



ADMINISTRATION GUIDE | PUBLIC

SAP Adaptive Server Enterprise 16.0 SP03

Document Version: 1.0 – 2019-06-06

Component Integration Services Users Guide

Content

- 1 Component Integration Services Overview. 5**
- 2 Component Integration Services Concepts. 7**
- 2.1 Proxy Tables. 9
 - Using the create table Command. 9
 - Using the create existing table Command. 10
 - Using the create proxy_table Command. 11
 - Remote Procedures as Proxy Tables. 14
 - Server Limits. 16
 - Cascading Proxy Tables. 18
- 2.2 Proxy Databases. 19
 - User Proxy Databases. 19
 - System Proxy Databases. 21
- 2.3 File System Access. 23
 - Directory Access. 23
 - Recursion Through Subordinate Directories. 26
 - File Access. 26
- 2.4 Remote Servers. 28
 - Connection Management. 29
 - Connect to Remote Servers Without the Interfaces File. 29
 - LDAP Directory Services. 30
 - Secure Communication with SSL. 31
 - Secure Communication Using Kerberos. 31
 - Security Issues 35
 - Remote Server Logins. 36
 - Map External Logins. 37
 - Configuring Failover for Remote Server Connections. 39
 - Remote Server Capabilities. 39
- 2.5 Query Processing. 40
- 2.6 RPC handling and Component Integration Services. 44
 - Site Handler and Outbound RPCs. 44
 - Configuring Component Integration Services for Outbound RPCs. 44
 - Text Parameters for RPCs. 46
 - Text Parameter Support for XJS/390. 47
- 2.7 Distributed Transaction Management. 47
 - DTM-Enabled Servers. 48
 - Pre-DTM servers. 48

	strict DTM enforcement.	49
	Enable Component Integration Services.	49
	Transactional RPCs.	49
	Restrictions on Transaction Management.	50
2.8	SAP ASE to SAP ASE update statistics.	50
2.9	Update Statistics on Non-SAP ASE Backends.	51
2.10	Java in the Database.	51
	Constraints on Java Class Columns.	52
	Java Abstract Datatypes (ADTs).	52
2.11	Datatypes.	53
	Unicode support.	53
	Datatype Conversions.	55
	text and image Datatypes.	55
2.12	Configuration and Tuning.	59
	Using sp_configure.	59
	Use Global Variables to Determine Status.	62
3	CIS Support for Connecting to an SAP HANA Server.	63
3.1	Creating SAP HANA as an ODBC Data Source on Windows.	63
3.2	Adding SAP HANA to the SAP ASE Interfaces File.	64
3.3	Configuring the PCI Bridge and PCA/ODBC.	64
3.4	Adding Server Class HANAODBC.	66
3.5	Datatype Mapping Between SAP ASE and HANA.	67
3.6	Restrictions.	69
4	SQL Reference.	72
4.1	dbcc Commands.	72
	dbcc Options.	72
	Trace Flags.	73
4.2	Functions.	75
	Aggregate Functions.	75
	Datatype Conversion Functions.	76
	Date Functions.	76
	Mathematical Functions.	77
	Security Functions.	78
	String Functions.	78
	System Functions.	79
	Text and Image Functions.	81
4.3	Transact-SQLCommands.	81
	alter table.	81
	case.	84
	connect to...disconnect.	85

	create existing table.	86
	create index.	93
	create table.	94
	delete.	96
	drop index.	97
	fetch.	98
	insert.	99
	readtext.	100
	select.	101
	truncate table.	102
	update.	103
	update statistics.	104
	writetext.	106
4.4	Passthrough Mode.	107
	The connect to Command.	107
	The sp_autoconnect System Procedure.	108
	The sp_passthru System Procedure.	109
	The sp_remotesql System Procedure.	109
4.5	Quoted Identifier Support.	110
4.6	Delimited Identifier Support.	111
4.7	auto identity Option.	111
4.8	Triggers.	111
5	Tutorial.	112
5.1	Adding a Remote Server.	112
	Creating Server Entries in System Tables.	113
	Adding an Alternate Login and Password.	113
	Verifying Connectivity.	114
5.2	Performing a Join Between Two Remote Tables.	114
6	Troubleshooting.	116
6.1	Problems Accessing Component Integration Services.	116
6.2	Problems Using Component Integration Services.	117
	Unable to Access Remote Server.	117
	Unable to Access Remote Objects.	118
	Problem Retrieving Data from Remote Objects	119
6.3	If You Need Help.	121

1 Component Integration Services Overview

Component Integration Services extends SAP Adaptive Server Enterprise (SAP ASE) capabilities, provides enhanced interoperability, and location transparency and functional compensation.

Location transparency means that Component Integration Services allows SAP ASE to present a uniform view of enterprise data to client applications. Enterprise-wide data from heterogeneous sources can be accessed as if it were local.

Functional compensation allows Component Integration Services to emulate all features of Transact-SQL, and interact with a data source only when actual data is needed. With this capability, the full range and power of Transact-SQL can be applied to any data source, whether or not the data source provides support for a particular feature of Transact-SQL. Examples of this capability are built-in functions and Java functions. Component Integration Services allows statements to use these functions even though the data on which these functions may operate is derived from external sources that cannot support the functions.

Component Integration Services, together with SQL Anywhere, SAP IQ and various DirectConnect interfaces, extends the reach of SAP ASE by enabling transparent access to database management systems anywhere in the enterprise. This transparent, extended reach of SAP ASE makes it easy for Enterprise Portal components to:

- Access data from anywhere, and present it as dynamic content to Web pages
- Execute transactions that span heterogeneous boundaries
- View an entire enterprise through a single view provided by the global metadata stored in the SAP ASE/Component Integration Services system catalogs

Component Integration Services allows users to access both SAP and non-SAP databases on different servers. These external data sources include host data files, tables, views, and RPCs (remote procedure calls) in database systems such as SAP ASE and Oracle.

Using Component Integration Services, you can:

- Access tables in remote servers as if the tables were local.
- Perform joins between tables in multiple remote, heterogeneous servers. For example, it is possible to join tables between an Oracle database management system (DBMS) and an SAP ASE, and between tables in multiple servers.
- Transfer the contents of one table into a new table on any supported remote server by means of a `select into` statement.
- Maintain referential integrity across heterogeneous data sources.
- Access native remote server capabilities using the Component Integration Services passthrough mode.

Component Integration Services can be used by anyone who needs to access multiple data sources or legacy data. It can also be used by anyone who needs to migrate data from one server to another.

A single server is often used to access data on multiple external servers. Component Integration Services manages the data regardless of the location of the external servers. Data management is transparent to the client application.

Component Integration Services, in combination with EnterpriseConnect™ and MainframeConnect™, provides transparent access to a wide variety of data sources, including:

- Oracle
- Informix
- Microsoft SQL Server
- SAP Adaptive Server Enterprise
- SAP SQL Anywhere
- SAP IQ
- Mainframe data, including:
 - ADABAS
 - IDMS
 - IMS
 - VSAM

To start Component Integration Services:

- Install DirectConnect server(s) or gateways for the external data sources you choose to access (for example, Oracle, Informix, Microsoft SQL Server).
- Configure the server to access remote objects.

Related Information

[Component Integration Services Concepts \[page 7\]](#)

2 Component Integration Services Concepts

Component Integration Services presents tables to a client application as if all the data in the tables were stored locally. Remote tables are mapped to local proxy tables which hold metadata. The ability to access remote (or external) tables as if they were local is a hallmark of Component Integration Services..

Remote tables are mapped to local proxy tables which hold metadata. Internally, when a query involving remote tables is executed, the storage location is determined, and the remote location is accessed so that data can be retrieved.

To achieve location transparency, tables must first be mapped to their corresponding external locations.

Access Methods

Access methods form the interface between the server and an external object. The access method used to retrieve remote data is determined by two attributes of the external object:

- The server class associated with the remote object. For each server class, there are separate access methods that handle all interaction between SAP ASE and remote servers of the same class and object type.
- The object type

Server Classes

A server class must be assigned to each server when it is added using `sp_addserver`. Server classes determine the access method used to interact with the remote server.

The server classes are:

- `<ASEnterprise>` – used if the server is SAP ASE. This is the default server class.
- `<ASAnywhere>` – used if the server is SAP SQL Anywhere version 6.0 or later. This server class should be used for SAP IQ versions earlier than SAP IQ 12.5.
- `<ASIQ>` – used if the server is SAP IQ version 12.5 and later.
- `<local>` – the local server. There can be only one.
- `<direct_connect >` – indicates that the server is an Open Server™ application that conforms to the interface requirements of a DirectConnect™ server. For access to Microsoft SQL Server, DB2, Oracle, or Informix, you must use a DirectConnect server.
- `<sds>` – indicates that the server conforms to the interface requirements of a Specialty Data Store.
- `RPCServer` – indicates that a server can only handle RPCs. It does not accept SQL statements, transaction control statements, or anything else.

Object Types

The server presents a number of object types to client applications as if they were local tables.

Supported object types are:

- `table` – the object in a remote server of any class is a relational table. This is the default type.
- `view` – the object in a remote server of any class is a view. Component Integration Services treats views as if they were local tables without any indexes.
- `remote procedure` – the object in a remote server of any class is a remote procedure. Component Integration Services treats the result set from the remote procedure as a read-only table.
- `file` – the object is an individual file within a file system.
- `directory` – the object is a file system directory.

Interface to Remote Servers

The interface between the server and remote servers is handled by the Open Client software, Client-Library™.

The Client-Library features that are used to implement the interface are dependent upon the class of server with which Component Integration Services is interacting.

For example, if the server class is `<direct_connect>`, a number of features such as cursor and dynamic requests are used.

Before the server can interact with a remote server, you must configure the following:

- Remote server addition to `directory services`
- Remote server definition
- Remote server login information
- Remote object definition

Before accessing remote tables with Component Integration Services, you must either have access to LDAP directory services, or an `interfaces` file (`sql.ini` file on Windows platforms).

Remote servers are defined by means of the stored procedure `sp_addserver`. This procedure is documented in *Reference Manual: Procedures*.

Once you have configured the remote server, you must provide login information. By default, Component Integration Services uses the names and passwords of SAP ASE clients whenever it connects to a remote server on behalf of those clients. However, this default can be overridden using `sp_addexternlogin`, which allows a system administrator to define the name and password for each user who connects to a remote server.

Using `connect to server_name`, you can verify that the server configuration is correct. This command establishes a passthrough mode connection to the remote server. Passthrough mode allows clients to communicate with remote servers in native syntax. This passthrough mode remains in effect until you issue a `disconnect` command.

Once you have configured a remote server, you cannot access objects in that remote server as tables until a mapping between them and a local object (proxy table) has been established.

You can create new tables on remote servers, and you can define the schema for an existing object in a remote server. The procedures for both are similar.

Related Information

[Connection Management \[page 29\]](#)

[Tutorial \[page 112\]](#)

2.1 Proxy Tables

A proxy table is a local table containing metadata which points to a remote object, and are the key to location transparency.

The remote table is mapped to the proxy table to make it appear as if it were a local table.

Related Information

[SQL Reference \[page 72\]](#)

2.1.1 Using the create table Command

The `create table` command creates a proxy table and a remote table at the same time.

Procedure

- Use this syntax:

```
create table <table_name> (<column_list>) [ [ external {table | file}] at  
"pathname" ]]
```

The remote location is specified with the `at pathname` clause. `create table` allows external object type `table` and `file`. The datatype of each column is passed to the remote server without conversion.

2.1.2 Using the create existing table Command

The `create existing table` command allows the definition of existing tables (proxy tables).

Context

The syntax for this option is similar to the `create table` command:

```
create existing table <table_name> (<column_list>)  
[[external {table | procedure | file}] at pathname]
```

The action taken by the server when it receives this command is quite different from the action it takes when it receives the `create table` command, however. A new table is not created at the remote location; instead, the table mapping is checked, and the existence of the underlying object is verified. If the object does not exist (either host data file or remote server object), the command is rejected with an error message.

If the object does exist, its attributes are obtained and used to update system tables `sysobjects`, `syscolumns`, and `sysindexes`.

- The nature of the existing object is determined.
- For remote server objects (other than RPCs), column attributes found for the table or view are compared with those defined in the `<column_list>`. Column names must match (case sensitive), column types and lengths must match, or at least be convertible, and the NULL attributes of the columns must match.
- Index information from the host data file or remote server table is extracted and used to create rows for the system table `sysindexes`. This defines indexes and keys in server terms and enables the query optimizer to consider any indexes that may exist on this table.

After successfully installing, and re-creating any proxy definitions, issue an `update statistics` command for the proxy tables. This allows the query optimizer to make intelligent choices regarding index selection and join order when proxy tables are involved.

Procedure

1. Define a server named `big_server`. Its server class is `<ASEnterprise>`, and its name in the interfaces file is `big_server`:

```
exec sp_addserver big_server, ASEnterprise, big_server
```

2. Define a remote login alias. This step is optional if the `<username>` and `<password>` are the same on both servers. User "sa" is known to remote server `big_server` as user "sa," password "timothy":

```
exec sp_addexternlogin big_server, sa, sa, timothy
```

3. Define the remote `authors` table:

```
create existing table authors  
(  
  au_id          varchar(11)      not null,
```

```

au_lname    varchar(40)    not null,
au_fname    varchar(20)    not null,
phone       char(12)      not null,
address     varchar(40)    null,
city        varchar(20)    null,
state       char(2)       null,
country     varchar(12)    null,
postalcode  char(10)      null
)
EXTERNAL TABLE at "big_server.pubs2.dbo.authors"

```

4. Update statistics on tables to ensure reasonable choices by the query optimizer:

```
update statistics authors
```

5. Execute a query to test the configuration:

```
select * from authors where au_lname = 'Carson'
```

2.1.2.1 Datatype Conversions

When you use the `create table` or `create existing table` commands, you must specify all datatypes, using recognized SAP ASE datatypes.

If the remote server tables reside on a class of server that is heterogeneous, the datatypes of the remote table are converted into the specified SAP ASE types automatically when the data is retrieved. If the conversion cannot be made, the `create table` or `create existing table` commands do not allow the table to be created or defined.

2.1.3 Using the `create proxy_table` Command

Using the `create proxy_table` command does not require a column list: Component Integration Services derives the column list from the metadata it obtains from the remote table.

Context

If the object does exist, `create proxy_table` updates `sysobjects`, `syscolumns`, and `sysindexes`.

You can add an optional clause to the `create existing table` statement to indicate the remote object is actually a stored (or other) procedure instead of a table. Without this clause, the remote object is assumed to be a table or view:

```

create existing table t1
(
    column_1    int,
    column_2    int
)
EXTERNAL PROCEDURE AT "big_server.mydb.dbo.p1"

```

If the remote object is type procedure, several processing differences occur:

- No indexes are created for objects of this type.
- You must provide a column list that matches the description of the remote procedure's result set. No verification of the list's accuracy is provided.
- You can use column names beginning with underscore ('_') to specify columns that are not part of the remote procedure's result set. These columns are referred to as parameter columns. For example:

```
create existing table t1
(
    a          int,
    b          int,
    c          int,
    _p1       int null,
    _p2       int null
)
external procedure
at "big_server.sybsystemprocs.dbo.myproc"
select a, b, c from t1
where _p1 = 10 and _p2 = 20
```

- In this example, the parameter columns `<_p1>` and `<_p2>` are not expected in the result set, but can be referenced in the query. Component Integration Services passes the search arguments to the remote procedure via parameters, using the names `<@p1>` and `<@p2>`.
- If a parameter column is included in the `select` list, its value is equivalent to the values specified for it in the `where` clause, if it was passed to the remote procedure as a parameter. If the parameter column did not appear in the `where` clause, or was not able to be passed to the remote procedure as a parameter, but was included in the `select` list, its value would be NULL.
- A parameter column can be passed to the remote procedure as a parameter if it is what the SAP ASE query processor considers to be a searchable argument. It is generally a searchable argument if it is not included in any "or" predicates. For example, the following query would prevent the parameter columns from being used as parameters.

```
select a, b, c from t1
where _p1 = 10 OR _p2 = 20
```

- Rules exist for the definition of parameter columns in the `create existing table` statement:
 - Parameter columns must allow NULL.
 - Parameter columns cannot precede normal, result columns (they must appear at the end of the column list).

Allowing the definition of remote procedures as local tables allows Component Integration Services to treat the result set of a remote procedure as a "virtual table," which can be sorted, joined with other tables, or inserted into another table via `insert/select` syntax. However, virtual tables are considered read-only:

- You cannot issue a `delete`, `update`, or `insert` command against a table of type `<procedure>`;
- You cannot issue a `create index`, `truncate table`, or `alter table` command against virtual tables.

If an object of the type `procedure` has been defined within the server, a query is not issued to the remote server on which the object resides. Instead, Component Integration Services issues an RPC and treats the results from the RPC as a read-only table.

When the following query is issued, Component Integration Services sends the RPC named `<myproc>` to server `big_server`. Row results are treated like the results from any other table; they can be sorted, joined with other tables, grouped, inserted into another table, and so forth:

```
create existing table rtable
(
  col1      int,
  col2      datetime,
  col3      varchar(30)
)
external procedure at "big_server...myproc "
select * from rtable
```

RPC parameters should represent arguments that restrict the result set. If the RPC is issued without parameters, the entire result set of the object is returned. If the RPC is issued with parameters, each parameter further limits the result set. For example, the following results in a single parameter, named `@col1`, that is sent along with the RPC. Its value is 10:

```
select * from rtable where col1 = 10
```

Component Integration Services attempts to pass as many of the search arguments as possible to the remote server, but depending on the SQL statement being executed, Component Integration Services might perform the result set calculation itself. Each parameter represents a search for an exact match, for example, the `=` operator.

The following rules define the parameters sent to the RPC. If an RPC is used as a Component Integration Services object, you should keep these rules in mind during development.

- Component Integration Services sends `=` operators in the `where` clause as parameters. For example, this query results in Component Integration Services sending two parameters:

```
select * from rpc1 where a = 3 and b = 2
```

Parameter `a` has a value of 3 and parameter `b` has a value of 2. The RPC is expected to return only result rows in which column `a` has a value of 3 and column `b` has a value of 2.

- Component Integration Services does not send any parameters for a `where` clause, or portion of a `where` clause, if there is not an exact search condition. For example:

```
select * from rpc1 where a = 3 or b = 2
```

Component Integration Services does not send parameters for `a` or `b` because of the `or` clause.

Another example:

```
select * from rpc1 where a = 2 and b < 3
```

Component Integration Services sends parameters for `<a>` and ``, and filters rows containing `` with values smaller than 3.

2.1.4 Remote Procedures as Proxy Tables

You can add an optional clause to the `create existing table` statement to indicate the remote object is actually a stored (or other) procedure instead of a table.

Without this clause, the remote object is assumed to be a table or view:

```
create existing table t1
(
    column_1    int,
    column_2    int
)
EXTERNAL PROCEDURE AT "big_server.mydb.dbo.p1"
```

If the remote object is type procedure, several processing differences occur:

- No indexes are created for objects of this type.
- You must provide a column list that matches the description of the remote procedure's result set. No verification of the list's accuracy is provided.
- You can use column names beginning with underscore ('_') to specify columns that are not part of the remote procedure's result set. These columns are referred to as parameter columns. For example:

```
create existing table t1
(
    a          int,
    b          int,
    c          int,
    _p1       int null,
    _p2       int null
)
external procedure
at "big_server.sybsystemprocs.dbo.myproc"
select a, b, c from t1
where _p1 = 10 and _p2 = 20
```

- In this example, the parameter columns `<_p1>` and `<_p2>` are not expected in the result set, but can be referenced in the query. Component Integration Services passes the search arguments to the remote procedure via parameters, using the names `<@p1>` and `<@p2>`.
- If a parameter column is included in the `select` list, its value is equivalent to the values specified for it in the `where` clause, if it was passed to the remote procedure as a parameter. If the parameter column did not appear in the `where` clause, or was not able to be passed to the remote procedure as a parameter, but was included in the `select` list, its value would be NULL.
- A parameter column can be passed to the remote procedure as a parameter if it is what the SAP ASE query processor considers to be a searchable argument. It is generally a searchable argument if it is not included in any "or" predicates. For example, the following query would prevent the parameter columns from being used as parameters.

```
select a, b, c from t1
where _p1 = 10 OR _p2 = 20
```

- Rules exist for the definition of parameter columns in the `create existing table` statement:
 - Parameter columns must allow NULL.
 - Parameter columns cannot precede normal, result columns (they must appear at the end of the column list).

Allowing the definition of remote procedures as local tables allows Component Integration Services to treat the result set of a remote procedure as a "virtual table," which can be sorted, joined with other tables, or inserted into another table via `insert/select` syntax. However, virtual tables are considered read-only:

- You cannot issue a `delete`, `update`, or `insert` command against a table of type `<procedure>`;
- You cannot issue a `create index`, `truncate table`, or `alter table` command against virtual tables.

If an object of the type `procedure` has been defined within the server, a query is not issued to the remote server on which the object resides. Instead, Component Integration Services issues an RPC and treats the results from the RPC as a read-only table.

When the following query is issued, Component Integration Services sends the RPC named `<myproc>` to server `big_server`. Row results are treated like the results from any other table; they can be sorted, joined with other tables, grouped, inserted into another table, and so forth:

```
create existing table rtable
(
    col1          int,
    col2          datetime,
    col3          varchar(30)
)
external procedure at "big_server...myproc "
select * from rtable
```

RPC parameters should represent arguments that restrict the result set. If the RPC is issued without parameters, the entire result set of the object is returned. If the RPC is issued with parameters, each parameter further limits the result set. For example, the following results in a single parameter, named `@col1`, that is sent along with the RPC. Its value is 10:

```
select * from rtable where col1 = 10
```

Component Integration Services attempts to pass as many of the search arguments as possible to the remote server, but depending on the SQL statement being executed, Component Integration Services might perform the result set calculation itself. Each parameter represents a search for an exact match, for example, the `=` operator.

The following rules define the parameters sent to the RPC. If an RPC is used as a Component Integration Services object, you should keep these rules in mind during development.

- Component Integration Services sends `=` operators in the `where` clause as parameters. For example, this query results in Component Integration Services sending two parameters:

```
select * from rpc1 where a = 3 and b = 2
```

Parameter `a` has a value of 3 and parameter `b` has a value of 2. The RPC is expected to return only result rows in which column `a` has a value of 3 and column `b` has a value of 2.

- Component Integration Services does not send any parameters for a `where` clause, or portion of a `where` clause, if there is not an exact search condition. For example:

```
select * from rpc1 where a = 3 or b = 2
```

Component Integration Services does not send parameters for `a` or `b` because of the `or` clause. Another example:

```
select * from rpc1 where a = 2 and b < 3
```

Component Integration Services sends parameters for <a> and , and filters rows containing with values smaller than 3.

2.1.5 Server Limits

SAP ASE configuration allows page sizes of 2K, 4K, 8K, or 16K bytes.

SAP ASE supports extended sizes of `char`, `varchar`, `univarchar`, `unichar`, `binary`, and `varbinary` datatypes. The limit depends on the page size of the server:

Table 1: Page and Column Limits

Page Size	Maximum Column Size
2048	2048
4096	4096
8192	8192
16384	16384

These sizes are approximate. The basic rule specifies that the limit is the maximum size that still allows a single row to fit on a page. These limits also vary depending on the locking scheme specified when the table is created. It is assumed that the bulk of proxy tables are created with the default locking scheme, which is `allpages` locking.

