



## **User-Defined Functions**

---

# **Sybase IQ 15.4**

DOCUMENT ID: DC01034-01-1540-01

LAST REVISED: November 2011

Copyright © 2011 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

# Contents

<b>Audience .....</b>	<b>1</b>
<b>Understanding User-Defined Functions .....</b>	<b>3</b>
Learning Roadmap: Types of UDFs .....	5
Learning Roadmap: Types of External C and C++ UDFs .....	6
User-Defined Functions Compliance with Sybase IQ Databases .....	7
Practices to Avoid .....	8
Naming Conventions for User-Defined Functions .....	8
SQL Data Types .....	9
<b>Building UDFs .....</b>	<b>15</b>
Design Basics of User-Defined Functions .....	15
Setting the Dynamic Library Interface .....	15
Upgrading to the v4 API .....	16
Library Version (extfn_get_library_version) .....	17
Library Version Compatibility (extfn_check_version_compatibility) .....	17
License Information (extfn_get_license_info) .....	18
Adding the extfn_get_license_info Method .....	19
Compile and Link Source Code to Build Dynamically Linkable Libraries .....	19
Compiling and Linking the Sample UDFs for Windows .....	21
Compiling and Linking the Sample UDFs for UNIX .....	21
AIX Switches .....	21
HP-UX Switches .....	22
Linux Switches .....	22
Solaris Switches .....	23
Windows Switches .....	24
Testing User-Defined Functions .....	25

Enabling and Disabling User-Defined Functions .....	25
Initially Executing a User-Defined Function .....	26
Controlling Error Checking and Call Tracing .....	27
Viewing Sybase IQ Log Files .....	28
Using Microsoft Visual Studio Debugger for User-Defined Functions .....	28
Modifying the UDF at Runtime .....	28
Granting and Revoking Permissions .....	29
Dropping User-Defined Functions .....	29
<b>Scalar and Aggregate UDFs .....</b>	<b>31</b>
Scalar and Aggregate UDF Restrictions .....	31
Creating a Scalar or Aggregate UDF .....	31
Creating a UDF Using SQL Anywhere Dialects in Sybase Central .....	32
Declaring and Defining Scalar User-Defined Functions .....	33
Declaring and Defining Aggregate UDFs .....	47
Calling Scalar and Aggregate UDFs .....	83
Scalar and Aggregate UDF Calling Patterns .....	83
Scalar and Aggregate UDF Callback Functions .....	84
Scalar UDF Calling Pattern .....	85
Aggregate UDF Calling Patterns .....	85
<b>Table UDFs and TPFs .....</b>	<b>97</b>
User Roles .....	97
Learning Roadmap for Table UDF Developers .....	97
Learning Roadmap for SQL Analysts .....	98
Table UDF Restrictions .....	99
Get Started .....	99
Sample Files .....	99
Understanding Producers Versus Consumers ..	101
Developing a Table UDF .....	103
Table UDF Implementation Examples .....	105
Query Processing States .....	121

Initial State .....	121
Annotation State .....	122
Query Optimization State .....	124
Plan Building State .....	127
Execution State .....	128
Row Block Data Exchange .....	128
Fetch Methods for Row Blocks .....	129
Using a Row Block to Produce Data .....	131
Row Block Allocation .....	132
Table UDF Query Plan Objects .....	134
Enabling Memory Tracking .....	135
Table Parameterized Functions .....	136
Learning Roadmap for TPF Developers .....	136
Developing a TPF .....	136
Consume Table Parameters .....	137
Order Input Table Data .....	140
Partitioning Input Data .....	140
TPF Implementation Examples .....	156
SQL Reference for Table UDF and TPF Queries .....	166
ALTER PROCEDURE Statement .....	167
CREATE PROCEDURE Statement (Table UDF) .....	169
CREATE FUNCTION Statement .....	172
DEFAULT_TABLE_UDF_ROW_COUNT Option	178
TABLE_UDF_ROW_BLOCK_SIZE_KB Option	179
FROM Clause .....	179
SELECT Statement .....	185
<b>API Reference for a_v4_extfn .....</b>	<b>195</b>
Blob (a_v4_extfn_blob) .....	195
blob_length .....	196
open_istream .....	197
close_istream .....	197
release .....	198
Blob Input Stream (a_v4_extfn_blob_istream) .....	199
get .....	199

