a different name, if applicable.
   To preview the database script that will be generated, choose the Preview button. The database script appears in the Script Preview window. To save the script, from the File menu, choose Save As. To return to the Generate SQL Scripts dialog box, choose the Scripts button.
3. When you are ready to save the script, choose the Generate button.
   The Save As dialog box appears.
4. Type a name for the file in which to save the Transact-SQL code.
5. Choose the OK button.
   The script is generated and stored in the selected file.

Current Bindings (Rule)

- The Datatype box lists the datatypes the rule is bound to.
- The Table Columns box lists the columns the rule is bound to.
- To unbind a rule from all its bindings, choose the Unbind All button.
- To print the rule bindings information, choose the Print button.
- To save the rule bindings information as a report, choose the Save As button.

Current Bindings (Default)

- The Datatype box lists the datatypes the default is bound to.
- The Table Columns box lists the columns the default is bound to.
- To unbind a default from all its bindings, choose the Unbind All button.
- To print the default bindings information, choose the Print button.
- To save the default bindings information as a report, choose the Save As button.

To set the user

1. In the User Name box, select the user whose identity in the database you want to assume. You must be the system administrator or the database owner to do this.
2. Choose the OK button.

To save as a view

1. In the New View Name box, type a name for the view.
2. Choose the OK button.

To create a new table

1. In the New Table Name box, type a name for the new table.
2. To create a column, select the Column Name cell and then, in the Properties box, type a name for the table column.
3. In the same row, select the Datatype cell and, then, in the Properties box, select the datatype.
   If you select a user-defined datatype and a rule or default is bound to it, the name of the rule or default appears in the Rule or Default cell.
4. If you chose a datatype of *char*, *varchar*, *binary*, or *varbinary*, select the Length cell in the same row and then type the maximum number of characters that can be stored in the column.
5. In the same row, select the Nulls cell to allow null values in the column.
6. If you want to bind a default or rule to the column, select the Default or Rule cell, and then select the default or rule in the Properties box.
7. Choose the Insert button.
8. Repeat steps 2 through 7 for each column of the table.
9. Choose the Create button.

To use the structure of an existing table to create a new table

1. Choose the Copy Table button.
   The Select Table to Copy dialog box appears.
2. In the Table Name box, select the table to copy.
3. Choose the OK button.
   The table structure appears in the Create New Table dialog box.
4. Edit and/or add new columns as necessary.
5. Choose the Create button.

To change column mapping

1. To specify a delimiter for the columns, under Delimiters, select the delimiter in the Column box.
2. To specify a delimiter for the last row of the column, under Delimiters, select the delimiter in the Row box.
3. To create a new delimiter, select the New button. In the Add New Delimiter dialog box, type a name and value for the delimiter, and then choose the OK button.
   The new delimiter is added to the Column and Row boxes.
4. To skip a column for data transfer, select the Skip cell of the column that you want to skip.
5. To change the datatype format of a column, select the Format cell of the column row and then, in the Properties box, select the datatype. If you pick an incorrect match, a message will appear.
   If necessary for that datatype, select the Length cell and type the length in the Properties box.
6. Choose the OK button.

To change column mapping

1. To specify a delimiter for the columns, under Delimiters, select the delimiter in the Column box.
2. To specify a delimiter for the last row of the column, under Delimiters, select the delimiter in the Row box.
3. To create a new delimiter, select the New button. In the Add New Delimiter dialog box, type a name and value for the delimiter, and then choose the OK button.
   The new delimiter is added to the Column and Row boxes.
4. To change the datatype format of a column, select the Format cell of the column row and then, in the Properties box, select the datatype. If you pick an incorrect match, a message will appear.
   If necessary for that datatype, select the Length cell and type the length in the Properties box.
5. Choose the OK button.

To rename an object

1. In the New Name box, type a new name for the object.
   When you change the name of any object, you must rebuild any views, stored procedures, or triggers that reference the object. Be aware than when you change the name of a stored procedure, trigger, view, rule, or default, the CREATE statement used to create that object is not updated. You must change the CREATE statement text for each object in its respective window if you plan to generate scripts based on the statement for those objects.
2. Choose the OK button.