- Limits on length of Transact-SQL variables and parameters – the size of `char`, `varchar`, `binary`, and `varbinary` variables are extended to equal the maximum size of columns of the same datatype for a given server. This allows variables to be passed to stored procedures (or RPCs) whose length exceeds the current limit of 255 bytes.
- Limits on number of columns per table – as many as 1024 columns per table are allowed, as long as the columns can still fit on a page. There is a limit of 254 variable-length columns (null columns are also considered variable length).
- Limits on the width of an index – the total width of an index within SAP ASE can be larger than in earlier versions, depending on server page size.

Earlier versions of SAP ASE had a limit of 255 bytes for `char/binary` columns.

Table 2: Maximum index width

Page Size	Index Width
2048	600
4096	1250
8192	2600
16384	5300

- Limits on the number of columns per index – 31 columns per index.
- Table names, column names and index names can be up to 255 bytes.
- Identifier names can now be up to 255 bytes.

create new proxy table Limits

the `create table` command allows columns of datatype `char`, `varchar`, `binary`, and `varbinary` to be specified with extended lengths, as described above. These datatypes and lengths are forwarded to the remote server on which the table is to be created.

create existing proxy table Limits

the `create existing table` command also allows columns to be specified with a length of greater than 255 bytes. This allows Component Integration Services to treat columns in remote databases as `char`, `varchar`, `binary`, or `varbinary` that previously had to be treated as `text` or `image` columns.

There is still an opportunity for column size mismatch errors. For example, in the case where the remote database contains a table with a column length of 5000 bytes, and the SAP ASE processing `create existing table` supports columns only up to 1900 bytes, a size mismatch error occurs. In this case, it is necessary to respecify the column as a `text` or `image` column.

When the proxy table column size exceeds that of the corresponding column in the remote table, a size mismatch error is detected and the command is aborted.

alter table Limits

If `alter table` operates on a proxy table, it is first processed locally, then forwarded to the remote server for execution. If the remote execution fails, the local changes are backed out and the command is aborted.

The remote server must process the command appropriately, or raise an error. If an error is produced, the Component Integration Services side of the command is aborted and rolled back.

select, insert, delete, update Limits

Component Integration Services handles large column values when proxy tables are involved in data manipulation language (DML) operations. Component Integration Services handles DML using one of several strategies:

- Tabular data stream (TDS)[™] language commands – if the entire SQL statement can be forwarded to a remote server, then Component Integration Services does so using TDS Language commands generated by CT-Library `ct_command` (`CS_LANG_CMD`).

The text of the language buffer may contain data for `long char` or `binary` values that exceeds 255 bytes, and remote servers must handle parsing of these command buffers.

- TDS dynamic commands – if Component Integration Services cannot forward the entire SQL statement to a remote server (for example, Component Integration Services is forced to provide functional compensation for the statement), then an `insert`, `update`, or `delete` may be handled by using TDS dynamic commands, with parameters as needed, using the CT-Library function `ct_dynamic` (`CS_PREPARE_CMD`, `CS_EXECUTE_CMD`, `CS_DEALLOC_CMD`).
The parameters for the dynamic command may be `CS_LONGCHAR_TYPE` or `CS_LONGBINARY_TYPE`.
- TDS cursor commands – CT-Library cursor operations can be used to handle proxy table operations for `select`, `update`, and `delete` if functional compensation has to be performed. For example, if you are updating a proxy table and there are multiple tables in the `from` clause, Component Integration Services may have to fetch rows from multiple data sources, and for each qualifying row, apply the `update` to the target table. In this case, Component Integration Services uses `ct_cursor` (`CS_DECLARE_CMD`, `CS_OPEN_CMD`, `CS_CURSOR_UPDATE_CMD`, `CS_CLOSE_CMD`, `CS_DEALLOC_CMD`).
After a cursor is prepared, parameters are specified. These parameters may include those of type `CS_LONGCHAR` or `CS_LONGBINARY`.
- Bulk insert commands – when performing a `select/into` operation, if the target server supports the bulk interface (only true of remote SAP ASE servers and DirectConnect for Oracle), then the remote server must be prepared to handle `char` and `binary` values greater than 255 (via `CS_LONGCHAR`, `CS_LONGBINARY` values).
Columns from remote servers may be returned to Component Integration Services as type `CS_LONGCHAR_TYPE` or `CS_LONGBINARY_TYPE`.

RPC Handling

RPCs sent to remote servers can contain parameters of types `CS_LONGCHAR` and `CS_LONGBINARY`. The Component Integration Services command `cis_rpc_handling` supports these types.

Sending long parameters to SAP ASE servers older than version 12.5 is not allowed, as earlier versions of SAP ASE do not support `CS_LONGCHAR` or `CS_LONGBINARY` data. Component Integration Services examines TDS capabilities for the remote server prior to sending the RPC, and if the remote server cannot accept these datatypes, an error results.

sp_tables

The SAP SQL Anywhere or ASIQ stored procedure `sp_tables` only returns user tables.

2.1.6 Cascading Proxy Tables

SAP ASE allows cascading proxy table configurations between any number of instances of Component Integration Services.

There are conditions where this can cause problems, such as circular references, or transactions in which the second proxy table references a local table on the same server as the first proxy table. In this case, application

deadlocks can result that are not detected by Component Integration Services. You must configure your systems to avoid these potential pitfalls.

2.2 Proxy Databases

There are two types of proxy databases: user and system.

2.2.1 User Proxy Databases

When a user proxy database is created, metadata for the proxy tables is imported automatically from the remote location that contains the actual tables. This metadata is then used to create proxy tables within the proxy database.

To create a proxy database, use:

```
create database <<dbname>>
  [create database options]
  [with default_location = 'pathname']
  [for proxy_update]
```

The use of the clause `with default_location` allows you to specify the storage location of any new tables, and the location from which metadata may be imported for automatic proxy table creation if the `for proxy_update` clause is also specified. `for proxy_update` establishes the database as a proxy database; `with default_location` defines the location from which proxy tables are imported. Without `for proxy_update`, the behavior of `with default_location` is the same as that provided by `sp_defaultloc` — a default storage location is established for new and existing table creation, but automatic import of proxy table definitions does not take place during the processing of the `create database` command.

The value of path name is a string identifier in the following format: `servername.dbname.owner`.

- `<servername>` – required field; represents the name of the server that owns the objects to be referenced by proxy tables. Must exist in `master.dbo.sys.servers.srvname`.
- `<dbname>` – optional. The name of the database within `<servername>` which contains objects to be referenced by proxy tables
- `<owner >` – optional. The name of the owner of objects to be referenced by proxy tables. This may be restrictive, so that if more than one user owns objects in `<dbname>`, specifying the owner selects only those objects owned by that user. Do not create proxy tables for objects owned by other users.

If `for proxy_update` is specified with `no default_location`, an error is reported.

When a proxy database is created (using the `for proxy_update` option), Component Integration Services functions are called upon to:

- Provide an estimate of the database size required to contain all proxy tables representing the actual tables/ views found in the primary server's database. This estimate is provided in terms of the number of database pages needed to contain all proxy tables and indexes. This size is used if no size is specified, and no database devices are specified.

i Note

If the database is created with specific size specifications [`on device_name = nn`], or if a device name is specified with no size [`on device_name`], then the size requirements for the proxy database are not estimated; it is assumed in this case that the user or data base administrator wants to override the default size calculated for the proxy database.

If you are importing metadata from another SAP ASE, remote database users are imported before proxy tables are created. Each imported database user must have a corresponding system user name in `syslogins`.

- Create all proxy tables representing the actual tables/views found in the companion server's database. Proxy tables are not created for system tables.
- Grant all permissions on proxy tables to "public."
- Add the "guest" user to the proxy database.
- Import database users from remote site (if SAP ASE).
- Grant `create table` permission to "public."
- Set the database status to indicate that this database is a user proxy database. This is done by setting a status field in `master.dbo.sysdatabases.status3` (0x0001, DBT3_USER_PROXYDB).

After the database has been created, it contains a proxy table for each table or view found in the default location. The behavior for a user proxy database is identical to prior database behavior. Users can create additional objects, such as `procedure`, `views`, `rules`, `defaults`, and so on, and both DDL and DML statements that operate on proxy tables behave as documented in this book.

The only exception to this is the `alter database` command. The syntax and capabilities of this command are described in the next section.

2.2.1.1 User Proxy Database Schema Synchronization

At times, it may be necessary for a database administrator to force resynchronization of the proxy tables contained within the proxy database.

This can be done using the `alter database` command:

```
alter database <<dbname>>
    [alter database options]
    [for proxy_update]
```

If the `for proxy_update` clause is entered with no other options, the size of the database is not extended; instead, the proxy tables, if any, are dropped from the proxy database and re-created from the metadata obtained from the `<pathname>` specified during `create database ... with default_location = 'pathname'`.

If `create database` is used with other options to extend the size of the database, the proxy table synchronization is performed after the size extensions are made.

The purpose of this `alter database` extension is to provide a DBA with an easy-to-use, single-step operation with which to obtain an accurate and up-to-date proxy representation of all tables at a single remote site.

This resynchronization is supported for all external data sources, and not just the primary server in a HA-cluster environment. Also, a database need not have been created with the `for proxy_update` clause. If a

default storage location has been specified, either through the `create database` command or using `sp_defaultloc`, the metadata within the database can be synchronized with the metadata at the remote storage location.

Certain behavior is implied by the use of `create/alter database` to specify a proxy database:

- Modification to the default location specified with the `create database` command is not allowed using `alter database`.
- Local tables cannot be created in the proxy database. `create table` commands result in the creation of proxy tables, and the actual table is created at the default location.
- The default location of the table may be specified in the `create table` command, using the `at 'pathname'` syntax. If the path name differs from the default location, then the `alter database` command will not synchronize the metadata for this table.
- To change the default location, drop the database, then re-create it with a new path name specified in the `default_location = 'pathname'` clause. If the location is changed using `sp_defaultloc`, then the new location is used to provide metadata synchronization, and proxy tables that were created with the prior location not be synchronized, and may be dropped and replaced if the name conflicts with that of tables at the new location.

2.2.2 System Proxy Databases

System proxy databases behave like user proxy databases, with some notable enhancements and exceptions. System proxy databases are only used in a high availability configuration.

System proxy databases allow customer-written applications to run on either node in a high-availability cluster. This does not imply “single-system image” capability; rather, it suggests an environment in which most user-written applications can execute on either node in the cluster. This means that both databases and user-created objects should be visible to both nodes.

A system proxy database has the same name as the database in the primary node it references, and contains handling for the user-defined objects that are necessary to support the application. Proxy tables are created for each user table and view found in the primary database, and stored procedures are converted to RPCs and forwarded to the node referenced by the proxy database.

Certain commands, when executed within a system proxy database, are rejected with an error:

- `create procedure` and `drop procedure`
- `create view` and `drop view`
- `create trigger` and `drop trigger`
- `create rule` and `drop rule`
- `create default` and `drop default`

The error generated in these cases is: `Msg 12818, Severity 16: Cannot create an object of this type in system-created proxy database.`

2.2.2.1 System Proxy Database Creation

A system proxy database is created automatically under the specific circumstances.

Including:

- The HA cluster is being configured through the use of the stored procedure `sp_companion` `ServerName, 'configure', with_proxydb`.
In this case, a system proxy database is created for each user database found in server indicated by `ServerName`.
- A `create database` command is issued in a server whose HA state is one of `MODE_APNC`, `MODE_SNC`, or `MODE_ASNC`.

When the creation of the system proxy database is complete, Component Integration Services functions are called upon to:

- `grant create table to public` – this allows table creation on the primary server to result in proxy table creation in the system proxy database.

2.2.2.2 Schema Synchronization When Current Database Has a System Proxy Database

In an high availability (HA) cluster, some of the changes to a primary server's database must be forwarded to the companion server to keep both servers synchronized.

Several DDL commands, when executed within a database that has a system proxy database, cause notification of the companion server and result in automatic synchronization of the resulting changes:

- `create table` and `drop table` – local operation executes, resulting in the local table being created or dropped. The command is then forwarded to the companion server, for execution in the system proxy database, so that a proxy table can be created or dropped
- `create index` and `drop index` – local operation executes, resulting in an index being created or dropped. The server owning the system proxy database is then notified, and the proxy table is dropped and re-created, allowing the change to the index to be represented within the proxy table.
- `create view` and `drop view` – the local operation succeeds, resulting in the local view being created or dropped. The server owning the system proxy database is then notified, and a proxy table is either created or dropped.

If these commands are executed within the system proxy database, similar behavior occurs:

- `create table` and `drop table` – local proxy table is created or dropped. The command is then forwarded to the primary server, so that a local table referenced by the proxy table can be created or dropped.
- `create index` and `drop index` – local operation on the proxy table executes, resulting in an index being created or dropped. The server owning the primary database is then notified, and an index is either created or dropped on the local table referenced by the proxy table
- `create view` and `drop view` – not allowed within a system proxy database.

2.2.2.3 Stored Procedure Execution Within a System Proxy Database

If a system stored procedure request is encountered when the current database is a system proxy database, Component Integration Services attempts to locate the stored procedure first in the local `sybsystemprocs` database, and execute it.

If it is not found in `sybsystemprocs`, Component Integration Services searches the `master` database. If the procedure is not a system stored procedure, or if it is but cannot be found locally, the stored procedure request is converted to an RPC and transmitted to the server referenced by the system proxy databases default location.

2.3 File System Access

SAP ASE provides access to the file system through the SQL language. With file system access, you can create proxy tables that are mapped to file system directories, or to individual files.

i Note

Directories and files mapped to proxy tables now have a file path limit of 255 bytes.

Only SAP ASE users with system administrator (sa) or System Security Officer (sso) roles are allowed to create proxy tables that are mapped to files or directories

This requirement addresses the concerns over the security aspects of accessing file system data from within the SAP ASE server process (which may have root permission as it runs).

2.3.1 Directory Access

Proxy tables can be created to reference file system directories.

The syntax is:

```
create proxy_table <table_name>  
external directory at "directory pathname[;R]"
```

The directory path name must reference a file system directory visible to and searchable by the SAP ASE process. A proxy table that maps column names to attributes of files that exist within the directory is created. If the `;'R'` (indicating "recursion") extension is added to the end of the path name, Component Integration

Services includes entries in all subordinate directories. This table describes the proxy table columns that are created when this command successfully completes:

Table 3: Proxy table columns

Column Name	Datatype	Description
id	numeric (24) – on 32-bit machines numeric (36) – on 64-bit machines	Identity value consisting of values from <code>st_dev</code> and <code>st_ino</code> . These two values are converted first to a single string (format: "%d%014ld"), and the string is then converted to a numeric value.
filename	varchar (n)	The name of the file within the directory specified in at 'pathname', or within directories subordinate to pathname. The total length (n) of filename is limited to 255 bytes.
size	int	For regular files – specifies the number of bytes in the file. For directories – block special or character special, this is not defined.
filetype	varchar (4)	The file type – legal values are: FIFO, for pipe files; DIR for directories; CHRS for character special files; BLKS for block special files; REG for ordinary files; UNKN for all other file types. Links are automatically expanded, and do not appear as a separate file type.
access	char (10)	Access permissions, presented in a more or less 'standard' UNIX format: "drwxrwxrwx"
uid	varchar (n)	The name of the file owner. The value of n is specified by the system definition <code>L_cuserid</code> , which is 9 on all UNIX systems. This value is 0 on Windows systems.
gid	varchar (n)	The name of the owning group. The value of n is specified by the system definition <code>L_cuserid</code> , which is 9 on all UNIX systems. This value is 0 on Windows systems.
atime	datetime	Date/time file data was last accessed.
mtime	datetime	Date/time when file was last modified.
ctime	datetime	Date/time when file status was last changed.
content	image	The actual physical content of the file (for regular files only). NULL if the file is not a regular file.

A proxy table that maps to a file system directory can support the following SQL commands:

- `select` – file attributes and content can be obtained from the proxy table using the `select` command. Built-in functions that are designed to handle text values are fully supported for the content column, (for example, `textptr`, `textvalid`, `patindex`, `pattern`).
- `insert` – new files or directories can be created using the `insert` command. The only columns that have meaning are `filename`, `filetype`, and `content`. The rest of the columns should be left out of the `insert` statement, and are ignored if they are located. The `content` column is ignored if file type is `DIR`, which indicates that a new directory is to be created.

To create a new directory, enter:

```
insert D1 (filename, filetype) values ("newdir", "DIR")
```

To create a new file, enter:

```
insert D1 (filename, content) values ("newdir/newfile", "This is an example.")
```

- `delete` – files or directories may be removed by the use of the `delete` command. A directory can be removed only if it is empty. For example:

```
/* delete the files only */
delete D1 where filename = 'newdir/newfile'
/* deletes the directory (if empty) */
delete D1 where filetype = 'DIR' and filename = 'newdir'
```

- `update` – only the name of a file may be changed using the `update` command.

i Note

Some file systems implement `update` as a deletion followed by the creation of a new directory entry; therefore the same filename could be updated multiple times if the update is not restricted. SAP recommends that you qualify the `update` to specific files by having the filename included in the `where` clause of the update. For example, this statement could cause multiple updates:

```
update t1 set filename=filename + 'old' where filetype = 'REG'
```

The problem would be avoided by adding a clause such as `"and filename like "%.c"`

- `readtext` – the contents of a file may be retrieved using the `readtext` command.
- `writetext` – the contents of a file may be modified using the `writetext` command.

No other SQL commands operate on proxy tables.

Regular file content is available only if the SAP ASE process has sufficient privileges to access and read the file, and if the file type indicates an “ordinary” file. In all other cases, the content column is NULL. For example:

```
select filename, size, content
   from directory_table
  where filename like '%.html'
```

returns the name, size and content of regular files with a suffix of “.html,” if the SAP ASE process has access privileges to the file. Otherwise, the content column will be NULL.

`create proxy_table` fails if the path name referenced by directory path name is not a directory, or is not searchable by the SAP ASE process.

If trace flag 11206 is turned on, messages are written to the error log that contain information about the contents of the directories and the query processing steps needed to obtain that information.

2.3.2 Recursion Through Subordinate Directories

If the path name specified in the `create proxy_table` statement contains the `;R` extension, Component Integration Services traverses all directories subordinate to the path name, and returns information for the contents of each subordinate directory.

When this is done, the file name returned by a query contains the complete name of the file relative to the path name. In other words, all subordinate directory names appear in the file name. For example, if path name specifies `"/work;R"`:

```
create proxy_table d1 external directory at "/work;R"  
select filename, filetype from d1
```

Table 4: Values for files

File Name	File Type
dir1	DIR
dir1/file1.c	REG
dir1/file2.c	REG
dir2	DIR
dir2/file1.c	REG

2.3.3 File Access

Another class of proxy tables allowed in SAP ASE enables SQL access to individual files within a file system.

The syntax is:

```
create proxy_table <table_name>  
external file at " pathname"
```

When this command is used, a proxy table with one column (named "record", type `varchar(255)`) is created. It is assumed in this case that the contents of the file are readable characters, and that individual records within the file are separated by the newline (`\n`) character.

You can also specify your own column names and datatypes, using the `create [existing] table` command:

```
create existing table fname (  
column1 int null,  
column2 datetime null,  
column3 varchar(1024) null  
etc. etc.  
) external file at "pathname" [column delimiter "<string>"]
```

Columns may be any datatype except `text`, `image`, or a Java ADT. The use of the `existing` keyword is optional, and has no effect on the processing of the statement. If the file referenced by path name does not

exist, it is created. If it does exist, its contents are not overwritten. There is no difference in behavior between the `create table` and `create existing table` commands.

When a proxy table is mapped to a file, these assumptions about the file and its contents are made:

- The file is not a directory, block special, or character special file.
- The SAP ASE process has at least read access to the file. If the file is to be created, the server process must have write access to the directory in which the file is to be created.
- The contents of an existing file are in human-readable form.
- Records within the file are delimited by a newline character.
- The maximum supported record size is 32767 bytes.
- Individual columns, except for the last one, are delimited by the `column delimiter` string, which can be up to 16 bytes long; the default is a single tab character.
- There is a correspondence between delimited values within each record of the file and the columns within the proxy table.

With proxy tables mapped to files, you can:

- Back up database tables to the file system using either `select/into` or `insert/select`. When an `insert` statement is processed, each column is converted to characters in the default character set of the server. The results of the conversion are buffered, and all columns (except the last) are delimited by a single tab. The last column is terminated by a newline character. The buffer is then written to the file, representing a single row of data.
- Provide a SQL alternative to using `bcp in` and `bcp out`. The use of a `select/into` statement can easily back up a table to a file, or copy a file's contents into a table.
- Query file content with the `select` statement, qualifying rows as needed with search arguments or functions. For example, you can read the individual records within the SAP ASE error log file:

```
create proxy_table errorlog
  external file at "/usr/big_server/ASE15_0/install/errorlog"
select record from errorlog where record like "%server%"
```

This query returns all rows from the file that match the `like` pattern. If the rows are longer than 255 bytes, they are truncated. You can specify longer rows by entering:

```
create existing table errorlog
(
  record varchar(512) null
)
external file at "/usr/big_server/ASE15_0/install/errorlog"
```

In this case, records up to 512 bytes in length are returned. Since the proxy table contains only one column, the actual length of each column is determined by the presence of a newline character.

Only the `select`, `insert`, and `truncate table` statements are supported for file access. `update` and `delete` result in errors if the file proxy is the target of these commands.

When inserting values into a file, all datatypes are first converted to `char` values and then delimited by the column delimiter.

Caution

`truncate table` sets the file size to 0.

Trace flag 11206 is used to log messages to the error log. These messages contain information about the stages of query processing that are involved with file access.

2.4 Remote Servers

Use `sp_addserver` to add entries to the `sys.servers` table for the local server and for each remote server that is to be called.

The syntax is:

```
sp_addserver <server_name> [, <server_class> [, <network_name>]]
```

where:

- `<server_name>` – is a unique name used to identify the server.
- `<server_class>` – is the type of server. The supported server classes with the types of servers that are in each class are described in the following sections. The default is server class `ASEnterprise`.

i Note

Component Integration Services does not support server class `db2`.

- `<network_name>` – is the server name in the interfaces file. This name may be the same as `<server_name>`, or it may differ. The `<network_name >` is sometimes referred to as the `<physical name>`. The default is the same name as `<server_name>`.

i Note

You need the same sort order and case sensitivity between servers.

The server classes are:

- `ASEnterprise` – SAP ASE uses server class `ASEnterprise`. When Component Integration Services first establishes a connection to a server in this class, Component Integration Services determines the SAP ASE version and establishes server capabilities based on the version found.
- `ASAnywhere` – a server with server class `ASAnywhere` is an instance of SQL Anywhere. The version must be SQL Anywhere 9.0 or later.
- `ASIQ` – a server with server class `ASIQ` is SAP IQ version 12.5 or later
- `direct_connect` – server with server class `direct_connect` is an Open Server-based application that conforms to the `direct_connect` interface specification
Open Server-based applications using server class `direct_connect` are the preferred means of accessing all external, non-SAP data sources.
SAP ASE with Component Integration Services enables interacts with clients and Open Server-based applications. For example, a client application interacts with CIS on SAP ASE, which interacts through the network to a DirectConnect, such as Oracle or Informix. The DirectConnect can also have direct access to the client application, and the client application to the DirectConnect.
- `sds` – server with server class `sds` conforms to the interface requirements of a Specialty Data Store. A Specialty Data Store is an Open Server application you design to interface with SAP ASE.

- `RPCServer` – a server configured with `RPCServer` is not capable of handling anything other than RPCs. Servers in this class cannot participate in distributed transactions, and Component Integration Services does not attempt to send SQL statements to a server configured with this class. To send RPCs to Backup Server or to XP Server, the `sp_serveroption negotiated logins` and `server logins` must be enabled. Typically, servers in this class are customer-written Open Server applications intended to either perform a customized operation, or to make data available to an SAP ASE application that is generated as a result of one or more RPCs. The only Open Server event handlers required for this type of application are:
 - `SRV_CONNECT` – handle and authenticate a login request from CIS or client application.
 - `SRV_DISCONNECT` – handle disconnect request from CIS or client application.
 - `SRV_ATTENTION` – handle CANCEL request from CIS or client application.
 - `SRV_RPC` – handle RPC from CIS. The handling of the RPC may produce a result set, which CIS will forward to the ASE client on whose behalf CIS forwarded the RPC.

Using this server class, it is possible to write an Open Server application that supports CIS proxy tables that map to RPC's:

```
create existing table myRPCTable
(
    <column description(s)
)
external procedure at 'myRpcServerName...rpcname'
```

2.4.1 Connection Management

When connecting to a remote server on behalf of a client, Component Integration Services uses Client-Library functions.

Once the first connection to a remote server is established for a given client, that connection remains open until the client disconnects from Component Integration Services.

i Note

When a connection uses Component Integration Services features for the first time, it becomes affiliated with one SAP ASE engine. The PSS flag `POMNI_AFFINITY_SET` is set and is not cleared automatically.

2.4.2 Connect to Remote Servers Without the Interfaces File

You can establish a connection to remote servers without using corresponding entries in directory services `ldap` or `interfaces` files.

This is accomplished through Component Integration Services's use of the CT-Library connection property `CS_SERVERADDR`, which allows a server to be specified in the form:

```
"hostname.domain.com:99999"
"hostname:99999"
"255.255.255.255:99999"
```

where 99999 is the port number, and `<hostname>` is expressed as a simple name, an IP address, or a complete domain name.

Enter names in this format using `sp_addserver` with the net name argument:

```
sp_addserver S1, ASEnterprise, "myhost.big_server.com:11222"
```

or:

```
sp_addserver S1, ASEnterprise, "192.123.321.101:11222"
```

There are some limitations to this usage of net names:

- SAP ASE site handler does not recognize this syntax.
- Replication Agent threads do not recognize this syntax.

If this syntax is used, CT-Library does not attempt to look up connection information from directory services, whether an interfaces file or LDAP server is configured.

If SSL is configured and you have a pointer to the SSL section in the server docs, you can use the optional SSL syntax:

```
"hostname.domain.com:99999:SSL"  
"hostname:99999:SSL"  
"255.255.255.255.99999:SSL"
```

For more information about configuring SAP ASE for SSL, see *Security Administration Guide > Confidentiality of Data*.

2.4.3 LDAP Directory Services

You can use the LDAP directory services instead of the interfaces file in both the client and the server.

SAP ASE supports LDAP services for obtaining server information, and so does Component Integration Services. When a connection to a remote server is attempted, Component Integration Services instructs Open Client software to reference either the interfaces file or an LDAP server unless the net name argument to `sp_addserver` contains a colon (:).

Component Integration Services uses LDAP services only when the configuration file (`libtcl.cfg`) specifies it. `libtcl.cfg` can be found at `$$SYBASE/$SYBASE_OCS/config/libtcl.cfg` or `$$SYBASE/$SYBASE_OCS/config/libtcl64.cfg` for 64-bit applications.

i Note

When an LDAP server is specified in `libtcl.cfg`, server information becomes accessible from the LDAP server only and SAP ASE and Component Integration Services ignore any (traditional) interfaces file.

2.4.4 Secure Communication with SSL

Using SSL, you can establish secure connections from Component Integration Services to any number of remote servers that support the SSL protocol (SAP ASE and some DirectConnects).

Component Integration Services handles SSL connections as follows:

- The location of the trusted roots file is established. If the current server is SSL-enabled, then all outbound Component Integration Services connections will use the same trusted roots file as SAP ASE.
- If the current server is SSL-enabled, then a connection property is established to define the Open Client callback that will be used to respond to a challenge from a remote SSL-enabled server. If the current server is not SSL-enabled, then the callback used fails any connection to a remote SSL-enabled server.

2.4.4.1 Using Trusted Roots Files

The trusted roots file contains certificates for other servers that the local server treats as trusted when properly added to the system.

If `$SYBASE_CERT` is defined, a trusted roots file is accessible by the local server (SAP ASE) in:

```
$SYBASE_CERT/trusted.txt
```

Otherwise, it is here on UNIX platforms:

```
$SYBASE/$SYBASE_ASE/certificates/servername.txt
```

And here on Windows platforms:

```
%SYBASE%\%SYBASE_ASE%\certificates\servername.txt
```

where `servername` is the name of the current SAP ASE.

2.4.5 Secure Communication Using Kerberos

Kerberos network-based authentication is a single sign on feature which allows Kerberos clients authenticated with Kerberos system, to be able to connect to any application that supports Kerberos authentication. With one centralized password stored, you need not specify a password to connect to an application that supports Kerberos.

Context

Kerberos version 5, the version supported by SAP ASE, also provides a feature called credential delegation or ticket forwarding, which allows a Kerberos client to delegate the credential when connecting to a server, allowing the server to initiate Kerberos authentication for further connections to other servers on behalf of Kerberos client.

The credential delegation feature is currently only certified with MIT Kerberos GSSAPI libraries version 4.x and later. Clients must obtain a credential from the Kerberos system that can be delegated (using the `kinit -f` option on UNIX systems) before connecting to SAP ASE.

A Kerberos client connected to SAP ASE can request a Remote Procedure Call (RPC) to SAP ASE, and for general distributed query processing requests to a remote Adapter Server through CIS by using the Kerberos credential delegation feature. Kerberos authentication is not supported for site handler based remote server connection.

Procedure

- To use Kerberos unified login, a System Security Office can use the following command to enable the Kerberos security mechanism for CIS to a remote SAP ASE.

```
sp_serveroption [<server>, <optname>, <optvalue>]
```

For example, the following command executed on local server S1 enables Kerberos authentication for connections to remote server S2 when the current logged in user is authenticated using Kerberos mechanism.

```
sp_serveroption s2, "security mechanism", csfkrb5
```

2.4.5.1 Configuring Kerberos security services

Once the Kerberos security mechanism is enabled for connections to a remote SAP ASE, you can use a number of security services provided by Kerberos.

Context

Including:

- Message confidentiality
Data is encrypted over the network to protect against unauthorized disclosure.
- Message integrity
Verifies that communications have not been modified during transport.
- Mutual authentication
Verifies the identity of the client and the server. The local server initiating the remote connection can request mutual authentication for all remote connection requests to target an SAP ASE. This allows the client to verify the identity of the remote server.

i Note

The optional security services provided by Kerberos are not enabled by default.

This command, executed on local server S1, sets mutual authentication for all connections to remote server S2 using Kerberos authentication.

```
sp_serveroption s2, "mutual authentication", true
```

Procedure

- The following command executed on local server S1 sets message confidentiality for all connections to remote server S2 using Kerberos authentication:

```
sp_serveroption s2, "use message confidentiality", true
```

- This command, executed on local server S1, sets message integrity for all connections to remote server S2 using Kerberos authentication

```
sp_serveroption s2, "use message integrity", true
```

- This command, executed on local server S1, sets mutual authentication for all connections to remote server S2 using Kerberos authentication:

```
sp_serveroption s2, "mutual authentication", true
```

2.4.5.2 Configuring the Remote SAP ASE Kerberos Principal Name

The SAP ASE principal name is the default name of the server.

Context

Since the principal name of the SAP ASE can be different than the server name, the System Security Officer can specify the server principal name for each remote server.

Procedure

- The following command specifies a remote SAP ASE principle name for remote server S2:

```
sp_serveroption S2, "server principal",  
ase2@myrealm.com
```

2.4.5.3 Configuring Component Integration Services Remote Procedure Calls

CIS uses persistent client-library connections to handle the RPC request.

Context

CIS handles outbound RPCs by determining whether the client already has a client-library connection to the server in which the RPC is intended. If no connection exists, it will be established.

To enable the CIS RPC handling mechanism, set the configuration option `cis_rpc_handling` to 1. When not enabled, the Kerberos user needs to temporarily enable CIS RPC for the current session to use this feature.

This command enables CIS RPC handling for the current login session:

```
set cis_rpc_handling on
```

In the following example, `user1` is a Kerberos user who logs into SAP ASE S1 and request RPC to the remote SAP ASE S2.

Procedure

1. Add an entry to the interfaces file or the Directory Service for both servers S1 and S2 and a `secmech` line for the Kerberos security mechanism.
2. Add a login for the Kerberos user if one does not exist.

```
create login user1 with password pwuser1
```

3. Enable the use of security mechanisms by setting the configuration option to on.

```
sp_configure "use security services", 1
```

4. On the local server S1, enable Kerberos Authentication for CIS to remote server S2.

i Note

This assumes that remote server S2 only receives CIS command requests from S1. However, if S2 can also request CIS commands to other servers and requires enabling Kerberos Authentication, than similar configuration on S2 will be required.

- a. On the local server S1, add the remote server S2.

```
sp_addserver S2
```

- b. Enable Kerberos security mechanism on S1 for outbound RPC requests to S2. The following command enables CIS RPC handling for the current login session.

```
sp_serveroption S2, "security mechanism",
```

2.4.6 Security Issues

When establishing a connection to a remote SAP ASE, Client-Library functions are used instead of a site handler when either `cis_rpc_handling` or `set transactional_rpc` is on.

This method of establishing connections prevents the remote server from distinguishing these connections from those of other clients. Thus, any remote server security configured on the remote server to allow or disallow connections from a given server does not take effect.

Another SAP ASE with Component Integration Services enabled cannot use *trusted mode* for remote server connections. This forces the SAP ASE to be configured with all possible user accounts if it is going to be used with Component Integration Services.

Passwords are stored internally in encrypted form.

2.4.6.1 Using Encrypted Columns in CIS

By default, encryption and decryption are handled by the remote SAP ASE.

CIS makes a one-time check for encrypted columns on the remote SAP ASE. If the remote SAP ASE supports encryption, CIS updates the local `syscolumns` catalog with the encrypted-column-related metadata as follows:

- `create proxy_table` automatically updates `syscolumns` with any encrypted-column information from the remote tables.
- `create existing table` automatically updates `syscolumns` with any encrypted-column metadata from the remote tables. The `encrypt` keyword is not allowed in the `<column_list>` for `create existing table`. CIS automatically marks columns as encrypted if it finds any encrypted columns on the remote table.
- `create table` at the location with encrypted columns is not allowed.
- `alter table` is not allowed on encrypted columns for proxy tables.
- `select into existing` brings the plain text from the source and inserts it into destination table. The local SAP ASE then encrypts the plain text before insertion into any encrypted columns.

The following columns are updated from the remote server's `syscolumns` catalog:

- `enctype` – type of data on disk.
- `enclen` – length of encrypted data.
- `status2` – status bits that indicate that column is encrypted.

2.4.6.2 Using On-Demand Encryption in CIS

CIS supports command encryption for some commands and system procedures.

If the local and remote servers are using version 16.0 SP03 and later, you can enable the following `set` command option to encrypt commands sent between the servers, even if the client does not support encryption:

```
set remote command encryption {on|off}
```

These commands and procedures are supported for encryption between the servers:

- `connect to< server_name>`
- `sp_autoconnect`
- `sp_passthru`
- `sp_remotesql`
- `execute <server_name>... <proc_name>`

For more information, see *Security Administration Guide > On-Demand Encryption*.

2.4.7 Remote Server Logins

To fully support remote logins, Client-Library provides connection properties that enable Component Integration Services to request a `<server connection>`.

This connection is recognized at the receiving server as a server connection (as opposed to an ordinary client connection), allowing the remote server to validate the connection through the use of `sysremotelogins` as if the connection were made by a site handler.

Server connections are not enabled automatically. Instead, the SSO or DBA must request it by executing `sp_serveroption`:

```
exec sp_serveroption <server_name>,  
    'server login', true | false
```

You cannot change the `server login` property if the current server's `@@servername` global variable is NULL.

If the `server login` option is true, then Component Integration Services uses Client-Library connection properties to establish connections to the specified server.

Remote passwords specified by the client application are passed unchanged to the remote server. The use of and rules associated with remote passwords in server logins are identical to those associated with site handler connections.

These connection properties are only established if:

- The server option `server login` is set to true.
- The remote server is configured with server class `ASEnterprise`.
- There is a local server name defined for the Component Integration Services-enabled server (in other words the query `select @@servername` returns something other than NULL).

i Note

You can use trusted mode with Component Integration Services connections if you set “server logins” for a remote server

2.4.7.1 Connecting to Backup Server and XP Server

Component Integration Services can send RPCs to Backup Server or XP Server.

Procedure

- You must first enable the `negotiated_logins` server option, allowing Component Integration Services to respond to the login challenge initiated by either of these SAP-provided servers:

```
exec sp_serveroption server_name, "negotiated_logins", true
```

2.4.8 Map External Logins

SAP ASE users who invoke Component Integration Services require login names and passwords to remote servers.

By default, the user name and password pair used by Component Integration Services to connect to a remote server is the same as is used by the client to connect to SAP ASE.

Component Integration Services supports a one-to-one mapping of SAP ASE login names and passwords to remote server login names and passwords.

For example, using the stored procedure `sp_addexternlogin`, it is possible to map SAP ASE user `steve`, password `greatpassword` to Oracle login name `login1`, password `password1`:

```
sp_addexternlogin Oracle, steve, login1, password1
```

You can provide a many-to-one mapping so that all SAP ASE users who need an Oracle connection can be assigned the same name and password:

```
sp_addexternlogin Oracle, NULL, login2, password2
```

One-to-one mapping has precedence, so that if user `steve` has an external login for Oracle, that would be used rather than the many-to-one mapping.

In addition, you can assign external logins to SAP ASE roles. With this capability, anyone with a particular role can be assigned a corresponding login name/password for any given remote server:

```
sp_addexternlogin Oracle, null, login3, password3, rolename
```

The role name identifies the name of a role, rather than the name of a user. When a user with this role active requires a connection to Oracle, the appropriate login name/password for the role is used to establish the connection. When establishing a connection to a remote server for a user who has more than one role active, each role is searched for an external login mapping, and the first mapping found is used to establish the login. This is the same order as displayed by `sp_activeroles`.

The general syntax for `sp_addexternlogin` is:

```
sp_addexternlogin
  <<servername>>,
  <<loginname>>,
  <<external_loginname>>,
  <<external_password>>
  [, <<rolename>>]
```

`<rolename>` is optional; if specified, then `<loginname>` is ignored.

Precedence for these capabilities are as follows:

- If one-to-one mapping is defined, it is used.
- If no one-to-one mapping is defined, and a role is active and a mapping for it can be found, the role mapping is used to establish a remote connection.
- If neither of the above are true, then many-to-one mapping is used if defined.
- If none of the above is true, then the SAP ASE login name and password are used to make the connection.

If role mapping is done, and a user's role is changed (via `set role`), any connections made to remote servers that used role mapping are disconnected.

`sp_helpexternlogin` has been updated to allow viewing the various types of external logins that have been added using `sp_addexternlogin`. The syntax for `sp_helpexternlogin` is:

```
sp_helpexternlogin [<servername> [,<loginname> [,<rolename>]]]
```

All three parameters are optional, and any of the parameters can be NULL.

The stored procedure `sp_dropexternlogin` also accepts the `<rolename>` argument. If `<role name>` is specified then the second argument, `<login name>`, is ignored.

2.4.9 Configuring Failover for Remote Server Connections

If the interfaces file (or LDAP directory service) is set up to define a failover configuration, then Component Integration Services takes advantage of it by automatically failing over connections to the failover server if a connection to the primary server fails.

Context

You can set up remote servers for failover after performing these configuration steps:

Procedure

1. Enable new server option `cis hafailover`:

```
exec sp_serveroption server_name, 'cis hafailover', true
```

2. Modify directory services (interfaces file or server entries in the LDAP server) to specify a failover server. For example, you can configure server S2 to serve as a failover server for S1, and vice-versa, by additions to the interfaces file:

```
S1
    master tcp ether host1 8000
    query tcp ether host1 8000
    hafailover S2

S2
    master ether host2 9000
    query ether host2 9000
    hafailover S1
```

See *Using Sybase Failover in a High Availability System*, Appendix C, for more discussion of the `CS_HAFILOVER` connection property. Component Integration Services uses the `ct_con_props()` API to set this property, if the `cis hafailover` server option is `true`.

2.4.10 Remote Server Capabilities

The first time SAP ASE establishes a connection to a remote server of class `sds` or `direct_connect`, it issues an RPC named `sp_capabilities` and expects a result set in return.

This result set describes functional capabilities of the remote server so that Component Integration Services can adjust its interaction with that remote server to take advantage of available features. Component Integration Services forwards as much syntax as possible to a remote server, according to its capabilities.

2.5 Query Processing

The query processing steps taken when Component Integration Services is enabled are similar to the steps taken by SAP ASE.

The exceptions are:

- If a client connection is made in passthrough mode, the SAP ASE query processing is bypassed and the SQL text is forwarded to the remote server for execution.
- When `select`, `insert`, `delete`, or `update` statements are submitted to the server for execution, additional steps may be taken by Component Integration Services to improve the query's performance, if local proxy tables are referenced.

The steps are:

1. Query Parsing
2. Query Normalization
3. Query Preprocessing
4. Decision Point
5. SAP ASE Optimization and Plan Generation
6. Component Integration Services Plan Generation
7. Component Integration Services Remote Location Optimizer

Query Parsing

The SQL parser checks the syntax of incoming SQL statements, and raises an error if the SQL being submitted for execution is not recognized by the Transact-SQL parser.

Query Normalization

During query normalization, each object referenced in the SQL statement is validated. Query normalization verifies the objects referenced in the statement exist, and the datatypes are compatible with values in the statement.

For example:

```
select * from t1 where c1 = 10
```

The query normalization stage verifies that table `t1` with a column named `c1` exists in the system catalogs. It also verifies that the datatype of column `c1` is compatible with the value `10`. If the column's datatype is `datetime`, for example, this statement is rejected.

Query Preprocessing

Query preprocessing prepares the query for optimization.

It may change the representation of a statement such that the SQL statement Component Integration Services generates is syntactically different from the original statement.

Preprocessing performs view expansion, so that a query can operate on tables referenced by the view. It also takes steps such as reordering expressions and transforming subqueries to improve processing efficiency. For example, subquery transformation may convert some subqueries into joins.

Decision Point

After preprocessing, a decision is made as to whether Component Integration Services or the standard SAP ASE query optimizer handles optimization.

Component Integration Services handles optimization (using a feature known as quickpass mode) when:

- Every table represented in the SQL statement resides within a single remote server.
- The remote server is capable of processing all the syntax represented by the statement.
Component Integration Services determines the query processing capabilities of the remote server by its server class. For example, Component Integration Services assumes that any server configured as server class `sql_server` is capable of processing all Transact-SQL syntax.
For remote servers with server class `direct_connect`, Component Integration Services issues an RPC to ask the remote server for its capabilities the first time a connection is made to the server. Based on the server's response to the RPC, Component Integration Services determines the syntax of the SQL it forwards to the remote server.
- The following is true of the SQL statement:
 - It is a `select`, `insert`, `delete`, or `update` statement.
 - If it is an `insert`, `update`, or `delete` statement, there are no `identity` or `timestamp` columns, or referential constraints.
 - It contains no `text` or `image` columns.
 - It contains no `compute by` clauses.
 - It contains no `for browse` clauses.
 - It is not a `select...into` statement.
 - It is not a cursor-related statement (for example, `fetch`, `declare`, `open`, `close`, `deallocate`, `update`, or `delete` statements that include `where current of cursor`).
- The remote connection does not include a cursor opened on the remote server.

If the above conditions are not met, quickpass mode cannot be used, and the standard SAP ASE query optimizer handles optimization.

SAP ASE Optimization and Plan Generation

SAP ASE optimization and plan generation evaluates the optimal path for executing a query and produces a query plan that tells the SAP ASE how to execute the query.

If the `update statistics` command has been run for the tables in the query, the optimizer has sufficient data on which to base decisions regarding join order. If `update statistics` has not been run, the SAP ASE defaults apply.

For more information on SAP ASE optimization, see *Performance and Tuning Series: Query Processing and Abstract Plans > Using Statistics to Improve Performance*.

Component Integration Services Plan Generation

If quickpass mode can be used, Component Integration Services produces a simplified query plan in which the entire statement is pushed to a remote server.

If quickpass mode cannot be used, the SAP ASE optimizer generates a plan for executing the entire statement. This plan is then examined and portions of the plan are chosen to be pushed off to remote servers. As much of the original plan is pushed off as is possible based on the locations of the tables and the capabilities of the remote servers. The remote statement may come very close to the original statement for a fully capable remote server. A more minimal statement may be produced for other servers with the local SAP ASE executing the portion of the plan that could not be sent.

For example, if a client entered the statement:

```
select a,b from table1 where cos(a) > 0 and sin(b) > 0
```

If the remote server that owned `table1` supported `cos()` but not `sin()`, the statement sent to the remote server would be:

```
select a,b from table1 where cos(a) > 0
```

The local server would then have a plan that would apply the check for `sin(b) > 0` to the result set returned by the remote server.

Component Integration Services Remote Location Optimizer

SAP ASE generates a query plan containing the optimal join order for a multi-table query without regard to the storage location of each table.

If remote tables are represented in the query, Component Integration Services performs additional optimization taking location into account and possibly rearranging the plan for a join order that allows part of the join to be executed remotely.

To make intelligent plan choices, statistics are required for all tables involved in the query, including proxy tables. These are obtained by executing `update statistics` for a specific table.

If `update statistics` has not been run, the SAP ASE defaults apply.

The SAP ASE optimizer incorporates the cost of network access to remote servers based on a “network exchange” unit which specifies the time required to execute the sequence:

- Open a cursor

- Fetch 50 rows
- Close a cursor

The cost of a single exchange is under the user's control, and is specified on a per-server basis, defaulting to 1000 milliseconds, by `sp_serveroption`:

```
sp_serveroption <servername>, "server cost", "nnnn"
```

where `<nnnn>` is a string of numeric digits representing the number of milliseconds to be used per exchange during the optimizer's calculation of network cost.

i Note

The server cost limit is 32767. If you exceed that limit, an arithmetic overflow error occurs.

When a new server is added to `sys.servers` using `sp_addserver`, the default cost, 1000 milliseconds, is stored in `sysattributes` for that server. `sp_serveroption` can be used to specify a greater or lesser cost for a given server. `sp_helpserver` shows the current network cost associated with the server.

Any command that can affect a table is checked by the server to determine whether the object has a local or remote storage location

If the storage location is remote, then the appropriate access method is invoked when the query plan is executed in order to apply the requested operation to the remote objects. The following commands are affected if they operate on objects that are mapped to a remote storage location:

- | | | |
|-------------------------|-----------------------|---------------------|
| • alter table | • drop table | • rollback |
| • begin transaction | • drop index | • select |
| • commit | • execute | • set |
| • create index | • fetch | • setuser |
| • create table | • insert | • truncate table |
| • create existing table | • open | • update |
| • deallocate table | • prepare transaction | • update statistics |
| • declare cursor | • readtext | • writetext |
| • delete | | |

2.6 RPC handling and Component Integration Services

When Component Integration Services is enabled, you can choose between the site handler or Component Integration Services to handle outbound remote procedure calls (RPCs).