Rule/Default Properties

The properties of the selected default or rule, including its owner, creation date, the contents of the default or rule, and its current bindings, are displayed.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

Stored Procedure Properties

The properties of the selected stored procedure, including its owner, creation date, objects it references, and parameters it includes, are displayed.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

User Datatype Properties

The properties of the selected datatype, including its owner, physical type, length, whether it allows null values, any rules or defaults bound to it, and the table columns it is associated with, are displayed.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

To view object dependencies

1. Under Subject Object, select the name of the object to view dependencies for.
   The objects that the object depends on and the objects that depend on that object are displayed.
2. To print the information, choose the Print button.
3. To save the information, choose the Save As button.
4. To view the properties of the selected object, choose the Properties button.
5. Choose the Close button.

View Properties

The properties of the selected view, including its owner, creation date, the tables and views that the view is based on, the columns of the view, and the keys of the view, are displayed.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

Trigger Properties

The properties of the selected trigger, including the trigger type, its creation date, the table the trigger is for, the table owner, and the referenced objects, are displayed.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

Table Properties

The properties of the selected table, including its creation date, its owner and size, the segment the table is on, the number of columns and rows, the row length, and the amount of reserved space, are displayed. In addition, information about the columns is displayed, as well as the indexes, keys, and triggers associated with the table.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

To read a BCP format from a file

1. In the File Name box, select the format file.
   If necessary, first choose the drive and directory.
2. Choose the OK button.

To write a BCP format to a file

1. In the File Name box, type a name for the format file.
   If necessary, specify the drive and directory in which to save the format file.
2. Choose the OK button.

To set data transfer options

1. In the Maximum Errors Allowed box, type the maximum number of errors the data transfer can generate before the transfer terminates.
2. In the Transfer Batch Size box, type the size of the batches in which the data should be transferred to the server.
3. Under Transfer Data, select All Rows to transfer all rows. Or, to transfer only a range, select Range and then, in the From box, type the first row to transfer, and in the To box, type the last row to transfer.

To specify a new index name

1. In the New Index Name box, type a name for the new index.
2. Choose the OK button.

To specify a new datatype name

1. In the New Type Name box, type a name for the new user-defined datatype.
2. Choose the OK button.

To add a new delimiter

1. In the Delimiter Name box, type a name for the delimiter.
2. In the Delimiter Value box, type the value for the delimiter.
3. Choose the OK button.

To select a table to copy for data transfer

1. In the Table Name box, select the name of the table to copy.
2. Choose the OK button.

To select a table for DBCC CheckTable

1. In the Table Name box, select the name of the table to check.
   DBCC CheckTable checks that index and data pages are correctly linked, that indexes are in properly sorted order, that all pointers are consistent, and that the data information on each page and page offsets are reasonable. If the log segment is on its own device, running DBCC CheckTable on the *syslogs* table reports the logs used and the amount of free space.
2. Choose the OK button.

To connect to a server

1. In the Active Servers box, select the server.
   The Active Servers box lists the network servers. You can refresh the list by choosing the Refresh button.
2. Choose the OK button.

To specify a new table name

1. In the New Table Name box, type a name for the table.
2. Choose the OK button.

To change servers

1. In the Servers box, select the name of the server.
2. Choose the OK button.

To change databases

1. In the Databases box, select the name of the database.
2. Choose the OK button.

To transfer data

1. Select either the Import into SQL Server or the Export from SQL Server check box.
2. Under Destination Table, select the table to import data into, or under Source Table, select the table to export data from.
   To define a new table to transfer data into, choose the Define New Table button.
3. Under File Type, select the type of file to transfer:

BCP character (*.CSV)	Comma delimited column separation
BCP character (*.TAB)	Tab delimited column separation
BCP defined (*.DAT)	Uses column mapping to define the columns
BCP format file (*.DAT)	Uses a BCP format file
BCP native (*.DAT)	Uses the data's native datatypes as defaults

4. Under Source File, select the file to import data from, or Under Destination File, select the file to export data to.
5. If you chose a file type of BCP defined, choose the Column Mapping button to define the columns.
6. To create or read a format file for the data transfer, choose the Format File button.
   If you chose BCP format file as the file type, you must choose the Format File button to specify which format file to read for the data transfer. If you chose any other file type, you can create a format file for the transfer by choosing the Format File button.
7. To set data transfer options, choose the Options button.
8. To start the data transfer, choose the Import or Export button.

This dialog box allows you to set the Select Into/Bulk Copy option so that data transfer into an existing table is faster. The Reset Off When Finished check box is selected by default. This turns Select Into/Bulk Copy off after the data transfer. Clear this check box to leave Select Into/Bulk Copy on after the data transfer is complete.