2.6.1 Site Handler and Outbound RPCs

Within an SAP ASE, outgoing RPCs are transmitted by means of a site handler, which multiplexes multiple requests through a single physical connection to a remote server.

The RPC is handled as part of a multistep operation:

1. Establish connection – the SAP ASE site handler establishes a single physical connection to the remote server. Each RPC requires that a logical connection be established over this physical connection. The logical connection is routed through the site handler of the intended remote server. The connection validation process for these connect requests is different from that of normal client connections. First, the remote server must determine if the server from which the connect request originated is configured in its `sys.servers` table. If so, then the system table `sys.remotelogins` is checked to determine how the connect request should be handled. If trusted mode is configured, password checking is not performed.
2. Transmit the RPC – the RPC request is transmitted over the logical connection.
3. Process results – all results from the RPC are relayed from the logical connection to the client.
4. Disconnect – the logical connection is terminated.

Because of the logical connect and disconnect steps, site handler RPCs can be slow.

2.6.2 Configuring Component Integration Services for Outbound RPCs

If Component Integration Services has been enabled, a client can use one of two methods to request that Component Integration Services handle outbound RPC requests.

Procedure

- Configure Component Integration Services to handle outbound RPCs as the default for all clients by issuing:

```
sp_configure "cis rpc handling", 1
```

If you use this method to set the `cis rpc handling` configuration parameter, all new client connections inherit this behavior, and outbound RPC requests are handled by Component Integration Services. This is a

server property inherited by all future connections. The client can, if necessary, revert back to the default SAP ASE behavior by issuing the command:

```
set cis_rpc_handling off
```

- Configure Component Integration Services to handle outbound RPCs for the current connection only by issuing:

```
set cis_rpc_handling on
```

This command enables `cis_rpc_handling` for the current thread only, and does not affect the behavior of other threads.

Results

When `cis_rpc_handling` is enabled, outbound RPC requests are not routed through the Adaptive Servers site handler. Instead, they are routed through Component Integration Services, which uses persistent Client-Library connections to handle the RPC request. Using this mechanism, Component Integration Services handles outbound RPCs as follows:

1. Determines whether the client already has a Client-Library connection to the server in which the RPC is intended. If not, establish one.
2. Sends the RPC to the remote server using Client-Library functions.
3. Relays the results from the remote server back to the client program that issued the RPC using Client-Library functions

RPCs can be included within a user-defined transaction. In fact, all work performed by Component Integration Services on behalf of its client can be performed within a single connection context. This allows RPCs to be included in a transaction's unit of work, and the work performed by the RPC can be committed or rolled back with the other work performed within the transaction.

The benefits of using Component Integration Services to handle outbound RPC requests are as follows:

- Client-Library connections are persistent so that subsequent RPC requests can use the same connection to the remote server. This can result in substantial RPC performance improvements, since the connect and disconnect logic is bypassed for all but the first RPC.
- Work performed by an RPC can be included in a transaction, and is committed or rolled back with the rest of the work performed by the transaction. This transactional RPC behavior is currently supported only when the server receiving the RPC is another SAP ASE or a DirectConnect which supports transactional RPCs.
- Connect requests appear to a remote server as ordinary client connections. The remote server cannot distinguish the connection from a normal application's connection, unless `server_logins` are enabled. This affects the remote server management capabilities of an SAP ASE, since no verification is performed against `sysremotelogins`, and all connections must have valid SAP ASE login accounts established prior to the connect request (trusted mode cannot be used in this case).

2.6.3 Text Parameters for RPCs

SAP ASE can send large chunks of data in a single remote procedure call.

This is done by treating certain parameters as text pointers, then dereferencing these text pointers to obtain the text values associated with them. The `text` data is then packaged into 16K chunks for SAP ASE and 32K chunks for all other servers, and handed to Client-Library as parameters to the RPC.

A text pointer is identified as a parameter of type `binary(16)`. The text value referenced by each text pointer parameter is obtained when the RPC is executed, and expanded into 16K chunks for SAP ASE and 32K chunks for all other servers, each of which is passed to Client-Library as a parameter of type `CS_LONGCHAR_TYPE`.

This behavior is triggered by this `set` command:

```
set textptr_parameters ON
```

When an RPC is requested (`cis_rpc_handling` must be on), text pointers are dereferenced in the Component Integration Services layer, and the text value obtained is used to construct one or more parameters for Client-Library.

For this to work, the text pointers must be preceded by a path name argument, which is used to identify the table from which the text pointers have been derived. For example:

```
declare @pathname varchar(90)
declare @textptr1 binary(16)
declare @textptr2 binary(16)
select @pathname = "mydatabase.dbo.t1",
       @textptr1 = textptr(c1),
       @textptr2 = textptr(c2)
from mydatabase.dbo.t1
where ... (whatever)
set textptr_parameters ON
exec big_server...myrpc @pathname, @textptr1, @textptr2
set textptr_parameters OFF
```

When the RPC named 'myrpc' gets sent to server `big_server`, the `<@pathname>` parameter is not actually sent, but is used to help locate the text values referenced by the textptr's `<@textptr1>` and `<@textptr2>`.

The `varchar` parameter `<@pathname>` must immediately precede the `binary(16)` parameter, otherwise `<@textptr1>` is considered an ordinary parameter and is transmitted to the server `big_server` as a normal `binary(16)` value.

The text will be broken into 16K or 32K chunks, each of which is a separate parameter of type `CS_LONGCHAR_TYPE`.

The current value of `<@@textsize>` is ignored.

This scheme is also designed to work with proxy tables mapped to remote procedures. For example:

```
create existing table myrpctable
(
  id int,          -- result column
  crdate datetime, -- result column
  name varchar(30), -- result column
  _pathname varchar(90), -- parameter column
  _textptr1 binary(16), -- parameter column
  _textptr2 binary(16), -- parameter column
) external procedure at 'big_server...myrpc'
go
```

```

declare @textptr1 binary(16)
declare @textptr2 binary(16)
select @textptr1 = textptr(c1), @textptr2 = textptr(c2)
from mydatabase.dbo.t1 where <whatever>
set textptr_parameters ON
select id, crdate, name
from myrpctable
where pathname = "mydatabase.dbo.t1" and
       _textptr1 = @textptr1 and
       _textptr2 = @textptr2

```

When the query against the proxy table `myrpctable` is processed, Component Integration Services sends an RPC named 'myrpc' to the server 'big_server'. The parameters will be derived from the search arguments contained in the `where` clause of the query. Since the `textptr_parameter` option has been set ON, the `textptrs` are expanded to `CS_LONGCHAR_TYPE`, as in the case of the RPC example shown previously.

2.6.4 Text Parameter Support for XJS/390

Because of the ability to forward large blocks of text as RPC parameters, Component Integration Services can interact with IBM mainframes using XJS/390. XJS/390 scripts (JavaScript-like syntax) stored within SAP ASE tables (or files accessible via proxy tables), and forwarded to the mainframe using an RPC.

The syntax of the script is analyzed and executed by XJS/390 facilities, and result sets are generated according to the procedural logic of the script.

Several features are enabled:

- Database events within SAP ASE can result in the generation of an MQ Series message. Since XJS/390 Mscript supports the generation of messages, an RPC can be sent to the mainframe to request that such a message be generated in response to a triggered event within the database.
- Component Integration Services users have access to VSAM, IMS, and MQSeries data without the need to install third-party middleware such as InfoHub.

Version 2.0 or later of XJS/390 is required for handling scripts as RPC parameters. See the XJS/390 specification for details.

2.7 Distributed Transaction Management

Distributed Transaction Management within SAP ASE tracks the state of a transaction in the local SAP ASE/Component Integration Services, as well as in all remote servers participating in transactions.

When a user application commits a transaction, the commit is propagated to all participating remote servers using SAP ASE Transaction Coordinator (ASTC). The management of multisite transactions is handled by ASTC in cooperation with Component Integration Services. Component Integration Services registers new participating servers for each transaction, then turns over control of the transaction coordination to ASTC, which calls back into Component Integration Services to execute various commands for transaction management.

Internally, ASTC views a server as either:

- DTM-enabled
- Pre-DTM

These types map to the three sets of callbacks used, and map to server classes as indicated here:

Table 5: Transaction capabilities

ASTC server type	Component Integration Services server class
DTM-enabled	<ul style="list-style-type: none"> • ASEnterprise (12.x or later) • DC/Oracle 12.5 or later
Pre-DTM	<ul style="list-style-type: none"> • ASEnterprise (pre-12.0) • ASAnywhere • ASIQ • other Direct Connect • sds

i Note

Before starting a distributed transaction, the local server must be named. `<@@servername >` cannot be null.

2.7.1 DTM-Enabled Servers

Remote servers that are “DTM-enabled” support the full two-phase commit service enabled by ASTC.

Servers that support this must allow a separate connection (or session) to either commit or roll back a transaction that was begun by another session. This capability is necessary if the commit coordinator (ASTC) is required to connect to a remote site and commit or roll back in-doubt transactions. SAP ASE 12.0 and later provide this support, as does DirectConnect for Oracle 12.5 or later.

2.7.2 Pre-DTM servers

Remote servers that are classified as “pre-DTM” are those that support transaction management statements such as `begin tran`, `commit tran`, `rollback tran`, but does not support one session’s ability to commit or rollback a transaction started by another session.

Component Integration Services makes every effort to manage user transactions for pre-DTM servers reliably. However, different access methods incorporated into the server allow varying degrees of support for this capability. The general logic described below is employed by server classes `ASEnterprise` (prior to 12.0), `ASAnywhere`, `ASIQ`, `direct_connect`, and `sds` if the Specialty Data Store supports transaction management.

The method for managing transactions involving remote servers uses a two-phase commit protocol. SAP ASE implements a strategy that ensures transaction integrity for most scenarios. However, there is still a chance

that a distributed unit of work will be left in an undetermined state. Even though two-phase commit protocol is used, no recovery process is included. The general logic for managing a user transaction is as follows:

Component Integration Services prefaces work to a remote server with a `begin transaction` notification. When the transaction is ready to be committed, Component Integration Services sends a `prepare transaction` notification to each remote server that has been part of the transaction. `prepare transaction` pings the remote server to determine whether the connection is still viable. If a `prepare transaction` request fails, all remote servers are told to roll back the current transaction. If all `prepare transaction` requests are successful, the server sends a `commit transaction` request to each remote server involved with the transaction. Any command preceded by `begin transaction` can begin a transaction. Other commands are sent to a remote server to be executed as a single, remote unit of work.

2.7.3 strict DTM enforcement

To ensure complete two-phase commit capability, ASTC uses the concept of `strict dtm enforcement`.

When enabled, `strict dtm enforcement` causes a transaction to abort if an attempt is made to include a pre-DTM server in the transaction.

2.7.4 Enable Component Integration Services

ASTC relies on Component Integration Services to handle all communication with remote servers.

ASTC uses the configuration option `enable xact coordination`. Since ASTC is enabled by default, Component Integration Services is also enabled by default.

This option, enabled by default, allows ASTC to manage all transactions involving remote servers. You must enable Component Integration Services before `xact coordination` is enabled. While `xact coordination` is enabled, Component Integration Services cannot be disabled. When `xact coordination` is enabled, `transactional_rpcs` are implicitly enabled.

2.7.5 Transactional RPCs

The server allows RPCs to be included within the unit of work initiated by the current transaction.

Before using transactional RPCs, issue the `set transactional_rpc on` command.

Assuming that the remote server can support the inclusion of RPCs within transactions, the following syntax shows how this capability might be used:

```
begin transaction
  insert into t1 values (1)
  update t2 set c1 = 10
  execute @status = big_server.pubs2.dbo.myproc
  if @status = 1
    commit transaction
  else
```

```
rollback transaction
```

In this example, the work performed by the procedure `<myproc>` in server `big_server` is included in the unit of work that began with the `begin transaction` command. This example requires that the remote procedure `<myproc>` return a status of "1" for success. The application controls whether the work is committed or rolled back as a complete unit.

The server that is to receive the RPC must allow RPCs to be included in the same transactional context as Data Manipulation Language (DML) commands (`select`, `insert`, `delete`, `update`). This is true for SAP ASE and is expected to be true for most DirectConnect products released by SAP. However, some database management systems may not support this capability.

2.7.6 Restrictions on Transaction Management

If nested `begin transaction` and `commit transaction` statements are included in a transaction that involves remote servers, only the outermost set of statements is processed.

The innermost set, containing the `begin transaction` and `commit transaction` statements, is not transmitted to remote servers.

2.8 SAP ASE to SAP ASE update statistics

When you perform `update statistics` on a remote server proxy table, if the relevant tables, index and column statistics are available, the table catalogs are imported to the local `sysabstats` and `sysstatistics`.

By default, `update statistics` for proxy tables always attempts to import the required statistics data. But when the statistics data is unavailable or incomplete on the remote table, Component Integration Services (CIS) reverts to the prior mechanism of gathering statistic data.

You can also force CIS to revert to the prior mechanism of gathering statistic data by turning on Traceflag 11229. This allows you to obtain all data from the database, then calculating the statistics.

i Note

This is the behavior if `update statistics` has not been run on the remote tables and there are no statistics available. This behavior excludes proxy tables mapped to RPCs, external files, and system directories.

2.9 Update Statistics on Non-SAP ASE Backends

The `update statistics` command helps the server make the best decisions about which indexes to use when it processes a query, by providing information about the distribution of the key values in the indexes.

`update statistics` does not automatically run when you create or re-create an index on a table that already contains data. It can be used when a large amount of data in an indexed column has been added, changed, or deleted. The crucial element in query optimization is the accuracy of the distribution steps. If there are significant changes in the key values in the index, re-run `update statistics` on that index.

Only the table owner or the system administrator can issue the `update statistics` command.

The syntax is:

```
update statistics <table_name> [<index_name>]
```

Because running `update statistics` is resource intensive, try to run `update statistics` at a time when the tables you specify are not heavily used. `update statistics` acquires locks on the remote tables and indexes as it reads the data. If you use trace flag 11209, tables are not locked.

You can set `update statistics` to run automatically at the time that best suits your site and avoid running it at times that hamper your system. For more information see *Performance and Tuning Guide: Monitoring and Analyzing > Using Statistics to Improve Performance*.

The server performs a table scan for each index specified in the `update statistics` command.

Since Transact-SQL does not require index names to be unique in a database, you must give the name of the table with which the index is associated.

After running `update statistics`, run `sp_recompile` so triggers and procedures that use the indexes use the new distribution:

```
sp_recompile authors
```

2.10 Java in the Database

Java in the database is supported for remote data access with Component Integration Services.

The following restrictions apply:

- Java is supported for remote SAP ASE 12.x and later only.
- Java is supported for language events only (no dynamic SQL can be used with remote tables.)

After installing your Java class files on the local server, install the required Java class files on the remote server.

- `@@textsize` – data is returned as a serialized Java object using the `image` datatype format and then deserialized on the local server. `@@textsize` must be set large enough to hold the serialized object. If `@@textsize` is set too small, the object is truncated, and the deserialization fails.
- `@@stringsize` – indicates the amount of `character` data to be returned from a `toString()` method. It is similar in behavior to `@@textsize`, except it applies only the `char` data returned by the Java

`Object.toString()` method. The default value is 50. The maximum value is 16384. A value of zero means “use the default.” This value can be modified by a `set` command:

```
set stringsize n
```

where `<n>` is an integer value between 0 and 16384. The value immediately displays in the global variable `@@stringsize`.

Related Information

[Java Class Definitions \[page 53\]](#)

2.10.1 Constraints on Java Class Columns

Constraints defined on Java columns of remote tables must be checked on the remote server.

If the constraint checking is attempted on the local server, it fails. Therefore, you must enable trace flag 11220 when you `insert`, `update`, or `delete` data for which constraint checking is done on Java datatypes.

There are two error messages that are specific to Java use with remote data access:

- Error 11275 – a statement referencing an extended datatype contained syntax that prevented it from being sent to the remote server. Rewrite the statement or remove the extended datatype reference.
- Error 11276 – an object in column '`<colname>`' could not be deserialized, possibly because the object was truncated. Check that the value of `@@textsize` is large enough to accommodate the serialized object.

Related Information

[Trace Flags \[page 73\]](#)

2.10.2 Java Abstract Datatypes (ADTs)

Java Classes in SQL (JCS) is the method of storing and using Java objects within the SAP ASE. Component Integration Services interaction in this implementation is needed to support Java objects and Java functions on remote servers.

Objects are passed between the local and remote servers in a serialized format that is a binary representation used to reinstantiate the object. Component Integration Services treats a serialized object as an `image` blob, using `text` and `image` handling functions to pass objects between servers. The object is reinstantiated on the destination server before processing continues.

When handling queries containing references to Java objects and functions on remote servers, Component Integration Services attempts to forward as much syntax as possible to the remote server. Any portion of the

query that cannot be passed to the remote server is handled on the local server, requiring the serialization and deserialization of all necessary remote objects. Due to the overhead associated with serializing and deserializing Java objects, performance of such queries is significantly less than comparable local access.

To facilitate the interchange of Java objects between servers, Component Integration Services issues:

```
set raw_object_serialization ON
```

to each `ASEnterprise` server that is Java-enabled. This allows Component Integration Services to easily deserialize the object obtained from the remote site.

2.10.2.1 Java Class Definitions

The Java class definitions on the local and remote servers must be compatible to facilitate passing objects between servers.

For this reason, Component Integration Services assumes that compatibility exists, and any errors in object definition are detected during deserialization efforts. Objects are considered compatible if the serialized form of the object on the remote server can be used to successfully instantiate an object on the local server, or vice versa. Also, any Java method referenced in the local server in conjunction with a remotely mapped object must be defined on the remote object as well.

It is the responsibility of the database administrator to ensure that class definitions on local and remote servers are compatible. Incompatible objects and invalid method references result in deserialization errors or Java exceptions that cancel the requesting query.

To improve overall performance, increase the `cis packet size` configuration variable to better facilitate passing serialized objects between servers. Serialized objects are passed between servers with an `image` datatype, and can vary in size from a few bytes to 2GB.

2.11 Datatypes

Component Integration Services supports Unicode, datatype conversions, and text and image datatypes.

2.11.1 Unicode support

SAP ASE contains formal support for the Unicode character set.

The datatypes provided are `unicar`, `univarchar`, and `unitext`. They comprise 2-byte characters expressed in Unicode. SAP ASE provides conversion functions between Unicode data and all other datatypes, consistent with current handling of `char` and `varchar` datatypes. By supporting these datatypes, Component Integration Services is able to present a view of all enterprise character data expressed in Unicode. Character data from mainframes and all other foreign or legacy systems is converted to Unicode when columns of type `unicar` or `univarchar` are used to defined columns in proxy tables.

The Component Integration Services features below are affected by these new datatypes:

- `create table` – may contain columns described using the new Unicode datatypes. If the table to be created is a proxy table, Component Integration Services forwards the entire command, including the Unicode datatype names (`unicchar`, `univarchar`, and `unitext`) to the remote server where the new table is to be created. If the remote server cannot handle the datatypes, it raises an error.
- `create existing table` – when comparing SAP ASE column types and lengths with the metadata obtained from a remote server, Unicode datatypes in the proxy table are allowed under the following circumstances:
 - The remote server datatype for a column is `unicchar`, `unitext`, or `univarchar` with equal length (expressed in characters, not bytes).
 - The remote server datatype for a given column is `char` or `varchar`. In this case, Component Integration Services performs conversions to Unicode on data fetched from the remote server, and conversions from Unicode to the default SAP ASE character set (UTF8) on data transmitted as part of DML commands (`select`, `insert`, `delete`, `update`).
 - The remote server datatype for a Unicode column is `binary` or `varbinary`. The length of the remote server column must be twice the length of the Unicode column. Component Integration Services performs conversions as required when transmitting data to or from the remote server. No other datatype mapping for Unicode datatypes is allowed when mapping a proxy table to a remote table. Other types result in a type mismatch error. You can convert data from legacy systems into Unicode simply by creating a proxy table that maps a Unicode column to an existing `char` or `varchar` column.

i Note

Unicode can only be mapped to `unitext` columns using the `create existing table` command.

- `create proxy_table` – by using `create proxy_table`, an SAP ASE user does not have to specify the column list associated with the proxy table. Instead, the column list is derived from column metadata imported from the remote server on which the actual table resides. Unicode columns from the remote server are mapped to Unicode columns in the proxy table only when the remote column is datatype `unicchar`, `unitext`, or `univarchar`.
- `alter table` – allows column types to be modified. With SAP ASE version 12.5 and later, a column's type can be modified to and from Unicode datatypes. If the command operates on a proxy table, the command is reconstructed and forwarded to the remote server that owns the actual table. If the remote server (or DirectConnect) cannot process the command, an error is expected, and the SAP ASE command is aborted.
If trace flag 11221 is on, `alter table` does not get forwarded to a remote server; adding, deleting, or modifying columns is done locally on the proxy table only.
Using the `alter table` command, `unitext` can be changed to `char`, `varchar`, `nchar`, `nvarchar`, `unicchar`, `univarchar`, `binary`, and `varbinary`. Any of these datatypes can be changed to `unitext`.
- `select`, `insert`, `update`, and `delete` statements – Unicode datatypes impact the processing of `select` statements in two ways when proxy tables are involved. The first involves the construction of SQL statements and parameters that are passed to remote servers; the second involves the conversion of data to Unicode when Component Integration Services fetches non-Unicode data.
A DML command involving a proxy table is handled using either TDS language requests or TDS cursor requests when interacting with the remote server. If a `select` statement contains predicates in the `where` clause that involve Unicode columns and constants, the Unicode constants must be handled in one of two ways, depending on whether language or cursor commands are used to process the statement:

1. TDS language – generate clear-text values that can be included in the language text buffer. This involves converting a constant Unicode value to clear text values that can be transmitted as part of a language request.
2. TDS cursor – generate Unicode parameters for CT-Library cursor requests. Parameter values may be Unicode data, requiring Component Integration Services to use parameter types of CS_UNICHAR_TYPE.

Component Integration Services handles an `insert` command involving a proxy table using either TDS language requests or TDS dynamic requests.

If the `insert` command can be processed in quickpass mode, then TDS language requests are used. If the command cannot be handled in quickpass mode, the `insert` is processed using TDS Dynamic requests.

In language requests, the issues are the same as with `select` – Unicode values must be converted to clear-text form so they can be transmitted with the rest of the SQL statement. In dynamic requests, Unicode data (along with all other data values) is transmitted as parameters to the dynamic command.

The receiving server is expected to process parameters of type CS_UNICHAR_TYPE.

The issues with `update` and `delete` commands are the same as for `select` and `insert`. Unicode values must be converted either to clear-text characters for transmission with the rest of the SQL statement, or they must be converted into parameters of type CS_UNICHAR_TYPE.

2.11.2 Datatype Conversions

Datatype conversion can take place whenever the server receives data from a remote source, be it SAP ASE, or an Open Server-based application.

Depending on the remote datatype of each column, data is converted from the native datatype on the remote server to a form that the local server supports.

Datatype conversions are made when the `create table`, `alter table` and `create existing table` commands are processed. The datatype conversions are dependent on the remote server's server class. See the `create table`, `alter table`, and `create existing table` commands for tables that illustrate the datatype conversions that take place for each server class when the commands are processed.

Related Information

[SQL Reference \[page 72\]](#)

2.11.3 text and image Datatypes

The `text` datatype is used to store printable character data, the column size of which depends on the logical page size of the SAP ASE.

The `image` datatype is used to store a number of bytes of hexadecimal-encoded binary data that, again, depends on the logical page size of the SAP ASE. The maximum length for `text`, `image`, and `unitext` data is defined by the server class of the remote server to which the column is mapped.

i Note

`unitext` with Component Integration Services is only supported for SAP ASE version 15.0 and higher.

`text`, `image`, and `unitext` columns cannot be used:

- As parameters to stored procedures, except when `set textptr_parameters` is on
- As local variables
- In `order by`, `compute`, or `group by` clauses
- In indexes
- In subqueries
- In `where` clauses, except with the keyword `like`
- In joins

Processing `text`, `unitext`, and `image` data (with remote servers only) can be logged by using trace flag 11207.

2.11.3.1 Limits of @@textsize

`select` statements return `text`, `image`, and `unitext` data up to the limit specified in the global variable `<@@textsize>`.

The `set textsize` command is used to change this limit. The initial value of `<@@textsize>` is 32K; the maximum value for `< @@textsize>` is 2147MB.

2.11.3.2 Odd Bytes Padded

`image` values of less than 255 bytes that have an odd number of bytes are padded with a leading zero (an insert of "0xaaabb" becomes "0x0aaabb").

It is an error to `insert` an `image` value of more than 255 bytes if the value has an odd number of bytes.

2.11.3.3 Converting text and image Datatypes

You can explicitly convert `text` values to `char` or `varchar` and `image` values to `binary` or `varbinary` with the `convert` function.

However, you are limited to the maximum length of the `character` and `binary` datatypes, which depends on the logical page size of the SAP ASE. If you do not specify the length, the converted value has a default length of 30 bytes. Implicit conversion is not supported.

2.11.3.4 Pattern Matching with text and unitext Data

Use the `patindex` function to search for the starting position of the first occurrence of a specified pattern in a `text`, `unitext`, `varchar`, or `char` column.

The `%` wildcard character must precede and follow the pattern (except when you are searching for the first or last character).

Use the `like` keyword to search for a particular pattern. This example selects each `text` data value from the `blurb` column of the `texttest` table that contains the pattern "Straight Talk%":

```
select blurb from texttest
       where blurb like "Straight Talk%"
```

You can use the keyword `like` to search a `unitext` column for a specific pattern. However, the `like` clause is not optimized when it is used with a `unitext` column. `like` pattern matching for `unitext` depends on the default Unicode sort order, which is also used for `like` pattern matching for `unichar` and `univarchar` datatypes.

2.11.3.5 Entering text and image Values

The DB-Library™ functions `dbwritetext` and `dbmoretext` and the Client-Library function `ct_send_data` are the most efficient ways to enter `text`, `unitext`, and `image` values.

2.11.3.6 readtext using bytes

Errors result if you use the `readtext using bytes` command on a `text` column, and the combination of size and offset result in the transmission of a partial character.

2.11.3.7 Restriction for text, image, and unitext with bulk copy

When you use `bulk copy` to copy `text`, `unitext`, and `image` values to a remote server, the server must store the values in data pages before sending them to the remote server.

Once the values have been issued to the remote server, the data pages are released. Data pages are allocated and released row by row. This is important because:

- The overhead of allocating and releasing data pages impacts performance.
- Data pages are allocated in the database where the table resides, so the database must be large enough to accommodate enough data pages for the largest `text`, `unitext`, and `image` values that exist for any given row.

2.11.3.8 Considerations for Using text, unitext, and image data with Server Classes

The ASEnterprise server class includes some restrictions for text, unitext, or image data.

- ASEnterprise:
 - A pointer in a text, unitext, or image column is assigned when the column is initialized. Before you can enter text, unitext, or image data into a column, the column must be initialized. This causes a 2K page to be allocated on the remote or SAP ASE. To initialize text, unitext, or image columns, use the update with a NULL or a non-null insert command.
 - Before you use writetext to enter text or unitext data or readtext to read it, the text or unitext, column must be initialized. Use update or insert non-null data to initialize the text column, and then use writetext and readtext.
 - Using update to replace existing text, unitext, and image data with NULL reclaims all of the allocated data pages, except the first page, in the remote server.
 - writetext, select into, DB-Library functions, or Client-Library functions must be used to enter text, unitext, or image values that are larger than 16KB.
 - readtext is the most efficient way to access text, unitext, and image data.
 - insert select and select into can be used to insert text, unitext, and image data to proxy tables, but a unique index is required.
- direct_connect:
 - Specific DirectConnect servers support text and image data to varying degrees. See the DirectConnect documentation for information on text, unitext, and image support.
 - The server uses the length defined in the global variable < @@textsize> for the column length. Before issuing create table, the client application should set < @@textsize> to the required length by invoking set textsize.
 - For DirectConnect servers that support text, unitext, and image datatypes but do not support text pointers, the following restrictions apply:
 - The writetext command is not supported.
 - The readtext command is not supported.
 - Client-Library functions that use text pointers are not supported.
 - DB-Library functions that use text pointers are not supported.
 - For DirectConnect servers that support text, unitext, and image datatypes but do not support text pointers, some additional processing is performed to allow the following functions to be used:
 - patindex
 - char_length
 - datalengthIf text pointers are supported, the server performs these functions by issuing an RPC to the DirectConnect server.
 - For DirectConnect servers that do not support text pointers, the server stores data in the sysattributes system table. Data pages are preallocated on a per column per row basis. The column size is determined by < @@textsize>. If this value is not sufficient an error is returned.
 - Specific DirectConnect servers may or may not support pattern matching against the text datatype. If a DirectConnect server does not support this pattern matching, the server copies the text value to internal data pages and performs the pattern matching internally. The best performance is seen when pattern matching is performed by the DirectConnect server.

- You must use `writetext`, `select into`, or `insert...select` to enter text, unitext, or image values that exceed 450 bytes.
- You can use `select into` and `insert...select` to insert text, unitext, or image values, but the table must have a unique index.

2.12 Configuration and Tuning

The system administrator or database owner may elect to use the server to optimize performance or to allow use by a required number of clients. Configuration choices might involve being able to review total numbers of reads and writes for a given SQL command.

Once an application is running, the system administrator should monitor performance and may choose to customize and fine-tune the system. The server provides tools for these purposes, including:

- Changing system parameters with the `sp_configure` procedure
- Using `update statistics` to ensure that Component Integration Services makes the best use of existing indexes
- Monitoring server activity with the `dbcc` command
- Setting trace flags
- Executing `ddlgen` and related backup and recovery issues
- Determining database size requirements

2.12.1 Using `sp_configure`

The configuration parameters in `sp_configure` control resource allocation and performance.

Context

The system administrator can reset these configuration parameters to tune performance and redefine storage allocation. In the absence of intervention by the system administrator, the server supplies default values for all the parameters.

A number of configuration parameters are unique to Component Integration Services

- `enable cis` – Use this parameter with `sp_configure` to enable Component Integration Services. Issuing `sp_configure "enable cis", 0` disables Component Integration Services after restarting the server.
- `enable file access` – enables access through proxy tables to eXternal File System.
- `enable full-text search` – enables Enhanced Full-Text Search services. Requires a license for ASE_EFTS
- `max cis remote connections` – a non-zero value indicates the number of connection data structures pre-allocated during server initialization. The default is zero.

- `cis bulk insert batch size` – determines how many rows from the source tables are to be bulk copied into the target table as a single batch using `select into`, when the target table resides in an SAP ASE or in a DirectConnect server that supports a bulk copy interface.
If left at zero (the default), all rows are copied as a single batch. Otherwise, after the count of rows specified by this parameter has been copied to the target table, Component Integration Services issues a bulk commit to the target server, causing the batch to be committed.
If a normal client-generated bulk copy operation (such as that produced by the `bcf` utility) is received, the client is expected to control the size of the bulk batch, and Component Integration Services ignores the value of this configuration parameter.
- `cis bulk insert array size` – When performing a bulk transfer of data from one SAP ASE to another, Component Integration Services buffers rows internally, and asks the Open Client bulk library to transfer them as a block. The size of the array is controlled by the configuration parameter `cis bulk insert array size`. The default is 50 rows, and the property is dynamic, allowing it to be changed without server reboot.
- `cis cursor rows` – allows users to specify the cursor row count for `cursor open` and `cursor fetch` operations. Increasing this value means more rows are fetched in one operation. This increases speed but requires more memory. The default is 50.
- `cis packet size` – allows you to specify the size of Tabular Data Stream™ (TDS) packets that are exchanged between Component Integration Services and a remote server when connection is initiated. The default packet size on most systems is 512 bytes, which is adequate for most applications. However, larger packet sizes may result in significantly improved query performance, especially when `text` and `image` or bulk data is involved.
If a packet size larger than the default is specified, then the target server must be configured to allow variable-length packet sizes. SAP ASE configuration parameters of interest in this case are:
 - `additional netmem`
 - `maximum network packet size`
 See the *System Administration Guide: Volume 1* for a complete explanation of these configuration parameters.
- `cis rpc handling` – determines whether Component Integration Services handles outbound RPC requests by default. When this is enabled using `sp_configure "cis rpc handling" 1`, all outbound RPCs are handled by Component Integration Services. When you use `sp_configure "cis rpc handling" 0`, the SAP ASE site handler is used. The thread cannot override it with `set cis_rpc_handling on`. If the global property is disabled, a thread can enable or disable the capability, as required.

The procedure for resetting configuration parameters is:

Procedure

1. Execute `sp_configure`, which updates the values field of the system table `master..sysconfigures`.
2. Restart the server if you have reset any of the static configuration parameters. The parameters listed below are dynamic.
 - `cis rpc handling`
 - `cis cursor rows`
 - `cis bulk insert batch size`

- `cis bulk insert array size`
- `cis packet size`

Related Information

[RPC handling and Component Integration Services \[page 44\]](#)

2.12.1.1 sysconfigures Table

The `master..sysconfigures` system table stores all configuration options.

It contains columns identifying the minimum and maximum values possible for each configuration parameter, as well as the configured value and run value for each parameter.

The `status` column in `sysconfigures` cannot be updated by the user. Status 1 means dynamic, indicating that new values for these configuration parameters take effect immediately. The rest of the configuration parameters (those with status 0) take effect only after the `reconfigure` command has been issued and the server restarted.

You can display the configuration parameters currently in use (run values) by executing `sp_configure` without giving it any parameters.

2.12.1.2 Changing the Configuration Parameters

`sp_configure` displays all the configuration values when it is used without an argument.

Context

When used with an option name and a value, the server resets the configuration value of that option in the system tables.

See the *System Administration Guide* for a complete discussion of `sp_configure` with syntax options.

Procedure

- To see the Component Integration Services options, enter:

```
sp_configure "Component Integration Services"
```

- To change the current value of a configuration parameter, execute `sp_configure` as follows:

```
sp_configure "<parameter>", <value>
```

2.12.2 Use Global Variables to Determine Status

Component Integration Services adds a number of global variables.

- `@@cis_rpc_handling`
- `@@transactional_rpc`
- `@@textptr_parameters`
- `@@stringsize`
- `@@bulkbatchsize` – contains the value of the current `cis bulk insert batch size` configured via `sp_configure`, or set through the `set bulk batch size` command.
- `@@bulkarraysize` – contains the value of the current `cis bulk insert array size` configured via `sp_configure` or set through the `set bulk array size` command.

These global variables show the current status of the corresponding configuration parameters. For instance, to see the status of `cis_rpc_handling`, issue the following command, which returns either 0 (off) or 1 (on):

```
select @@cis_rpc_handling
```

3 CIS Support for Connecting to an SAP HANA Server

CIS includes the native ODBC interface, which lets you connect directly to an SAP HANA server from SAP ASE. You must install the SAP HANA client package, which includes the ODBC drivers, on the same machine that is running SAP ASE.

See *Client Installation and Update Guide* in your SAP HANA database documentation.

3.1 Creating SAP HANA as an ODBC Data Source on Windows

The ODBC Administrator provides a central place for creating and managing ODBC data sources on the Windows platform.

Context

When using 32-bit ODBC drivers, use the 32-bit ODBC Administrator to manage data sources; for 64-bit clients, use the 64-bit ODBC Administrator.

Procedure

1. From the Windows Start window, select *Settings > Control Panel > Administrative Tools > Data Sources (ODBC)*.
2. Click *Add*.
3. Select *HDBODBC* from the list of drivers.
4. Click *Finish*.
5. In the ODBC Configuration window, enter the Data Source Name and Server:Port . For example, to connect to a server named "HANA_DB" on host "hanaserver" with address 157.133.66.75 using TCP/IP protocol and port 1870, select TCP/IP and enter:

```
host=hanaserver:1870
```

To use a host network address, enter:

```
host=157.133.66.75:1870
```

6. Click *Ok*.

3.2 Adding SAP HANA to the SAP ASE Interfaces File

You must add the HANA server entry to the SAP ASE interfaces file before you can connect to it.

Procedure

Using vi or a similar text editor, add an entry for the SAP HANA server to the interfaces file or sql.ini file.

This example adds an entry on UNIX for the HANA_DB server using port number 1870 and the libodbcHDB.so SAP HANA ODBC driver:

```
HANA_DB
query tcp ether 157.133.66.75 1870 libodbcHDB.so
```

This example adds an entry on Windows for the HANA_DB server using the HDBODBC (HDBODBC32 for Windows 32-bit) driver, which is registered in the Windows registry:

```
[HANA_DB]
Query=NLWNSCK,157.133.66.75,1870,HDBODBC
```

i Note

Because PCI Bridge uses HDBODBC as the default value, you need not include an HDBODBC entry in the interfaces file (sql.ini) for Windows.

3.3 Configuring the PCI Bridge and PCA/ODBC

SAP ASE uses Pluggable Component Interface (PCI) Bridge, which implements on-demand software dispatching, to load shared objects when it invokes a target function.

Context

ODBC Driver Manager is typically bootstrapped from the pluggable component adapter for ODBC (PCA/ODBC), which is configured with PCI Bridge. PCA/ODBC acts as a broker, managing service requests between the SAP ASE and the ODBC Driver Manager. PCA/ODBC forwards and controls requests in both directions— from the SAP ASE to ODBC Driver Manager, and vice versa.

ODBC Driver Manager component is independent of SAP ASE. You can change or upgrade these software components at any time without upgrading or rebuilding SAP ASE.

Note

You must enable PCI Bridge for PCA/ODBC before it can use ODBC Driver Manager. SAP ASE requires the `sybpcidb` database when you enable PCI Bridge. The `sybpcidb` contains all configuration data for PCI Bridge and its associated PCAs, such as PCA/ODBC.

Procedure

1. As system administrator, create the `sybpcidb` database using the `srvbuild` or `sybconfig` installation and configuration utilities, or by running the `${SYBASE}/${SYBASE_ASE}/scripts/installpcidb` script (see the installation guide).

2. Enable PCI Bridge:

```
sp_configure 'enable pci', 1
```

See *Reference Manual: Configuration Parameters*.

3. Restart SAP ASE.

4. Use `sp_odbcconfig` to configure these PCA/ODBC parameters:

- `pca_odbc_load_dir` – (UNIX only) sets the absolute path to the location for the ODBC drivers. The `pca_odbc_load_dir` makes an array of multiple locations. That is, if there are multiple ODBC drivers installed on the machine, `pca_odbc_load_dir` creates an array of them.

The default value for `pca_odbc_load_dir` is `/usr/sap/hdbclient`.

1. Change to the `sybpcidb` database:

```
use sybpcidb
```

2. Set the path for `pca_odbc_load_dir`:

```
sp_odbcconfig 'add', 'pca_odbc_load_dir', '<path_to_driver_location>'
```

3. Restart SAP ASE after making changes to `pca_odbc_load_dir`.

To remove drivers that are no longer used:

```
use sybpcidb
go
sp_odbcconfig 'delete', 'pca_odbc_load_dir', '<path_to_driver_location>'
```

Note

On Windows, ODBC drivers are administered and maintained by the system registry during driver component installation. See your Windows platform documentation and specific driver vendor documentation.

- `pca_odbc_root_dir` – UNIX systems store the ODBC data sources in the `odbc.ini` file. `pca_odbc_root_dir` indicates the path to the directory that contains `odbc.ini`. The default value for `pca_odbc_root_dir` is `$SYBASE`.

The format for an entry in `odbc.ini` for the HANA_DB server is:

```
[HANA_DB]
DRIVER = libodbcHDB.so
```

```
SERVERNODE = hanaserver:1870
```

i Note

On Windows, use the ODBC registry to configure the ODBC data sources. See your Windows documentation.

On UNIX systems, use `sp_odbcconfig` to set the value for `pca_odbc_root_dir`:

1. Change to the `sybpcidb` database:

```
use sybpcidb
```

2. Enter the new path:

```
sp_odbcconfig 'update', 'pca_odbc_root_dir', '<old_path>', '<new_path>'
```

where `<new_path>` indicates a relative location to `$SYBASE`. Preface `<new_path>` with a forward slash to indicate an absolute location to the root directory.

i Note

You must restart SAP ASE if ODBC Driver is already loaded under control of PCA/ODBC. SAP ASE consults the DSN records stored under `odbc.ini` only after you enable `pca_odbc_dsn_enabled`.

- `pca_odbc_dsn_enabled` – enables or disables the DSN mode for the PCA/ODBC subsystem. Values are:
 - 0 – disables DSN (the default). SAP ASE automatically gathers DSN information from the SAP directory control layer (DCL) from the interfaces file.
 - 1 – enables DSN. SAP ASE consults the DSN records stored under `odbc.ini`.
5. (UNIX only) Add `$SYBASE/DataAccess64/ODBC/dm/lib64` to the `LD_LIBRARY_PATH`.

3.4 Adding Server Class HANAODBC

SAP ASE includes the HANAODBC server class, which enables you to connect to the SAP HANA server using ODBC.

Context

You must install the HANA ODBC driver on the same machine that is running SAP ASE.

Procedure

1. To add the HANAODBC server class, execute:

```
sp_addserver <logical_server_name>, HANAODBC, <ODBC_DSN_name>
```

This example adds `big_hana_server` as the `Big_company` DSN:

```
sp_addserver big_hana_server, HANAODBC, Big_company
```

2. Use `sp_addexternlogin` to map the SAP ASE login account and password.

This example configures the system administrator to use the name 'HANA_login' with "HANA_login_password" to log in to the remote server named "big_hana_server":

```
sp_addexternlogin big_hana_server, sa, HANA_login, HANA_login_password
```

3.5 Datatype Mapping Between SAP ASE and HANA

When you create proxy tables, SAP ASE datatypes are mapped to corresponding HANA datatypes.

HANAODBC maps these SAP ASE datatypes to HANA datatypes when you use `create table` to create proxy tables:

SAP ASE Datatype	HANA Datatype
tinyint	tinyint
smallint	smallint
int	integer
integer	integer
bigint	bigint
unsigned smallint	integer
unsigned int	bigint
unsigned bigint	Not supported
numeric (precision, scale)	decimal (precision, scale)
decimal (precision, scale)	decimal (precision, scale)
float (precision)	float (precision)
double (precision)	float
real	real
smallmoney	decimal (10,4)
money	decimal (19,4)

SAP ASE Datatype	HANA Datatype
smalldatetime	timestamp
datetime	timestamp
date	date
time	time
char (n)	char (n)
varchar (n)	varchar (n)
unichar (n)	nchar (n)
univarchar (n)	nvarchar (n)
nchar (n)	char (n)
nvarchar (n)	varchar (n)
text	clob
unitext	nclob
binary (n)	binary (n)
varbinary (n)	varbinary (n)
image	blob
bit	tinyint
bigdatetime	timestamp
bigtime	Not supported

These datatypes are allowed when you map remote HANA columns to local proxy table columns:

Remote HANA Datatype	SAP ASE Datatype
tinyint	tinyint
smallint	smallint
int	integer
integer	integer
bigint	bigint
decimal (precision, scale)	decimal (precision, scale)
real	real
double	double precision
float (precision)	double precision
date	date
time	time
seconddate	datetime

Remote HANA Datatype	SAP ASE Datatype
timestamp	bigdatetime
varchar (n)	varchar (n)
nvarchar (n)	univarchar (n)
alphanum (n)	univarchar (n)
char (n)	char (n)
nchar (n)	unichar (n)
varbinary (n)	varbinary (n)
clob	text
binary (n)	binary (n)
blob	image
nclob	unitext

3.6 Restrictions

CIS support for HANA includes some restrictions.

- CIS for HANA is supported only in SAP ASE servers that are configured for threaded kernel mode. You cannot use CIS for HANA in SAP ASE servers that are configured for process mode.
- CIS for HANA does not support `update` or `delete` commands that include a `where current of <cursor_name>` clause.
- CIS for HANA does not support sensitive cursors.
- CIS for HANA does not support large object (LOB) parameters.
- You cannot issue queries that include text pointers to the remote HANA server because HANA LOBs do not support text pointers.
- If you issue `create index...desc` to create an index on a proxy table, the `desc` parameter is not generated for the `create index` statement. Instead, the index in the remote HANA table is created with an `asc` parameter.
- `alter table` statements to remote tables include restrictions. You cannot:
 - Include partition clauses
 - Include referential integrity constraints
 - Add constraints
 - Replace columns
 - Change locking schemes
 - Add columns with `check` conditions
 - Add columns with unique or primary key constraints
- A single `alter table` command cannot include more than one of these parameters:
 - `add <column_name>`
 - `modify <column_name>`

- `drop <column_name>`

- You must define a unique index on the base table for positioned `update` and `delete` statements, and the unique index must have at least one nonnullable column. SAP ASE issues an error message if this unique index is unavailable.
- You cannot use distributed transaction management (DTM) with two-phase commit on HANA proxy tables.
- You cannot include `save tran` and `rollback to savepoint` statements in transactions that involve HANA proxy tables.
- `rollback tran` operations issued on transactions involving HANA proxy tables roll back all statements up to the first statement in the current transaction. All statements after the `rollback tran` operation execute normally.
- You cannot issue `rollback trigger` on triggers that include HANA proxy tables.
- Search clauses can contain only the input parameters (and none of the resultant columns) for queries using proxy tables that are mapped to an external HANA procedure. For example, if you create this table:

```
create existing table rpc_table
(
    col1 int,
    col2 int
) external procedure at "odbchana..SYSTEM.MYPROC"
```

Then issue the following query; the search results are not narrowed, because the query specifies `col1`:

```
select * from rpc_table where col1 > 100
```

- You cannot use null as a value for input parameters for proxy tables that are mapped to an external HANA procedure. For example, this is not allowed:

```
select * from rpc_table where _p1 = null
```

- CIS for HANA does not support bulk operations.
- Implicit conversions may cause a loss of precision for millisecond values for SAP ASE `bigdatetime` and `bigtime` datatypes that are mapped to a HANA `timestamp` datatype. SAP ASE stores the millisecond portion of `bigdatetime` and `bigtime` values using 6 digits (for example `<hh:mm:ss.zzzzzz>`). However, HANA stores the millisecond portion of `bigdatetime` and `bigtime` values in 7 digits in its `timestamp` datatype. Transmitting values using ODBC is precise up to 3 digits. Implicitly converting values from ODBC to the SAP ASE internal format may further depreciate the precision, because the available conversion formats support only up to 3 digits.
- You must fully pad binary values when you specify them in a search clause. In the first select statement below, the value of `c1` (0x123) is not fully padded, and therefore returns no results. The second select statement uses a fully padded value (0x012300) for `c1`, and thereby returns a valid result:

```
1> create table cb(c1 binary(3)) at 'odbchana..SYSTEM.cb'
2> go
1> insert into cb values(0x123)
2> go
(1 row affected)
1> select * from cb where c1 = 0x123
2> go
c 1
-----
(0 rows affected)
3> select * from cb where c1 = 0x012300
c1
-----
0x012300
```

(1 row affected)

Restrictions for quoted identifiers when using CIS for HANA include:

- Object names in object definitions included with DDL statements must include at least one uppercase letter.
- Quoted identifiers for DML statements must include at least one uppercase letter.
- CIS for HANA raises an error for DDL statements that include an object definition that contains only lowercase characters. For example, this raises an error: `<hanaserver>..<schema>.table1`.
- SAP ASE creates object names in proxy tables without quotes. Consequently, SAP ASE cannot determine whether to add quotes to a quoted identifier when executing a DML statement if the remote quoted identifier contains only lowercase characters (for example `column1`). That is, SAP ASE creates the object name in the same way for `column1` as for `COLUMN1`, because SAP ASE defaults to lowercase. As a workaround, use trace flag 11242, which quotes all lowercase identifiers are quoted.
- CIS for HANA does not support bracketed identifiers (that is, [and]).