This dialog box allows you to suspend the indexes on a table that you are transferring data into so that data transfer is faster. If you choose Yes, the indexes are dropped and then rebuilt after the data transfer. Select the Write Indexes to File check box to generate the scripts used to create the indexes. This option ensures that if you have a power failure or if you try to rebuild a unique index and have a nonunique value in your data file, you will be able to fix the index.

Confirm Unique Delimiters

This dialog box warns that you have specified a delimiter for a non-character column (for example, *binary* or *varbinary*) and there is a chance that the delimiter may occur in the column's data values, which could cause errors with your file. To replace the delimiter for non-character columns with None, select the Replace with 'None' check box and choose the OK button. If you are sure that the delimiter you have chosen will not conflict with your data values, choose OK but be aware that errors could occur.

To specify options for the index

- To specify how full SQL Server should make each page when creating a new index on existing data, select the Fillfactor check box and type a FILLFACTOR for the index.

  Note that it is not advisable to change the FILLFACTOR because the FILLFACTOR percentage affects SQL Server performance. SQL Server must take the time to split pages when they fill up.

- If the data you are creating the index on is sorted, select the Sorted Data check box.

  This eliminates the CREATE INDEX statement sort. SQL Server verifies that the data has been sorted by checking each index value to see if it is higher than the previous one. If any row fails this check, the CREATE INDEX statement is canceled.

User Permissions

This dialog box allows you to view a user's permissions for a table, view, stored procedure, or all database objects.

- To print the information, choose the Print button.
- To save the information as a report, choose the Save As button.
- To close the dialog box, choose the Close button.

Login ID

The unique name a user uses when logging in to a SQL Server. With integrated security for Windows NT, you do not need to provide a login ID.

Password

A confidential word the user supplies to gain access to a SQL Server. Because a password is confidential, it does not appear on the screen when typed. With integrated security for Windows NT, you do not need to provide a password.

Default Database

The database a user can access when first logging in to a SQL Server.

Datatype

An identifier that specifies what type of information a column holds and how the data is stored. Microsoft SQL Server provides 17 system-supplied datatypes; user-defined datatypes can also be created.

Segment

A subset of a device, available to a particular database to store objects on. Before using a segment, see your system administrator or the database owner for a list of segments that you can use. Certain segments can be allocated to specific tables or indexes for performance or other reasons.

Table

A collection of rows that have associated columns.

Column

The logical equivalent of a field. A column contains an individual data item within a row or a record.

Rule

A database object that is bound to a column or to a user-defined datatype that controls what data can be entered in that column.

Default

A database object bound to a column or to a user-defined datatype that causes SQL Server to automatically supply a specified value if one is not explicitly provided.

Index

A database object that provides access to data in the rows of a table, based on key values. Indexes provide quick access to data and can enforce uniqueness on the rows in a table.

Clustered Index

An index in which the physical order and the logical (indexed) order are the same.

Nonclustered Index

An index that stores key values and pointers to data.

Permissions

The authorization that allows a user to perform certain actions on certain database objects or to use certain statements.

Key

A field used to identify a record; often used as the index field for a table.

Trigger

A special form of stored procedure that is executed when a user issues a change statement such as INSERT, DELETE, or UPDATE to a specific table or column.

Stored Procedure

A pre-compiled collection of Transact-SQL statements and optional control-of-flow statements stored under a name. SQL Server-supplied stored procedures are called system procedures.

Database

A set of related tables and other database objects that are organized to facilitate searching, sorting, and recombination of data.

View

An alternative way of looking at the data in one or more tables. A view is usually created as a subset of columns from one or more tables.

Primary Key

A column or columns whose values uniquely identify a row in a table.

Foreign Key

A column or columns whose values are required to match a primary key in some other table.

Common Key

A key created to make explicit a logical relationship between two tables in a database.

User-defined Datatype

A definition of the type of data a column can contain; it is created by the user. User-defined datatypes are defined in terms of existing system datatypes.

Object

One of the components of a database: a table, index, trigger, view, key, rule, default, user-defined datatype, or stored procedure. Also called a database object.

Base Tables

The permanent tables on which a view is based. Also called underlying tables.

Query

A request for data retrieval.

Server

A computer on a local area network that controls access to resources such as files, printers, and communication devices. When using a SQL Server, you can access only servers designated as SQL Server servers.

ANSI -> OEM Conversion

If SQL Server is not using the ANSI (ISO) character set, and you need to display extended characters, AutoANSItoOEM conversion activates character conversion. This corrects problems with extended characters displaying as graphics instead of letters with diacritical marks. You must disconnect from SQL Server and reconnect for the selection to take effect. Be aware that this option affects all Windows- and Windows NT-based applications using the DB-Library API, not just SQL Object Manager (unless the DB-Library application explicitly overrides this setting).