4 SQL Reference

Component Integration Services supports a number of classes.

Each server class has a set of unique characteristics that system administrators and programmers need to know about in order to configure the server for remote data access. These properties are:

- Types of servers that each server class supports
- Datatype conversions specific to the server class
- Restrictions on Transact-SQL statements that apply to the server class

4.1 dbcc Commands

All `dbcc` commands used by Component Integration Services are available with a single `dbcc` entry point.

The syntax for `dbcc cis` is:

```
dbcc cis ("s<subcommand>" [, <vararg1>, <vararg2>...])
```

If Component Integration Services is not configured or loaded, the command results in a runtime error.

The use of the `dbcc cis` command is unrestricted.

4.1.1 dbcc Options

Component Integration Services includes a number of `dbcc` parameters.

- `remcon` – displays a list of all remote connections made by all Component Integration Services clients. It takes no arguments.
- `srvdes` – returns a formatted list of all in-memory `SRVDES` structures (if no argument is provided). If an argument is provided, this command syncs the in-memory version of a `SRVDES` with information found in `sys.servers`. The command takes an optional server ID (`<srv_id>`) argument as follows:

```
srvdes, [<srv_id>]
```

- `showcaps` – displays a list of all capabilities for `<server_name>` by capability name, ID, and value. The command is as follows:

```
dbcc cis("showcaps", "<server_name>")
```

When you run the above command with a valid `<server_name>`, the output returns `CapValue` details for the server, along with basic information. `CapValue` indicates the internal capability value of the associated

capability on the server. The following example displays a list of capabilities for the server JARVIS_AI, arranged by capability ID, capability name, and the associated capability values:

```
dbcc cis("showcaps","IWS_IQ")
SrvID  SrvName          CapID  CapName          CapValue
-----
    20  JARVIS_AI          555  SQL syntax       1
    20  JARVIS_AI          666  join handling    6
```

In the above example, the CapValue for the join handling capability is 6. Valid values for the join handling capability, along with descriptions for each are described in the table below:

Table 6: Join handling capabilities

Join name	Value	Description
CAP_JOIN_NONE	0	No joins possible
CAP_JOIN_PARTIAL	1	Joins possible, but not inner/outer joins
CAP_JOIN_TSQL	2	TSQL inner/outer joins
CAP_JOIN_ORACLE	3	Oracle type inner/outer joins
CAP_JOIN_ANSI	4	Support ANSI outer joins
CAP_JOIN_DB2_ANSI	5	Support DB2 ANSI outer joins
CAP_JOIN_TSQL_ANSI	6	TSQL and ANSI inner/outer joins
CAP_JOIN_ORACLE_ANSI	7	Oracle and ANSI outer joins

4.1.2 Trace Flags

The `dbcc traceon` option allows the system administrator to turn on trace flags within Component Integration Services.

Trace flags enable the logging of certain events when they occur within Component Integration Services. Each trace flag is uniquely identified by a number. Some are global to Component Integration Services, while others are `spid`-based and affect only the user who enabled the trace flag. `dbcc traceoff` turns off trace flags.

The syntax is:

```
dbcc traceon (<traceflag> [, <traceflag>...])
```

Trace Flag	Description
11201	Logs client connect events, disconnect events, and attention events. (global). Events sent to error log only.

Trace Flag	Description
11202	Logs client language, cursor declare, dynamic prepare, and dynamic execute-immediate text (global). Text sent to error log only.
11203	Logs client RPC events (global). Events sent to error log only.
11204	Logs all messages routed to client (global). Messages sent to only the error log.
11205	Logs all interaction with remote server (global). Interaction sent to error log only.
11206	Logs file/directory processing steps. (global).
11207	Logs text and image processing. (global).
11208	Prevents the <code>create index</code> and <code>drop index</code> statements from being transmitted to a remote server. <code>sysindexes</code> is updated anyway. (spid).
11209	Instructs <code>update statistics</code> to obtain only row counts rather than complete distribution statistics, from a remote table. (spid).
11211	Prevents the <code>drop table</code> syntax from being forwarded to remote servers if the table was created using the <code>create table at location</code> syntax.
11212	Prevents escape on underscores ("_") in table names. (spid).
11213	Prevents generation of column and table constraints. (spid).
11214	Disables Component Integration Services recovery at start-up. (global).
11216	Disables quickpass. (spid).
11217	Disables quickpass. (global).
11218	Makes cursors involving Component Integration Services tables updateable by default.
11220	Disables constraint checking of remote tables on the local server. This avoids duplicate checking. Setting this trace flag on ensures that queries are not rejected by the quickpass mode because of constraints. (spid).
11221	Disables <code>alter table</code> commands to the remote server when ON. This allows users to modify <code><type></code> , <code><length></code> , and <code><nullability></code> of columns in a local table without changing columns in the remote table. Use trace flag 11221 with caution. It may lead to tables that are "out of sync." (spid).
11223	Disables proxy table index creation during <code>create existing table</code> or <code>create proxy_table</code> command execution. If this flag is set on, no index metadata is imported from the remote site referenced by the proxy table, and no indexes for the proxy table are created. This trace flag should be used with care and turned off when no longer necessary. (global).
11228	Disables proxy table mapping to RPCs.

Trace Flag	Description
11229	Instructs Component Integration Services to use pre-SAP ASE version 12.5.3 methods of gathering statistics data.
11277	Send messages to client containing text of SQL <code>select/insert/delete/update</code> that is being sent to a remote server.
11299	Allows connection information to be logged when a connection to a remote server fails.

4.2 Functions

When a SQL statement such as a `select`, `insert`, `delete`, or `update` contains a built-in function, Component Integration Services must determine if the function can be forwarded to the remote server, or if it must be evaluated within the local server using remote data.

Functions are only sent to a remote server if the statement containing them can be handled by quickpass mode.

In the tables shown below, support for function by server class is indicated by a 'Y'; 'N' indicates no support is provided, and 'C' indicates support for it is determined by capabilities of the underlying DBMS (often the case for DirectConnects).

4.2.1 Aggregate Functions

The aggregate functions generate summary values that appear as new columns in the query results.

Table 7: Server Class Support for Aggregate Functions

Function	ASE	ASA	ASIQ	dir_con
avg	Y	Y	Y	C
count	Y	Y	Y	C
max	Y	Y	Y	C
min	Y	Y	Y	C
sum	Y	Y	Y	C
count_big	Y	N	N	N

4.2.2 Datatype Conversion Functions

Datatype conversion functions change expressions from one datatype to another and specify new display formats for date/time information.

Table 8: Server Class Support for Datatype Conversion Functions

Function	ASE	ASA	ASIQ	dir_con
<code>convert()</code>	Y	Y	Y	C
<code>inttohex()</code>	Y	N	N	N
<code>hextoint()</code>	Y	N	N	N
<code>biginttohex()</code>	Y	N	N	N
<code>hextobigint()</code>	Y	N	N	N
<code>bintostr</code>	Y	N	N	N
<code>strtobin</code>	Y	N	N	N

4.2.3 Date Functions

The date functions manipulate values of the datatypes `datetime` or `smalldatetime`.

The `getdate()` function is always expanded by the local server; the presence of this builtin function does not cause a query to be eliminated from quickpass mode optimizations, however.

Table 9: Server Class Support for Date Functions

Function	ASE	ASA	ASIQ	dir_con
<code>dateadd</code>	Y	Y	Y	C
<code>datediff</code>	Y	Y	Y	C
<code>datetime</code>	Y	Y	N	C
<code>datepart</code>	Y	Y	Y	C

4.2.4 Mathematical Functions

Mathematical functions return values commonly needed for operations on mathematical data. Mathematical function names are not keywords.

Each function also accepts arguments that can be implicitly converted to the specified type. For example, functions that accept approximate numeric types also accept integer types. SAP ASE automatically converts the argument to the desired type

Table 10: Server Class Support for Mathematical Functions

Function	ASE	ASA	ASIQ	dir_con
abs	Y	Y	Y	C
acos	Y	Y	N	C
asin	Y	Y	N	C
atan	Y	Y	N	C
atn2	Y	Y	N	C
ceiling	Y	Y	Y	C
cos	Y	Y	N	C
cot	Y	Y	N	C
degrees	Y	Y	N	C
exp	Y	Y	N	C
floor	Y	Y	Y	C
log	Y	Y	N	C
log10	Y	Y	N	C
pi	Y	Y	N	C
power	Y	Y	N	C
radians	Y	Y	N	C
rand	Y	Y	Y	C
round	Y	Y	N	C
sign	Y	Y	N	C
sin	Y	Y	N	C

Function	ASE	ASA	ASIQ	dir_con
sqrt	Y	Y	Y	C
tan	Y	Y	N	C

4.2.5 Security Functions

Security functions return security-related information.

Table 11: Server Class Support for Security Functions

Function	ASE	ASA	ASIQ	dir_con
ic_sec_service_on()	N	N	N	N
show_sec_services()	N	N	N	N

4.2.6 String Functions

String function operate on binary data, character strings, and expressions.

Table 12: Server Class Support for String Functions

Function	ASE	ASA	ASIQ	dir_con
ascii	Y	Y	N	C
char	Y	Y	N	C
charindex	Y	Y	N	C
char_lengt	Y	Y	N	C
difference	Y	Y	Y	C
lower	Y	Y	Y	C
ltrim	Y	Y	Y	C
patindex	N	N	N	N
replicate	Y	Y	N	C
reverse	Y	N	N	Y

Function	ASE	ASA	ASIQ	dir_con
right	Y	Y	Y	C
rtrim	Y	Y	Y	C
soundex	Y	N	Y	C
space	Y	Y	N	C
str	Y	Y	N	C
stuff	Y	Y	N	C
substring	Y	Y	Y	C
upper	Y	Y	Y	C

4.2.7 System Functions

System functions return special information from the database.

Table 13: Server Class Support for System Functions

Function	ASE	ASA	ASIQ	dir_con
col_length	Y	Y	N	C
col_name	Y	Y	N	C
curunreservedpgs	N	N	N	N
data_pgs	N	N	N	N
datalength	Y	Y	N	C
db_id	N	N	N	N
db_name	N	N	N	N
getdate	Y	N	N	N
getutcdate	Y	N	N	N
host_id	N	N	N	N
host_name	N	N	N	N

Function	ASE	ASA	ASIQ	dir_con
index_col	N	N	N	N
isnull	Y	Y	N	N
lct_admin	N	N	N	N
mut_excl_roles	N	N	N	N
object_id	N	N	N	N
object_name	N	N	N	N
proc_role	N	N	N	N
ptn_data_pgs	N	N	N	N
reserved_pgs	N	N	N	N
role_contain	N	N	N	N
role_id	N	N	N	N
role_name	N	N	N	N
rowcnt	N	N	N	N
show_role	N	N	N	N
suser_id	N	Y	Y	N
suser_name	N	Y	Y	N
tsequal	Y	Y	N	N
used_pgs	N	N	N	N
user	Y	Y	Y	N
user_id	Y	Y	Y	N
user_name	Y	Y	Y	N
valid_name	N	N	N	N
valid_user	N	N	N	N

4.2.8 Text and Image Functions

Text and image functions operate on `text` and `image` data.

Table 14: Server Class Support for text and image Functions

Function	ASE	ASA	ASIQ	dir_con
<code>textptr()</code>	Y	Y	N	C
<code>textvalid()</code>	Y	Y	N	C

4.3 Transact-SQL Commands

If Component Integration Services does not pass all of a command's syntax to a remote server (such as all clauses of a `select` statement), the syntax that is passed along is described for each server class.

Each command has several sections that describe it:

- Description – contains a brief description of the command.
- Syntax – contains a description of the full Transact-SQL syntax of the command.
- Usage – contains a general, server class-independent description of handling by Component Integration Services.
- Server class `ASEnterprise` – contains a description of handling specific to server class `ASEnterprise`. This includes syntax that is forwarded to a remote server of class `ASEnterprise`.
- Server class `ASAnywhere` – contains a description of handling specific to server class `ASAnywhere`. This includes syntax that is forwarded to a remote server of class `ASAnywhere`.
- Server class `ASIQ` – contains a description of handling specific to server class `ASIQ`. This includes syntax that is forwarded to a remote server of class `ASIQ`.
- Server class `direct_connect` – contains a description of handling specific to server class `direct_connect`. This includes syntax that is forwarded to a remote server of class `direct_connect`. In this release, all comments that apply to server class `direct_connect`, also apply to server class `sds`.

4.3.1 alter table

Component Integration Services includes changes to the `alter table` command.

When the server receives the `alter table` command, it passes the command to an appropriate access method if:

- The object on which the command is to operate has been associated with a remote or external storage location.
- The command consists of an `add column` request. Requests to add or drop constraints are not passed to the access methods; instead, they are handled locally.

`alter table` is passed to remote servers as a language request.

See `alter table` in *Reference Manual: Commands*.

These sections describe the `alter table` command in different server classes.

Server Class ASEnterprise

Component Integration Services forwards the following syntax to a server configured as class ASEnterprise:

```
alter table [database.[owner].]table_name
  {add column_name datatype [{identity | null}]
  {[, next_column]}...}
| [drop column_name [, column_name]]
| modify column_name [data_type] [NULL] |
  [not null]] [, column_name]
```

- When a user adds a column with the `alter table` command, Component Integration Services passes the datatype of each column to the remote server without type name conversions.
- For ASEnterprise class servers only, the `lock` clause is also forwarded, if contained in the original query, if the version of SAP ASE is 11.9.2 or later.

Server Class ASAnywhere

Handling of `alter table` by servers in this class is the same as for ASEnterprise servers.

Server Class ASIQ

- Handling of `alter table` by servers in this class is the same as for ASEnterprise servers.
- `text` and `image` datatypes are fully supported by server class ASIQ.

Server Class direct_connect

- Component Integration Services forwards the following syntax to a remote server configured as class `direct_connect`:

```
alter table [<database>.<owner>].<table_name>
  add <column_name> <datatype> [{identity | null}]
  {[, <next_column>]}...
```

- Although Component Integration Services requests a capabilities response from a server with class `direct_connect`, support for `alter table` is not optional. Component Integration Services forwards `alter table` to the remote server regardless of the capabilities response.

- The behavior of the server with class `direct_connect` is database dependent. The Transact-SQL syntax is forwarded, and errors may or may not be raised, depending on the ability of the remote database to handle this syntax.
- Server class `direct_connect` does not support `bigint`, `unsigned tinyint`, `unsigned smallint`, `unsigned int`, `unsigned bigint`.
- If the syntax capability of the remote server indicates Transact-SQL, SAP ASE datatypes are sent to the remote server. If the syntax capability indicates DB2 SQL, DB2 datatypes are sent. Direct Connect does not support `bigint`, `unsigned tinyint`, `unsigned smallint`, `unsigned int`, `unsigned bigint`.

Table 15: DirectConnect datatype conversions for alter table

SAP ASE Datatype	DirectConnect Default Datatype
<code>binary(<n>)</code>	<code>binary(<n>)</code>
<code>bit</code>	<code>bit</code>
<code>char</code>	<code>char</code>
<code>date</code>	<code>date</code>
<code>datetime</code>	<code>datetime</code>
<code>decimal(<p>, <s>)</code>	<code>decimal(<p>, <s>)</code>
<code>float</code>	<code>float</code>
<code>image</code>	<code>image</code>
<code>int</code>	<code>int</code>
<code>money</code>	<code>money</code>
<code>numeric(<p>, <s>)</code>	<code>numeric(<p>, <s>)</code>
<code>nchar(<n>)</code>	<code>nchar(<n>)</code>
<code>nvarchar(<n>)</code>	<code>nvarchar(<n>)</code>
<code>real</code>	<code>real</code>
<code>smalldatetime</code>	<code>smalldatetime</code>
<code>smallint</code>	<code>smallint</code>
<code>smallmoney</code>	<code>smallmoney</code>
<code>time</code>	<code>time</code>
<code>timestamp</code>	<code>timestamp</code>

SAP ASE Datatype	DirectConnect Default Datatype
<code>tinyint</code>	<code>tinyint</code>
<code>text</code>	<code>text</code>
<code>unichar</code>	<code>unichar</code>
<code>unitext</code>	<code>unitext</code>
<code>varbinary(<n>)</code>	<code>varbinary(<n>)</code>
<code>varchar(<n>)</code>	<code>varchar(<n>)</code>

4.3.2 case

Component Integration Services includes changes to the `case` expression.

See *case Expression* in the *Transact-SQL Users Guide*.

Server Class ASEnterprise

The presence of a `case` expression in the original query syntax will not cause the query optimizer to reject quickpass mode.

Server Class ASAnywhere

The presence of a `case` expression in the original query syntax will not cause the query optimizer to reject quickpass mode.

Server Class ASIQ

The ability to handle `case` expressions is not set for servers in this class. When a SQL statement containing a `case` expression is optimized, the presence of the `case` expression causes the Component Integration Services quickpass optimization to reject the statement. When this happens, the `case` expression must be evaluated by the local SAP ASE after retrieving data from the remote server.

Server Class `direct_connect`

The ability to handle `case` expressions is determined by the result set from the RPC `sp_capabilities`. If `direct_connect` indicates that it can handle case expressions, then Component Integration Services forwards them to the `direct_connect` when quickpass mode is used to handle the query.

4.3.3 `connect to...disconnect`

Component Integration Services includes changes to the `connect to...disconnect` command.

See `connect to...disconnect` in *Reference Manual: Commands*.

Server Class `ASEnterprise`

When `disconnect` is issued, Component Integration Services forwards `disconnect` to the remote server, to take it out of passthrough mode. If not in passthrough mode, syntax errors may occur, but they are ignored by Component Integration Services and not forwarded to the client.

Server Class `ASAnywhere`

No interaction occurs with `ASAnywhere` when `connect` or `disconnect` are issued.

Server Class `ASIQ`

No interaction occurs with `ASIQ` when `connect` or `disconnect` are issued.

Server Class `direct_connect`

When `connect` is issued using a server in class `direct_connect`, the `direct_connect` is sent an RPC:

```
sp_thread_props "passthru mode", 1
```

When `disconnect` is issued, and the server for which a passthrough-mode connection has been established is a `direct_connect`, the `direct_connect` is sent an RPC:

```
sp_thread_props "passthru mode", 0
```

4.3.4 create existing table

Component Integration Services includes changes to the `connect to...disconnect` command.

When a `create existing table` command is received, it is interpreted as a request to import metadata from the remote or external location of the object for updating system catalogs. Importing this metadata is performed by means of three RPCs sent to the remote server with which the object has been associated:

- `sp_tables` – verifies that the remote object actually exists.
- `sp_columns` – obtains column attributes of the remote object for comparison with those defined in `create existing table`.
- `sp_statistics` – obtains index information to update the local system table, `sysindexes`.

See `create existing table` in *Reference Manual: Commands*.

Server Class ASEnterprise

This table describes the allowable datatypes that can be used when mapping remote SAP ASE columns to local proxy table columns:

Table 16: SAP ASE datatype conversions for create existing table

Remote SAP ASE Data-type	Allowable SAP ASE Datatypes
<code>binary(<n>)</code>	<code>image</code> , <code>binary(<n>)</code> , and <code>varbinary(<n>)</code> ; if not <code>image</code> , the length must match
<code>bit</code>	<code>bit</code> , <code>decimal</code> , <code>float</code> , <code>int</code> , <code>money</code> , <code>numeric</code> , <code>real</code> , <code>smallint</code> , <code>smallmoney</code> , and <code>tinyint</code>
<code>char(<n>)</code>	<code>text</code> , <code>nchar(<n>)</code> , <code>nvarchar(<n>)</code> , <code>char(<n>)</code> , <code>varchar(<n>)</code> , <code>unichar</code> , <code>univarchar</code> ; if not <code>text</code> , the length must match
<code>datetime</code>	<code>datetime</code> , <code>smalldatetime</code> , <code>char</code> , and <code>varchar</code>
<code>decimal(<p>, <s>)</code>	<code>bit</code> , <code>decimal</code> , <code>float</code> , <code>int</code> , <code>money</code> , <code>numeric</code> , <code>real</code> , <code>smallint</code> , <code>smallmoney</code> , and <code>tinyint</code>
<code>float</code>	<code>bit</code> , <code>decimal</code> , <code>float</code> , <code>int</code> , <code>money</code> , <code>numeric</code> , <code>real</code> , <code>smallint</code> , <code>smallmoney</code> , and <code>tinyint</code>
<code>image</code>	<code>image</code>
<code>int</code>	<code>bit</code> , <code>decimal</code> , <code>float</code> , <code>int</code> , <code>money</code> , <code>numeric</code> , <code>real</code> , <code>smallint</code> , <code>smallmoney</code> , and <code>tinyint</code>
<code>money</code>	<code>bit</code> , <code>decimal</code> , <code>float</code> , <code>int</code> , <code>money</code> , <code>numeric</code> , <code>real</code> , <code>smallint</code> , <code>smallmoney</code> , and <code>tinyint</code>

**Remote SAP ASE Data-
type**

Allowable SAP ASE Datatypes

nchar(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>); if not text, the length must match
numeric(<p>, <s>)	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
nvarchar(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>), unichar, univarchar; if not text, the length must match
real	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
smalldatetime	datetime, smalldatetime, char, and varchar
smallint	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
smallmoney	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
text	text, unitext
timestamp	timestamp
tinyint	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
unichar	char, varchar, unichar, univarchar, text, datetime, and smalldatetime
univarchar	char, varchar, unichar, univarchar, text, datetime, and smalldatetime
unitext	unitext
varbinary(<n>)	image, binary(<n>), and varbinary(<n>); if not image, the length must match
varchar(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>), unichar, univarchar; if not text, the length must match
date	
time	
bigint	Implicit: binary, varbinary, bit, tinyint, smallint, int, decimal, numeric, float, real, money, smallmoney Explicit: char, varchar, unichar, univarchar

**Remote SAP ASE Data-
type**

Allowable SAP ASE Datatypes

unsigned tinyint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
unsigned smallint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
unsigned int	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
unsigned bigint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime

i Note

Component Integration Services only supports `unitext` with SAP ASE version 15.0 and higher.

Server Class ASAnywhere

This table describes the allowable datatypes that can be used when mapping remote SAP ASE columns to local proxy table columns:

Table 17: SQL Anywhere datatype conversions for create existing table

Remote SQL Anywhere Data-type	Allowable SQL Anywhere Datatypes
binary(<n>)	image, binary(<n>), and varbinary(<n>); if not image, the length must match
bit	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
char(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>), unichar, univarchar; if not text, the length must match
datetime	datetime and smalldatetime
decimal(<p>, <s>)	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
float	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
image	image
int	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
money	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
numeric(<p>, <s>)	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
real	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
smalldatetime	datetime and smalldatetime
smallint	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
smallmoney	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
text	text

Remote SQL Anywhere Data-type**Allowable SQL Anywhere Datatypes**

timestamp	timestamp
tinyint	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
varbinary(<n>)	image, binary(<n>), and varbinary(<n>), unichar, unitext, univarchar; if not image, the length must match
varchar(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>), unichar, univarchar; if not text, the length must match
date	
time	
bigint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
unsigned tinyint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
unsigned smallint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
unsigned int	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime

Remote SQL Anywhere Data-type

Allowable SQL Anywhere Datatypes

unsigned bigint	Implicit: binary, varbinary, bit, tinyint, smallint, unsigned smallint, int, unsigned int, bigint, unsigned bigint, decimal, numeric, float, real money, smallmoney Explicit: char, varchar, unichar, univarchar Unsupported: text, image, date, time, datetime, smalldatetime
nchar (<n>)	text, nchar (<n>), nvarchar (<n>), char (<n>), varchar (<n>), unichar, univarchar; if not text, the length must match
nvarchar (<n>)	text, nchar (<n>), nvarchar (<n>), char (<n>), varchar (<n>), unichar, univarchar; if not text, the length must match

Server Class ASIQ

- text and image datatypes are supported by ASIQ version 12.6 and requires a license.
- Behavior is the same as for server class ASAnywhere.

Server Class direct_connect

- The RPC `sp_columns` queries the datatypes of the columns in the existing table.
- Local column datatypes do not need to be identical to remote column datatypes, but they must be convertible as shown in the table below. If not, a column type error is raised and the command is aborted.

Table 18: DirectConnect datatype conversions for create existing table

DirectConnect Data-type	Allowable SAP ASE Datatypes
binary(<n>)	image, binary(<n>), varbinary(<n>); if the length does not match, the command is aborted
binary(16)	timestamp
bit	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
char(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>) and varchar(<n>), unichar, univarchar; if the length does not match, the command is aborted

DirectConnect Data-type	Allowable SAP ASE Datatypes
datetime	datetime, smalldatetime
decimal(<p>,<s>)	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
float	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
image	image
int	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
money	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
nchar(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>) and varchar(<n>), unichar, univarchar; if the length does not match, the command is aborted
numeric(<p>,<s>)	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
nvarchar(<n>)	text, nchar(<n>), nvarchar(<n>), char(<n>) and varchar(<n>), unichar, univarchar; if the length does not match, the command is aborted
real	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
smalldatetime	datetime and smalldatetime
smallint	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
smallmoney	bit, decimal, float, int, money, numeric, real, smallint, smallmoney, and tinyint
text	text
timestamp	timestamp, binary(8), varbinary(8)
unichar	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>), unichar, univarchar; if not text, the length must match
univarchar	text, nchar(<n>), nvarchar(<n>), char(<n>), varchar(<n>), unichar, univarchar; if not text, the length must match

DirectConnect Data-type	Allowable SAP ASE Datatypes
date	
time	
bigint	UDB and DC/Microsoft support bigint.

- Datatype information is passed in the CS_DATAFMT structure associated with the parameter. The following fields of the structure contain datatype information:
 - `<datatype>` – the CS_Library datatype representing the SAP ASE datatype. For example, CS_INT_TYPE.
 - `<usertype>` – the native DBMS datatype. `sp_columns` passes this datatype back to Component Integration Services during a `create existing table` command as part of its result set (see `sp_columns` in the *Reference Manual*). SAP ASE returns this datatype in the `<usertype >` field of parameters to assist the DirectConnect in datatype conversions.

4.3.5 create index

Component Integration Services includes changes to the `create index` command.

When the server receives the `create index` command, it passes the command to an appropriate access method, if the object on which the command is to operate has been associated with a remote or external storage location.

The command is reconstructed using a syntax appropriate for the class and is passed to the remote server for execution.

`create index` is passed to remote servers as a language request.

See also `create index` in *Reference Manual: Commands*.

Server Class ASEnterprise

Component Integration Services forwards everything except the `on <segment_name>` clause to the remote server.

Server Class ASAnywhere

Component Integration Services forwards everything except the `on <segment_name>` clause to the remote server.

Server Class ASIQ

Component Integration Services forwards everything except the `on <segment_name>` clause to the remote server.

Server Class `direct_connect`

When the language capability is set to "Transact-SQL", Component Integration Services forwards all syntax except the `max_rows_per_page` and `on <segment_name>` clauses to the remote server.

4.3.6 create table

Component Integration Services includes changes to the `create table` command.

When the server receives a `create table` command, the command is interpreted as a request for new table creation. The server invokes the access method appropriate for the server class of the table that is to be created. If it is remote, the table is created. If this command is successful, system catalogs are updated, and the object appears to clients as a local table in the database in which it was created.

`create table` is reconstructed in a syntax that is appropriate if the server class. For example, if the server class is `direct_connect` and the remote server is DB2, the command is reconstructed using SQL Anywhere syntax before being passed to the remote server. Datatype conversions are made for datatypes that are unique to the SAP ASE environment.

Some server classes have restrictions on what datatypes can and cannot be supported.

`create table` is passed to remote servers as a language request.

See also `create table` in *Reference Manual: Commands*.

Server Class ASEnterprise

Component Integration Services passes the datatype of each column to the remote server without conversion.

Server Class ASAnywhere

Component Integration Services passes the datatype of each column to the remote server without conversion.

Server Class ASIQ

Component Integration Services passes the datatype of each column to the remote server without conversion.

Server Class direct_connect

- Component Integration Services reconstructs `create table` and passes commands to the targeted DirectConnect. The gateway transforms the commands into a form that the underlying DBMS recognizes.
- Direct Connect does not support `bigint`, `unsigned tinyint`, `unsigned smallint`, `unsigned int`, `unsigned bigint`.
- SAP ASE datatypes are converted to either the DirectConnect syntax mode datatypes shown below.

Table 19: DirectConnect datatype conversions for create table

SAP ASE Datatype	DirectConnect Default Datatype
<code>binary<n></code>	<code>binary<n></code>
<code>bit</code>	<code>bit</code>
<code>char</code>	<code>char</code>
<code>datetime</code>	<code>datetime</code>
<code>decimal<p, s></code>	<code>decimal<p, s></code>
<code>float</code>	<code>float</code>
<code>image</code>	<code>image</code>
<code>int</code>	<code>int</code>
<code>money</code>	<code>money</code>
<code>numeric<p, s></code>	<code>numeric<p, s></code>
<code>nchar(n)</code>	<code>nchar(n)</code>
<code>nvarchar(n)</code>	<code>nvarchar(n)</code>
<code>real</code>	<code>real</code>
<code>smalldatetime</code>	<code>smalldatetime</code>
<code>smallint</code>	<code>smallint</code>
<code>smallmoney</code>	<code>smallmoney</code>

SAP ASE Datatype	DirectConnect Default Datatype
timestamp	timestamp
tinyint	tinyint
text	text
unichar(n)	unichar
univarchar(n)	char<(n)> for bit data
varbinary<(n)>	varbinary<(n)>
varchar<(n)>	varchar<(n)>
date	
time	
bigint	UDB and DC/Microsoft support bigint

4.3.7 delete

Component Integration Services includes changes to the `delete` command.

See `delete` in *Reference Manual: Commands*.

Server Class ASEnterprise

If Component Integration Services cannot forward the original query without alteration, it performs the delete using method 2.

Server Class ASAnywhere

If Component Integration Services cannot forward the original query without alteration, it performs the delete using method 2.

Server Class ASIQ

If Component Integration Services cannot forward the original query without alteration, you get an error because ASIQ does not support updatable cursors.

Server Class `direct_connect`

- The syntax forwarded to servers of class `direct_connect` is dependent on the capabilities negotiation, which occurs when Component Integration Services first connects to the remote DirectConnect. Examples of negotiable capabilities include: subquery support, `group by` support, and built-in support.
- Component Integration Services passes data values as parameters to either a cursor or a dynamic SQL statement. Language statements can also be used if the DirectConnect supports it. The parameters are in the datatype native to SAP ASE and must be converted by the DirectConnect into formats appropriate for the target DBMS.

4.3.8 drop index

Component Integration Services includes changes to the `drop index` command.

When the server receives the `drop index` command, it passes the command to an appropriate access method, if the object on which the command is to operate has been associated with a remote or external storage location.

`drop index` is reconstructed using a syntax appropriate for the class and is passed to the remote server for execution.

This command is passed to remote servers as a language request.

See `drop index` in *Reference Manual: Commands*.

Server Class ASEnterprise

Component Integration Services forwards the following `drop index` syntax to a remote server configured as class ASEnterprise:

```
drop index <table_name.index_name>
```

Component Integration Services precedes this statement with a `use database` command since the `drop index` syntax does not allow you to specify the database name.

Server Class ASAnywhere

Component Integration Services forwards the following `drop index` syntax to a remote server configured as class `ASAnywhere`:

```
drop index <table_name.index_name>
```

Component Integration Services precedes this statement with a `use database` command since the `drop index` syntax does not allow you to specify the database name.

Server Class ASIQ

Component Integration Services forwards the following `drop index` syntax to a remote server configured as class `ASIQ`:

```
drop index <table_name.index_name>
```

Component Integration Services precedes this statement with a `use database` command since the `drop index` syntax does not allow you to specify the database name.

Server Class direct_connect

Component Integration Services forwards the following `drop index` syntax to a remote server configured as class `direct_connect`:

```
drop index <table_name.index_name><e>
```

4.3.9 fetch

Component Integration Services includes changes to the `fetch` command.

See `fetch` in *Reference Manual: Commands*.

Server Class ASEnterprise

If the cursor is read only, Component Integration Services sends a language request to the remote server when the first `fetch` is received after the cursor is opened. Otherwise, Component Integration Services declares a cursor to the remote server by means of Client-Library.

Server Class ASAnywhere

Handling of the `fetch` statement is the same as for `ASEnterprise`.

Server Class ASIQ

Component Integration Services sends a language request to the remote server when the first `fetch` is requested after the cursor is opened.

Server Class direct_connect

Component Integration Services treats servers in this class the same as servers in `ASEnterprise`.

4.3.10 insert

Component Integration Services includes changes to the `insert` command.

See `insert` in *Reference Manual: Commands*.

Server Class ASEnterprise

- `insert` commands using the `values` keyword are fully supported.
- `insert` commands using a `select` command are supported for all datatypes except `text` and `image`. `text` and `image` columns are only supported when they contain null values.
- If all `insert` and `select` tables reside on the same remote server, the entire statement is forwarded to the remote server for execution. This is referred to as quickpass mode. Quickpass mode is not used if `select` does not conform to all the quickpass rules for a `select` command.
- If the `select` tables reside on one remote server, and the `insert` table resides on a different server, Component Integration Services selects each row from the source tables, and inserts the row into the target table.
- You cannot `insert` into a computed column.

Server Class ASAnywhere

Handling of the `insert` statement is the same as for `ASEnterprise`.

Server Class ASIQ

Handling of the insert statement is the same as for ASEnterprise.

Server Class direct_connect

- `insert` commands using the `values` keyword are fully supported.
- `insert` commands using a `select` command are fully supported, but the table must have a unique index if the table has `text` or `image` columns. When using `insert` with a `select` command, the entire command is sent to the remote server if:
 - All tables referenced in the command reside on the remote server.
 - The capability's response from the DirectConnect indicates that `insert-select` commands are supported.
 - If you use the TopN feature, you must have an `order by` clause.

If both conditions are not met, Component Integration Services selects each row from the source tables, and inserts the row into the target table.

- Component Integration Services passes data values as parameters to either a cursor or a dynamic SQL statement. Language statements can also be used if the DirectConnect supports it. The parameters are in the datatype native to SAP ASE and must be converted by the DirectConnect into formats appropriate for the target DBMS.

4.3.11 readtext

Component Integration Services includes changes to the `readtext` command.

See `readtext` in *Reference Manual: Commands*.

Server Class ASEnterprise

Component Integration Services forwards the following syntax to the remote server when the underlying table is a proxy table:

```
readtext [[<database>.<owner>.<table_name>.<column_name><e
t><ext_pointer> <offset> <size><e
>[using {chars | characters}]
```

Server Class ASAnywhere

Handling of the `readtext` statement is the same as for ASEnterprise.

Server Class ASIQ

Handling of the `readtext` statement is the same as for `ASEnterprise`.

Server Class `direct_connect`

- If the `DirectConnect` does not support text pointers, `readtext` cannot be sent and its use results in errors.
- If the `DirectConnect` does support text pointers, Component Integration Services forwards the following syntax to the remote server:

```
readtext
  [[<database>.]<owner>.]<table_name>.<column_name>
  ><t><ext_pointer> <offset> <size>
  [using {chars | characters}]
```

- `readtext` is issued anytime text or image data must be read. `readtext` is called when a `select` command refers to a text or image column in the select list, or when a `where` clause refers to a text or image column.

For example, you have a proxy table `books` that is mapped to the `books` table on the remote server `foo`. The columns are `id`, `name`, and the text column `blurb`. When the following statement is issued:

```
select * from books
```

Component Integration Services sends the following syntax to the remote server:

```
select id, name, textptr(blurb) from foo_books
```

```
readtext foo_books.blurb @p1 0 0 using chars
```

4.3.12 `select`

Component Integration Services includes changes to the `select` command.

See `select` in *Reference Manual: Commands*.

Server Class `ASEnterprise`

- All syntax is supported. Since the remote server is assumed to have all capabilities necessary to process Transact-SQL syntax, all elements of a `select` command, except those mentioned above, are forwarded to a remote server, using quickpass mode.
- A bulk copy transfer is used to copy data into the new table when a `select...into` command is issued and the `into` table resides on a remote SAP ASE. Both the local and remote databases must be configured with `dboption` set to `select into / bulkcopy`.

Server Class ASAnywhere

- All syntax is supported. Since the remote server is assumed to have all capabilities necessary to process Transact-SQL syntax, all elements of a `select` command, except those mentioned above, are forwarded to a remote server, using quickpass mode.
- If the `select...into` format is used and the `into` table is accessed through the `ASAnywhere` interface, bulk inserts are not used. Instead, Component Integration Services uses Client-Library to prepare a parameterized dynamic `insert` command, and executes it for each row returned by the `select` portion of the command.

Server Class ASIQ

All syntax is supported. Since the remote server is assumed to have all capabilities necessary to process Transact-SQL syntax, all elements of a `select` command, except those mentioned above, are forwarded to a remote server, using quickpass mode.

Server Class direct_connect

- The first time Component Integration Services requires a connection to a server in class `direct_connect`, a request for capabilities is made of the `DirectConnect`. Based on the response, Component Integration Services determines the parts of a `select` command to forward to the `DirectConnect`. In most cases, this is determined by the capabilities of the DBMS with which the `DirectConnect` is interfacing.
- If the entire statement cannot be forwarded to the `DirectConnect` using quickpass mode, Component Integration Services compensates for the functionality that cannot be forwarded. For example, if the remote server cannot handle the `order by` clause, quickpass is not used and Component Integration Services performs a sort on the result set.
- Component Integration Services passes data values as parameters to either a cursor or a dynamic SQL statement. Language statements can also be used if the `DirectConnect` supports it. The parameters are in the datatype native to SAP ASE and must be converted by the `DirectConnect` into formats appropriate for the target DBMS.
- The `select...into` command is supported, but the table must have a unique index if the table has `text` or `image` columns.
- If the `select...into` format is used and the `into` table is accessed through a `DirectConnect`, bulk inserts are not used. Instead, Component Integration Services uses Client-Library to prepare a dynamic `insert` command, and executes it for each row returned by the `select` portion of the command.

4.3.13 truncate table

Component Integration Services includes changes to the `truncate table` command.

See `truncate table` in *Reference Manual: Commands*.

Server Class ASEnterprise

Component Integration Services forwards the `truncate table` command to servers of class ASEnterprise.

Server Class ASAnywhere

Component Integration Services forwards the `truncate table` command to servers of class ASAnywhere.

Server Class ASIQ

Component Integration Services forwards the `truncate table` command to servers of class ASIQ.

Server Class direct_connect

Transact-SQL syntax is sent:

```
truncate table [[<database>.]<owner>.]<table_name>
```

4.3.14 update

Component Integration Services includes changes to the `update` command.

See `update` in *Reference Manual: Commands*.

Server Class ASEnterprise

- If Component Integration Services cannot pass the entire statement to a remote server, a unique index must exist on the table.
- The `update` command is fully supported for all datatypes except `text` and `image`. `text` and `image` data cannot be changed with the `update` command, except when setting the `text` or `image` value to null. Use the `writetext` command instead.
- If quickpass mode is not used, data is retrieved from the source tables, and the values in the target table are updated using a separate cursor designed for handling a positioned `update`.

Server Class ASAnywhere

Handling of the `update` statement is the same as for `ASEnterprise`.

Server Class ASIQ

Handling of the `update` statement is the same as for `ASEnterprise`.

If Component Integration Services cannot forward the original query without alteration, you get an error because ASIQ does not support updatable cursors.

Server Class `direct_connect`

- The following syntax is supported by servers of class `direct_connect`:

```
update [[<database>.]<owner>.]<table_name> | <view_name>
set [[<database>.]<owner>.]<table_name>.|<view_name>.|
< column_name><1> =
    {<expression><1>|NULL| (<select_statement>)}
[, <column_name><2> =
    {<expression><2>|NULL| (<select_statement>)}]...
```

```
[where <search_conditions>]
```

`update` commands that conform to this syntax use quickpass mode, if the capabilities response from the remote server indicates that all elements of the command are supported. Examples of negotiable capabilities include: subquery support, `group by` support, and built-in support.

- If the remote server does not support all elements of the command, or the command contains a `from` clause, Component Integration Services issues a query to obtain the values for the `set` clause, and then issues an `update` command to the remote server.
- Component Integration Services passes data values as parameters to either a cursor or a dynamic SQL statement. Language statements can also be used if the `DirectConnect` supports it. The parameters are in the datatype native to SAP ASE and must be converted by the `DirectConnect` into formats appropriate for the target DBMS.

4.3.15 update statistics

Component Integration Services includes changes to the `update statistics` command.

See `update statistics` in *Reference Manual: Commands*.

Server Class ASEnterprise

- If the table on which the statistics are requested has no indexes, Component Integration Services issues the following command:

```
select count(*) from <table_name>
```

It is also the only command issued when trace flag 11209 is on.

- If the table has an index and the index is specified in the command, Component Integration Services issues the following commands:

```
select count(*) from <table_name><e>
```

```
select count(*) <column_name> [,<column_name>, ...]  
from <table_name>  
group by <column_name> [,<column_name>, ..]
```

The column name(s) represent the column or columns that make up the index.

For example, when the following command is issued:

```
update statistics customers ind_name
```

Component Integration Services issues:

```
select count(*) from customers
```

```
select count(*) last_name, first_name  
from customers  
group by last_name, first_name
```

- If the table has one or more indexes but no index is specified in the statement, Component Integration Services issues the `select count (*)` once, and the `select/order by` commands for each index.
- You must have the `sa_role` to run `update statistic` on a proxy table using a remote login.
- In SAP ASE version 15.0 or later, if a proxy table points to a partitioned table, only global statistics are imported. These are aggregated statistics since proxy tables in SAP ASE version 15.0 are not partitioned.

Server Class ASAnywhere

The processing of `update statistics` in this server class is identical to pre-SAP ASE version 15.0 servers.

Server Class ASIQ

The processing of `update statistics` in this server class is identical to pre-SAP ASE version 15.0 servers.

Server Class `direct_connect`

- The processing of `update statistics` in this server class is identical to that of server class `ASEnterprise` described above.
- If the `direct_connect` indicates that it cannot handle the `group by` or the `count (*)` syntax, statistics are not collected for the `direct_connect`.

4.3.16 `writetext`

Component Integration Services includes changes to the `writetext` command.

See `writetext` in *Reference Manual: Commands*.

Server Class `ASEnterprise`

The `writetext` command is processed using a separate connection to the remote server.

Server Class `ASAnywhere`

The `writetext` command is processed using a separate connection to the remote server.

Server Class `ASIQ`

The `writetext` command is processed using a separate connection to the remote server.

Server Class `direct_connect`

If the `DirectConnect` supports text pointers, Component Integration Services treats the `DirectConnect` as if it were a server in class `ASEnterprise`.

4.4 Passthrough Mode

Passthrough mode is provided within Component Integration Services as a means of enabling a user to perform native operations on the server to which the user is being “passed through.”

For example, requesting passthrough mode for an Oracle server allows you to send native Oracle SQL statements to the Oracle DBMS. Results are converted into a form that is usable by the Open Client application and passed back to the user.

The Transact-SQL parser and compiler are bypassed in this mode, and each language batch received from the user is passed directly to the server to which the user is connected in passthrough mode. Results from each batch are returned to the client.

There are several ways to use passthrough mode:

- `connect to`
- `sp_autoconnect`
- `sp_passthru`
- `sp_remotesql`

4.4.1 The connect to Command

The `connect to` command enables users to specify the server to which a passthrough connection is required.

The syntax of the command is:

```
connect to <server_name>
```

where `<server_name>` is the name of a server added to the `sys.servers` table, with its server class and network name defined. See `sp_addserver` in the *Reference Manual: Procedures*.

When establishing a connection to `<server_name>` on behalf of the user, the server uses:

- A remote login alias set using `sp_addexternlogin`, or
- The name and password used to communicate with the SAP ASE.

In either case, if the connection cannot be made to the server specified, the reason is contained in a message returned to the user.

Once a passthrough connection has been made, the Transact-SQL parser and compiler are bypassed when subsequent language text is received. Any statements received by the server are passed directly to the specified remote server.

i Note

Some database management systems do not recognize more than one statement at a time and produce syntax errors if, for example, multiple `select` statements were received as part of a single language text buffer.

After statements are passed to the requested server, any results are converted into a form that can be recognized by the Open Client interface and sent back to the client program.

To exit from passthrough mode, issue the `disconnect` or `disc`, command. Subsequent language text from this client is then processed using the Transact-SQL parser and compiler.

Permission to use `connect to` must be explicitly granted by the system administrator. The syntax is:

```
grant connect to <user_name>
```

To revoke permission to use `connect to`, the syntax is:

```
revoke connect from <user_name>
```

The `connect to` permissions are stored in the `master` database. To globally grant or revoke permissions to “public”, the system administrator sets the permissions in the `master` database; the effect is server-wide, regardless of what database is being used. The system administrator can only grant or revoke permissions to or from a user if the user is a valid user of the `master` database.

The system administrator can grant or revoke “all” permissions to or from “public” within any database. If the permissions are in the `master` database, “all” includes the `connect to` command. If they are in another database, “all” does not include the `connect to` command.