Identifiers

A word or name, from 1 to 30 characters. The first character must be a letter or an underscore (_), at sign (@), or pound sign (#). An identifier that begins with # denotes a temporary table. An identifier that begins with @ denotes a variable. Spaces are not allowed in identifiers.

Unique Index

An index that prohibits duplicate index or key values. The system checks for duplicate key values when the index is created and checks each time data is added with an INSERT or UPDATE statement.

Composite Index

Indexes used when two are more columns are best searched as a unit.

Sorted Data

Selecting the Sorted Data check box eliminates the CREATE INDEX statement sort. SQL Server verifies that the data has been sorted by checking each index value to see if it is higher than the previous one. If any row fails this check, the CREATE INDEX statement is canceled.

Update Statistics

Updates information about the distribution of key values in specified indexes. If there was no data in the table when the index on the table was created, update the index statistics after data is added.

Binding a Rule

Links the rule to a table column and/or to a user-defined datatype.

Binding a Default

Links the default to a table column and/or to a user-defined datatype.

Base Datatype

A SQL Server-supplied datatype from which a user-defined datatype will be built.

Row

A set of related columns that describe a specific entity. Also called a record.

Referential Integrity

Referential integrity involves managing corresponding data values between tables when the foreign key (FK) or a table contains the same values as the primary key of another table. For example, in the *pubs* database, the author_id column is the primary key of the Authors table. The author_id column in the AuthorTitle table is a foreign key to the Authors table. Values in the AuthorTitle.author_id column must match values in the Authors.author_id column to maintain referential integrity. Referential integrity is enforced with triggers during INSERT, UPDATE, and DELETE operations. With cascading referential integrity, changes to an existing primary key value are cascaded to existing foreign key values in other tables. With restrictive referential integrity, operations that violate referential integrity between the primary key and foreign key values are restricted.

Object Manager Keyboard Shortcuts

To get information about the keys for SQL Object Manager, choose from these topics:

SQL Object Manager Keys
Cursor Movement Keys
Dialog Box Keys
Menu Keys
Editing Keys
Help Key
System Keys
Text Selection Keys
Windows Keys

Editing Keys

Use the following keys to edit text:

Key(s)	Function
BACKSPACE	Deletes the character to the left of the insertion point, or deletes the selected text.
DEL	Deletes the character to the right of the insertion point, or deletes the selected text.
SHIFT+DEL	Deletes the selected text and places it on the Clipboard.
SHIFT+INS	Inserts text from the Clipboard to the active window.
CTRL+INS	Copies the selected text to the Clipboard.
ALT+BACKSPACE	Undoes the previous editing operation.

System Keys

The following keys can be used from any window, regardless of which application you are using:

Key(s)	Function
CTRL+ESC	Switches to the Task List.
Alt+ESC	Switches to the next application or minimized icon, including full-screen programs.
ALT+TAB	Cycles through running applications.
PRTSC	Copies an image of the entire screen contents to the Clipboard.
ALT+PRTSC	Copies an image of the SQL Object Manager window to the Clipboard.
ALT+F4	Closes the application.
CTRL+F4	Closes the active window.
F1	Displays Help information in SQL Object Manager dialog boxes.

Dialog Box Keys

Use the following keys within dialog boxes:

Key(s)	Function
TAB	Moves from option to option (left to right and top to bottom).

SHIFT+TAB	Moves from option to option in reverse order.
ALT+letter	Moves to the option or group whose underlined letter matches the one you press.
Arrow keys	Moves the selection cursor from option to option within a group of options. Or moves the cursor left, right, up, or down within a list or text box.
HOME	Moves to the first item or character in a list or text box.
END	Moves to the last item or character in a list or text box.
PAGE UP or PAGE DOWN	Scrolls up or down in a list box, one window at a time.
ALT+Up or Down arrow	Opens a drop-down list box and selects an item in a drop-down list box.
SPACEBAR	Selects or clears a check box.
SHIFT+Arrow key	Selects text in a text box.
SHIFT+HOME	Selects text from the cursor point to the first character in a text box.
SHIFT+END	Selects text from the cursor point to the last character in a text box.
ENTER	Executes a command button, or chooses the selected item in a list box and executes the command.
ESC or ALT+F4	Closes a dialog box without completing the command (same as the Cancel button).