In this example, the system administrator wants to revoke permission from “public” and wants only the user “fred” to be able to execute the `connect to` command. “fred” must be made a valid user of `master`. To do this, the system administrator issues the following commands in `master`:

```
revoke connect from public
sp_adduser fred
grant connect to fred
```

4.4.2 The `sp_autoconnect` System Procedure

Some users may always require a passthrough connection to a given server. If this is the case, Component Integration Services can be configured so that it automatically connects these users to a specified remote server in passthrough mode when the users connect to the server.

This feature is enabled and disabled by `sp_autoconnect` using this syntax:

```
sp_autoconnect <server_name>, true|false [,<loginname>]
```

Before using `sp_autoconnect`, add the `<server_name>` to `sys.servers` using `sp_addserver`.

A user can request automatic connection to a server using `sp_autoconnect`, but only the system administrator can enable or disable automatic passthrough connection for another user. Thus, only the system administrator can specify a third argument to this procedure.

If the second argument is `true`, the `autoconnect` feature is enabled for the current user (or the user specified in the third argument). If the second argument is `false`, `autoconnect` is disabled.

When a user connects to the server, that user’s `autoconnect` status in `syslogins` is checked. If enabled, the `<server_name>`, also found in `syslogins` (placed there by `sp_autoconnect`), is checked for validity. If the server is valid, the user is automatically connected to that server, and a passthrough status is established. Subsequent language statements received by the server from this user are handled exactly as if the user

explicitly entered the `connect` command. This user then views the server similar to a passthrough gateway to the remote server.

When an “autoconnected” user executes a `disconnect`, she or he is returned normally to the server.

If the remote server cannot be reached, the user (unless the user is assigned the “sa” role) will not be connected to the local SAP ASE. A “login failed” error message is returned.

4.4.3 The `sp_passthru` System Procedure

`sp_passthru` allows the user to pass a SQL command buffer to a remote server.

The syntax of the SQL statements being passed is assumed to be the syntax native to the class of server receiving the buffer; no translation or interpretation is performed. Results from the remote server are optionally placed in output parameters. The syntax for `sp_passthru` is:

```
sp_passthru <server, command, errcode, errmsg, rowcount>
> [<, arg1, arg2, ... arg>n]
```

where:

- `<server>` is the name of the server that is to receive the SQL command buffer; the datatype is `varchar(255)`.
- `<command>` is the SQL command buffer; the datatype is `varchar(255)`.
- `<errcode>` is the error code returned by the remote server; the datatype is `int` output.
- `<errmsg>` is the error message returned by the remote server; the datatype is `varchar(1024)` output.
- `<rowcount>` is the number of rows affected by the last command in the command buffer; the datatype is `int` output.
- `<arg1 - argn>` are optional parameters. If provided, these output parameters will receive the results from the last row returned by the last command in the command buffer. The datatypes may vary. All must be output parameters.

This example returns the date from the Oracle server in the output parameter `@oradate`. If an Oracle error occurs, the error code is placed in `@errcode` and the corresponding message is placed in `@errmsg`. The `@rowcount` parameter is set to 1.

```
sp_passthru ORACLE, "select date from dual", @errcode output, @errmsg output,
@rowcount output, @oradate output
```

4.4.4 The `sp_remotesql` System Procedure

`sp_remotesql` allows you to pass native syntax to a remote server. The procedure establishes a connection to a remote server, passes a query buffer, and relays the results back to the client.

The syntax for `sp_remotesql` is:

```
sp_remotesql server_name, <query_buf1>
>[, <query_buf2>, ... , <query_buf254>]
```

where:

- `<server_name>` is the name of a server that has been defined using `sp_addserver`.
- `<server_name>` is a `varchar(255)` field. If `<server_name>` is not defined or is not available, the connection fails, and the procedure is aborted. This parameter is required.
- `<query_buf1>` is a query buffer of type `char` or `varchar` with a maximum length of 255 bytes. This parameter is required.

Each additional buffer is `char` or `varchar` with a maximum length of 255 bytes. If supplied, these optional arguments are concatenated with the contents of `<query_buf1>` into a single query buffer.

In this example, the server passes the query buffer to `<freds_server>`, which interprets the `select @@version` syntax and returns version information to the client. The returned information is not interpreted by the server.

```
sp_remotesql <freds_server>, "select @@version"
```

4.5 Quoted Identifier Support

Quoted identifiers are forwarded to remote servers that support them.

This is triggered by a `set` command:

```
set quoted_identifier on
```

If this thread property is enabled, Component Integration Services quotes identifiers before sending SQL statements to remote servers.

Remote servers must have the ability to support quoted identifiers. There is a capability in the `sp_capabilities` result set reserved for this purpose:

- Capability ID: 135
- Capability name: quoted identifier
- Capability value: 0 = no support; 1 = supported

The capability defaults to 0 for DirectConnects that do not provide a value for this capability.

4.6 Delimited Identifier Support

The behavior of bracketed identifiers is identical to quoted identifiers, with the exception that you do not need to set `quoted_identifier` on in order to use them.

4.7 auto identity Option

When the `auto_identity` database option is enabled, an `IDENTITY` column is added to any tables that are created in the database.

The column name is `CIS_IDENTITY_COL`, for proxy tables, or `SYB_IDENTITY_COL`, for local tables. In either case, the column can be referenced using the `syb_identity` keyword.

4.8 Triggers

Component Integration Services allows triggers on proxy tables; however, their usefulness is limited. You can create a trigger on a proxy table and the trigger is invoked just as it would be for a normal SAP ASE table.

However, before and after `image` data is not written to the log for proxy tables because the `insert`, `update`, and `delete` commands are passed to the remote server. The `inserted` or `deleted` tables, which are actually views into the log, contain no data for proxy tables. Users cannot examine the rows being inserted, deleted, or updated, so a trigger with a proxy table has limited value.

In SAP ASE version 15.0 and later, there is no support for the `updated` function with triggers.

5 Tutorial

As a system administrator, you perform tasks like configuring the server to access remote data sources.

This tutorial includes instructions for:

- Adding a remote server
- Mapping remote objects to local proxy tables
- Performing joins between remote tables

Routine system administration tasks such as starting and stopping SAP ASE, creating logins, creating groups, adding users, granting permissions, and password administration are explained in the SAP ASE documentation.

i Note

This tutorial assumes that the `pubs2` database has been installed.

5.1 Adding a Remote Server

You can use the server to access data on remote servers.

Procedure

1. Add the remote server to the interfaces file.
Use the `dsedit` or `dscp` utility to edit the interfaces file (called `sql.ini` on Windows), located in the `§SYBASE` directory. See the *Configuration Guide* for your platform.
2. Add the name, server class, and network name of the remote server to system tables.
Optionally, assigning an alternate login name and password.

5.1.1 Creating Server Entries in System Tables

Use `sp_addserver` to add entries to the `sys.servers` table.

Context

`sp_addserver` creates entries for the local server and an entry for each remote server that is to be called. The `sp_addserver` syntax is:

```
sp_addserver <server_name> [,server_class [,network_name]]>
```

where:

- `<server_name>` is the name used to identify the server. It must be unique.
- `<server_class>` is one of the supported server classes. The default value is `ASEnterprise`. If `server_class` is set to `local`, `<network_name>` is ignored.
- `<network_name>` is the server name in the interfaces file. This name may be the same as `<server_name>`, or it may differ. The `<network_name>` is sometimes referred to as the *physical name*.

You must reboot SAP ASE after your add a local server.

Procedure

- The following examples create entries for the local server named `big_server` and for the remote server `CTOSDEMO` with server class `ASEnterprise`.

```
sp_addserver big_server, local
sp_addserver CTOSDEMO, ASEnterprise, CTOSDEMO
```

5.1.2 Adding an Alternate Login and Password

Use `sp_addexternlogin` to assign an alternate login name and password to be used when communicating with a remote server.

The syntax for `sp_addexternlogin` is:

```
sp_addexternlogin <remote_server>, <login_name>, <remote_name>< >[,
<remote_password>]
```

where:

- `<remote_server>` is the name of the remote server. The `<remote_server>` must be known to the local server by an entry in the `master.dbo.sys.servers` table.

- `<login_name>` is an account known to the local server. `<login_name>` must be represented by an entry in the `master.dbo.syslogins` table. The “sa” account, the “sso” account, and the `<login_name>` account are the only users authorized to modify remote access for a given local user.
- `<remote_name>` is an account known to the `<remote_server>` and must be a valid account on the node where the `<remote_server>` runs. This is the account used for logging in to the `<remote_server>`.
- `<remote_password>` is the password for `<remote_name>`.

This example allows the local server to gain access to remote server FRED using the remote name “system” and the remote password “sys_pass” on behalf of user “sa”:

```
sp_addexternlogin FRED, sa, system, sys_pass
```

This example tells the local server that when the login name “bobj” logs in, access to the remote server OMNI1012 is by the remote name “jordan” and the remote password “hitchpost”. Only the “bobj” account, the “sa” account, and the “sso” account have the authority to add or modify a remote login for the login name “bobj”.

```
sp_addexternlogin OMNI1012, bobj, jordan, hitchpost
```

5.1.3 Verifying Connectivity

The `connect to` command establishes a passthrough mode connection to the remote server. This passthrough mode remains in effect until you issue a `disconnect` command.

Use the `connect to <server_name>` command to verify that the configuration is correct. `connect to` requires that “sa” explicitly grant connect authority to users other than “sa.”

5.2 Performing a Join Between Two Remote Tables

Component Integration Services allows you to perform joins across remote tables.

Procedure

1. Edit the interfaces file using `dsedit`.
2. Use `sp_addserver` to add entries to the `sys.servers` system table. On the server originating the call, there must be an entry for each remote server that is to be called. The following examples create entries for the local server named `big_server` and for the remote server `big_server` of class `ASEnterprise`

```
sp_addserver big_server, local
sp_addserver CTOSDEMO, ASEnterprise, big_server
```

3. Use `create existing table` to enable the definition of existing (proxy) tables. When the server processes this command, it does not create a new table. Instead, it checks the table mapping and verifies

the existence of the underlying object. If the object does not exist (either host data file or remote server object), the server rejects the command and returns an error message to the client.

After you define an existing table, issue an `update statistics` command for that table. This helps the query optimizer make intelligent choices regarding index selection and join order.

This demonstrates creating remote SAP ASE tables `publishers` and `titles` in the sample `pubs2` database mapped to a local server:

1. Define a server named `big_server`. Its server class is `ASEnterprise`, and its name in the interfaces file is `big_server`:

```
exec sp_addserver big_server, ASEnterprise, big_server
```

2. Define a remote login alias. This step is optional. User "sa" is known to remote server `big_server` as user "sa," password "timothy":

```
exec sp_addexternlogin big_server, sa, sa, timothy
```

3. Define the remote `publishers` table:

```
create proxy_table publishers  
at "big_server.pubs2.dbo.publishers"
```

4. Define the remote `titles` table:

```
create proxy_table titles  
at "big_server.pubs2.dbo.titles"
```

4. Use the `select` statement to perform the join. For example:

```
select Publisher = p.pub_name, Title = t.title  
from publishers p, titles t  
where p.pub_id = t.pub_id  
order by p.pub_name
```

6 Troubleshooting

Troubleshooting requires that you have enough information about certain error conditions that you can resolve problems without help from Technical Support, and lists of information that you can gather before calling Technical Support, which may help resolve your problem more quickly

The *Troubleshooting and Error Messages Guide* should also be used for troubleshooting. While this appendix provides troubleshooting tips for most frequently asked Component Integration Services questions, lists all error messages with a one-line recovery procedure; the *Troubleshooting and Error Messages Guide* provides tips on SAP ASE problems that are not specific to Component Integration Services.

6.1 Problems Accessing Component Integration Services

SAP ASE issues an error message if you issue a command that accesses a remote object and Component Integration Services is not found.

Procedure

- SAP ASE issues this message:

```
cis extension not enabled or installed
```

Verify that the `enable cis` configuration parameter is set to 1 by running:

```
sp_configure "enable cis"
```

- `sp_configure` returns the following row for the `enable cis` parameter:

name	min	max	config value	run value
enable cis	0	1	1	1

Both “config value” and “run value” should be 1. If both values are 0, set the `enable cis` configuration parameter to 1, and restart the server. Use the syntax:

```
sp_configure "enable cis" 1
```

If “config value” is 1 and “run value” is 0, the `enable cis` configuration parameter is set, but does not take effect until the server is restarted.

i Note

Component Integration Services is enabled by default beginning with SAP ASE version 12.0.

6.2 Problems Using Component Integration Services

Problems you may encounter when using Component Integration Services include being unable to access the remote server or objects, and problems retrieving data from remote objects.

6.2.1 Unable to Access Remote Server

SAP ASE issues error message 11206 when you cannot access a remote server.

SAP ASE returns this message when you cannot access a remote server:

```
11206 Unable to connect to server server_name.
```

The message is preceded by one of these Client-Library messages:

Requested server name not found

The server is not defined in the interfaces file when the following messages display:

```
11206 Unable to connect to server server_name.
```

When a remote server is added using `sp_addserver`, the interfaces file is not checked. It is checked the first time you try to make a connection to the remote server. To correct this problem, add the remote server to the interfaces file that is being used by Component Integration Services.

Driver call to connect two endpoints failed

If the remote server is defined in the interfaces file, but no response was received from the connect request, the following messages are displayed:

```
Driver call to connect two endpoints failed
11206      Unable to connect to server <server_name>.
```

- Verifying that your environment is set up correctly. To test this, try to connect directly to the remote server using `isql` or a similar tool:
 1. Log in to the machine where Component Integration Services is running.
 2. Set the SYBASE environment variable to the same location that was used when Component Integration Services was started. Component Integration Services uses the interfaces file in the directory specified by the SYBASE environment variable, unless it is overridden in the `runserver` file by the `-i` argument.
These first two steps are important to ensure that the test environment is the same environment that Component Integration Services was using when you could not connect to the remote server.

3. Use `isql` or a similar tool to connect directly to the remote server.

If the environment is set up correctly and the connection fails, continue through this list. If the connection is made, there is a problem with the environment being used by Component Integration Services.

- Verify that the remote server is up and running. Log in to the machine where the remote server is located to verify that the server is running. If the server is running, continue through this list. If the server is not running, restart the server and try your query again.
- Verifying that the entry for the remote server in the interfaces file is correct.
 - Verify that the machine name is the correct name for the machine the software is loaded on.
 - Verify that if the interfaces file is a text file, the query and master lines start with a tab and not spaces.
 - Verify that the port number is available. Check the `services` file in the `/etc` directory to ensure that the port number is not reserved for another process.

Login Failed

If you can access the remote server, but the login name and password are correct, the following messages display:

```
Login failed
11206 Unable to connect to server <server_name>.
```

See if there is an external login established for the remote server by executing:

```
exec sp_helpexternlogin <server_name>
```

If no external login is defined, Component Integration Services uses the user login name and password that was used to connect to SAP ASE. For example, if the user connected to SAP ASE using the "sa" account, Component Integration Services uses the login name "sa" when making a remote connection. Unless the remote server is another SAP ASE, the "sa" account probably does not exist, and an external login must be added using `sp_addexternlogin`.

If an external login is defined, verify that the user's login name is correct. Remote server logins are case sensitive. Is the case correct for the user login name you are using and the entry in `externlogins`?

If the login name is correct, the password might be incorrect. You cannot display the password. If the user login name is incorrect or if the password might be incorrect, drop the existing external login and redefine it by executing the commands:

```
exec sp_dropexternlogin <server_name>, <login_name>
go
exec sp_addexternlogin <server_name>, <login_name>, <remote_login>,
<remote_password>
go
```

6.2.2 Unable to Access Remote Objects

SAP ASE issues error message 11214 when you are unable to access a remote object.

```
Error 11214 Remote object <object> does not exist.
```

The problem may be in the local proxy table definition or in the table itself on the remote server.

Verify that:

- The object has been defined in Component Integration Services.

To confirm, run:

```
sp_help <object_name>
```

If the object does not exist, create the object in Component Integration Services.

- If the object has been defined in Component Integration Services, the definition is correct. Table names can have four parts with the format `server.dbname.owner.tablename`. The `dbname` part is not valid for Oracle, or InfoHUB servers. If the object definition is incorrect, delete it using `sp_dropobjectdef`, and define it correctly using `sp_addobjectdef`.
- If the local object definition is correct, check the table on the remote server, to verify that:
 - Permissions are set to allow access to both the database and table.
 - The database has been marked suspect.
 - The database is available.
 - You can access the remote table using a native tool (for example, SQL*Plus on Oracle).

6.2.3 Problem Retrieving Data from Remote Objects

When you receive error messages pertaining to mismatches in remote objects, the Component Integration Services object definition does not match the remote object definition.

This happens if:

- The object definition was altered outside of Component Integration Services.
- An index was added or dropped outside of Component Integration Services.

Object is Altered Outside Component Integration Services

Once an object is defined in Component Integration Services, alterations made to an object at the remote server are not made to the local proxy object definition. If an object is altered outside of Component Integration Services, the steps to correct the problem differ, depending on whether `create existing table` or `create table` was used to define the object.

To determine which method was used to define the object, run:

```
sp_help <object_name>
```

If:

- You defined the object via `create existing table`, the following message is returned in the result set:

```
Object created with 'existing' option
```

- This message is not displayed, the object was defined via `create table`.

- You used `create existing table` to create the table in Component Integration Services:
 1. Use `drop table` in Component Integration Services.
 2. Create the table again in Component Integration Services using `create existing table`. This creates the table using the new version of the table on the remote server.
- You created the table in Component Integration Services using `create table`, the remote object is dropped when you use `drop table`. To prevent this:
 1. Rename the table on the remote server so the table is not deleted when you use `drop table`.
 2. Create a table on the remote server using the original name.
 3. Use `drop table` in Component Integration Services to drop the table in Component Integration Services and on the remote server.
 4. Rename the saved table in step 1 with its original name on the remote server.
 5. Create the table again in Component Integration Services using `create existing table`.

⚠ Caution

Do not use `drop table` in Component Integration Services before renaming the table on the remote server, or you will delete the table on the remote server.

A good rule to follow is to create the object on the remote server, and then execute `create existing table` to create the object in Component Integration Services. This enables you to correct mismatch problems with fewer steps and with no chance of deleting objects on the remote server.

Index Is Added or Dropped Outside Component Integration Services

Component Integration Services is unaware of indexes that are added or dropped outside Component Integration Services. Verify that the indexes used by Component Integration Services are the same as the indexes used on the remote server. Use `sp_help` to see the indexes used by Component Integration Services. Use the appropriate command on your remote server to verify the indexes used by the remote server.

If:

- The indexes are not the same, the steps to correct the problem differ, depending on whether `create existing table` or `create table` was used to define the object.
To determine which method was used to define the object, run:

```
sp_help <object_name>
```

- The object was defined via the `create existing table` command, the following message is returned in the result set:

```
Object created with 'existing' option
```

- This message is not displayed, the object was defined via `create table`.
- The `create existing table` command was used to create the object:
 1. Use `drop table` in Component Integration Services.
 2. Re-create the table in Component Integration Services using `create existing table`. This updates the indexes to match the indexes on the remote table.
 3. If `create table` was used to create the object:

1. Use `drop index` to drop the index from the remote table.
 2. Re-create the index in Component Integration Services using `create index`. This creates the index in Component Integration Services and the remote server.
- If `create table` was used to define the object, an alternative method is to turn on trace flag 11208. This trace flag prevents `create index` from transmitting to the remote server. To use trace flag 11208, follow these steps:
 1. Turn on trace flag 11208:

```
dbcc traceon(11208)
```

2. Create the index using `create index`.
3. Turn off trace flag 11208:

```
dbcc traceoff(11208)
```

6.3 If You Need Help

If you encounter a problem that you cannot resolve using the manuals, ask the designated person at your site to contact Technical Support.

Context

Gather the following information prior to calling Technical Support to help resolve your problem more quickly.

- If a problem occurs while you are trying to access remote data, execute the same script against a local table. If the problem does not exist on the local table, it is specific to Component Integration Services and you should continue through this list.
- Find out what version of SAP ASE you are using:

```
select @@version
```

- Note the SQL script that reproduces the problem. Include the script that was used to create the tables.
- Find the processing plan for your query. This is generated using `set showplan`. An example of this is:

```
set showplan, noexec on
go
select au_lname, au_fname from authors
where au_id = 'A1374065371'
go
```

The output for this query looks like this:

```
set showplan, noexec on
go
select au_lname, au_fname from authors where au_id = 'A1374065371'
go
The Abstract Plan (AP) of the final query execution plan:
( remote_sql )
To experiment with the optimizer behavior, this AP can be modified and then
```

```

passed to the optimizer using the PLAN clause:
SELECT/INSERT/DELETE/UPDATE ...
PLAN '( ... )
QUERY PLAN FOR STATEMENT 1 (at line 1).
1 operator(s) under root
The type of query is SELECT.
ROOT:EMIT Operator
  |LE_REMSCANOP Operator
  |   SELECT "au_lname" , "au_fname" FROM pubs2.dbo."authors"
WHERE "au_
  |   id" = 'A1374065371'

```

The `noexec` option compiles the query, but does not execute it. No subsequent commands are executed until `noexec` is turned off.

- Obtain the event logging when executing the query by turning on trace flags 11201 – 11205. These trace flags log the following:
 - 11201 – client connect, disconnect, and attention events.
 - 11202 – client language, cursor declare, dynamic prepare, and dynamic execute-immediate text.
 - 11203 – client RPC events.
 - 11204 – messages routed to client.
 - 11205 – interaction with remote servers.
 - 11206 – logs file and directory processing steps.
 - 11207 – logs text and image processing.

After executing the script with the trace flags turned on, the logging is found in the error log in the `$$SYBASE/install` directory. For example:

```

dbcc traceon (11201,11202,11203,11204,11205)
go
select au_lname, au_fname from authors
where au_id = 'A1374065371'
go
dbcc traceoff (11201,11202,11203,11204,11205)
go

```

The error log output is as follows (the timestamps printed at the beginning of each entry have been removed to improve legibility):

```

server TDS_LANG, spid 15: command text:
select au_lname, au_fname from authors where au_id = 'A1374065371'
server RemoteAccess constructed
server EXECLANG, spid 15, server huntington0_19442, quickpass statement:
ELECT "au_lname" , "au_fname" FROM pubs2.dbo."authors" WHERE "au_id" =
'A1374065371'
server BINDCOLS, spid 15: column 1, name au_lname, fmt.type 'CHAR',
fmt.maxlen 40, fmt.stat 16, con.type 'VARCHAR', con.maxlen 40
server BINDCOLS, spid 15: column 2, name au_fname, fmt.type 'CHAR',
fmt.maxlen 20, fmt.stat 16, con.type 'VARCHAR', con.maxlen 20
server BINDCOLS, spid 15: bind array size 50, total memory required is 4304
bytes
server FETCH , spid 15: cursor C1; ct_fetch() returned 0 rows; status -204
server RemoteAccess deleted

```



This tracing is global, so once the trace flags are turned on, any query that is executed is logged; therefore, turn tracing off once you have your log. Also, clean out the error log periodically by bringing the server down, renaming the error log, and restarting the server. This creates a new error log.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Gender-Related Language

We try not to use gender-specific word forms and formulations. As appropriate for context and readability, SAP may use masculine word forms to refer to all genders.

© 2019 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.