Help Key

Use the following key to get Help information:

Key	Function
F1	Within a dialog box, pressing F1 displays the dialog box Help information. If the Help window is already open, pressing F1 displays information about how to use Help.

Cursor Movement Keys

Use the following keys to move the cursor (insertion point) in text boxes and other places where you can

type text:

Key(s)	Moves the insertion point
Up arrow	Up one line.
Down arrow	Down one line.
Right arrow	Right one character.
Left arrow	Left one character.
CTRL+Right arrow	Right one word.
CTRL+Left arrow	Left one word.
HOME	To the beginning of the line.
END	To the end of the line.
PAGE UP	Up one window.
PAGE DOWN	Down one window.
CTRL+HOME	To the beginning of the window or text area.
CTRL+END	To the end of the window or text area.

Text Selection Keys

Use the following keys to select text:

Key(s)	Function
SHIFT+Left or Right arrow	Selects text one character at a time to the left or right, or, if the character is already selected, cancels the selection.
SHIFT+Down or Up arrow	Selects one line of text up or down, or, if the line is already selected, cancels the selection.
SHIFT+PAGE UP	Selects text up one window, or, if the previous window is already selected, cancels the selection.
SHIFT+PAGE DOWN	Selects text down one window, or, if the next window is already selected, cancels the selection.
SHIFT+HOME	Selects text to the beginning of the line.
SHIFT+END	Selects text to the end of the line.
CTRL+SHIFT+Left arrow	Selects the previous word.
CTRL+SHIFT+Right arrow	Selects the next word.
CTRL+SHIFT+HOME	Selects text to the beginning of the document.
CTRL+SHIFT+END	Selects text to the end of the document.

Menu Keys

Use the following keys to select menus and to choose commands:

Key(s)	Function
ALT or F10	Selects the leftmost menu on the menu bar.
ALT + letter	Chooses the menu or menu item whose underlined letter matches the one you press.
Left or Right arrow	Moves among menus.
Up or Down arrow	Moves among menu items.
ENTER	Chooses the selected menu item.
ESC	Cancels the selected menu.

Windows Keys

Use the following keys to navigate in the Microsoft Windows or Windows NT operating systems:

Key(s)	Function
ALT+SPACEBAR	Opens the System menu for an application.
ALT+- (Hyphen)	Opens the System menu for a document window.
ALT+F4	Closes an application.
ALT+ESC	Switches to the next application or minimized icon, including full-screen programs.
ALT+TAB	Cycles through running applications.
ALT+ENTER	Switches an application for an operating system other than Windows between running in a window and running full screen.
Direction keys	Moves a window when you have chosen Move from the System menu, or changes the size of a window when you have chosen Size from the System menu.

SQL Object Manager Keys

The following keys are specific to SQL Object Manager:

Key(s)	Function

F8	Refreshes current window.
F3	Closes SQL Object Manager.
F4	Closes the current window.
CTRL+N	Displays the SQL Server Login dialog box.
CTRL+O	Displays the Database Objects window.
CTRL+T	Displays the Data Transfer window.
CTRL+S	Displays the Generate SQL Scripts dialog box.
CTRL+Q	Displays the Query window.
CTRL+U	Selects the Current Connections list box.
CTRL+D	Selects the Current Database list box.

The following keys are specific to the SQL Object Manager Query window:

Key(s)	**Function**
ALT+X or CTRL+E	Executes a query.
CTRL+Z	Clears the Query window.
CTRL+C	Cancels the query
F1	Displays TransactSQL Help information for the selected syntax.
ALT+F1	Displays Help (sp_help) information for the selected object.
F6	Moves between the Query and result windows.
ALT+S	Activates the resizing mode for the splitter bar. Use the Up and Down Arrow keys to resize the windows and then press ENTER to set the new size.

Using Microsoft SQL Object Manager

Microsoft SQL Object Manager is a tool that allows you to manage Microsoft SQL Server database objects through the Microsoft Windows and Windows NT operating systems. Database objects are the components of a database--its tables, indexes, triggers, views, keys, rules, defaults, datatypes, and stored procedures. SQL Object Manager allows you to create, alter, and delete database objects, as well as to grant and revoke object permissions. Using SQL Object Manager you can transfer data to and from a SQL Server, generate database scripts, and execute data retrieval and data modification statements.

The following procedures will get you started using Microsoft SQL Object Manager. Click the topics below to get more information. For more information about the tasks you can do with SQL Object Manager, from the Help menu choose Contents.

Creating a Table

Creating Indexes

Creating a Trigger

Creating a Rule

Creating a Default

Granting Object Permissions