Column Data (a_v4_extfn_column_data) .....	200
Column List (a_v4_extfn_column_list) .....	202
Column Order (a_v4_extfn_order_el) .....	202
Column Subset (a_v4_extfn_col_subset_of_input) ...	203
Describe API .....	204
*describe_column_get .....	204
*describe_column_set .....	221
*describe_parameter_get .....	238
*describe_parameter_set .....	257
*describe_udf_get .....	273
*describe_udf_set .....	275
Describe Column Type	
(a_v4_extfn_describe_col_type) .....	277
Describe Parameter Type	
(a_v4_extfn_describe_parm_type) .....	278
Describe Return (a_v4_extfn_describe_return) .....	280
Describe UDF Type (a_v4_extfn_describe_udf_type)	
.....	282
Execution State (a_v4_extfn_state) .....	283
External Function (a_v4_extfn_proc) .....	284
_start_extfn .....	285
_finish_extfn .....	285
_evaluate_extfn .....	286
_describe_extfn .....	286
_enter_state_extfn .....	287
_leave_state_extfn .....	287
External Procedure Context	
(a_v4_extfn_proc_context) .....	287
get_value .....	289
get_value_is_constant .....	292
set_value .....	292
get_is_cancelled .....	293
set_error .....	294
log_message .....	295
convert_value .....	295

get_option .....	296
alloc .....	297
free .....	298
open_result_set .....	298
close_result_set .....	299
get_blob .....	300
set_cannot_be_distributed .....	301
License Information (a_v4_extfn_license_info) .....	301
Optimizer Estimate (a_v4_extfn_estimate) .....	302
Order By List (a_v4_extfn_orderby_list) .....	302
Partition By Column Number	
(a_v4_extfn_partitionby_col_num) .....	303
Row (a_v4_extfn_row) .....	304
Row Block (a_v4_extfn_row_block) .....	305
Table (a_v4_extfn_table) .....	306
Table Context (a_v4_extfn_table_context) .....	306
fetch_into .....	308
fetch_block .....	311
rewind .....	313
get_blob .....	314
Table Functions (a_v4_extfn_table_func) .....	314
_open_extfn .....	316
_fetch_into_extfn .....	317
_fetch_block_extfn .....	318
_rewind_extfn .....	318
_close_extfn .....	319
<b>API Troubleshooting for a_v4_extfn .....</b>	<b>321</b>
Generic describe_column Errors .....	321
Generic describe_udf Errors .....	322
Generic describe_parameter Errors .....	322
Missing UDF Returns an Error .....	323
<b>External Environment for UDFs .....</b>	<b>325</b>
Executing UDFs from an External Environment .....	326
External Environment Restrictions .....	327
The CLR External Environment .....	327

## Contents

The ESQL and ODBC External Environments .....	330
The Java External Environment .....	339
Java External Environment in a Multiplex .....	345
Java External Environment Restrictions .....	346
Java VM Memory Options .....	346
SQL Data Type Conversions for Java UDFs .....	347
Creating a Java Scalar UDF .....	349
Creating a Java Scalar UDF Version of the SQL substr Function .....	351
Creating a Java Table UDF .....	352
Java External Environment SQL Statement Reference .....	355
The PERL External Environment .....	362
The PHP External Environment .....	366
<b>Index</b> .....	<b>373</b>



# Audience

The *User-Defined Functions Guide* is intended for SQL analysts, C developers, C++ developers, and Java developers who want to extend the functionality of Sybase® IQ.

As a developer, use the tasks, concepts, and API reference material to program non-SQL external user-defined functions.

As a SQL analyst, use this guide to develop SQL queries that reference non-SQL external user-defined functions.

Audience

# Understanding User-Defined Functions

Learn how user-defined functions are used within Sybase IQ.

Sybase IQ allows user defined functions (UDFs), which execute within the database container. The UDF execution feature is available as an optional component for use within Sybase IQ or within the RAPStore component of RAP – The Trading Edition Enterprise®.

You must be specifically licensed to use these external C/C++ UDFs interfaces.

These external C/C++ UDFs differ from the Interactive SQL UDFs available in earlier versions of Sybase IQ. Interactive SQL UDFs are unchanged and do not require a special license.


UDFs that execute within Sybase IQ take advantage of the extreme performance of the server, while also providing users the flexibility of analyzing their data with the flexibility of a programmatic solution. User-Defined Functions consist of two components:

- UDF declaration, and
- UDF executable code

A UDF is declared in the SQL environment through a SQL function or stored procedure which describes the parameters and provides a reference to the external library.

The actual executable portion of the UDF is contained within an external (shared object or dynamic load) library file, which is automatically loaded by the server upon the first invocation of a UDF Declaration function or stored procedure associated with that library. Once loaded, the library remains resident in the server for rapid access through subsequent invocations of SQL functions or stored procedures that reference the library.

The Sybase IQ user-defined function architecture is represented in the diagram below.



Sybase IQ supports high-performance in-process external C/C++ user-defined functions. This style of UDF supports functions written in C or C++ code that adhere to the interfaces described in this guide.

The C/C++ source code for the UDFs is compiled into one or more external libraries that are subsequently loaded into the server's process space when needed. The UDF calling mechanism is defined to the server through a SQL function. When the SQL function is invoked from a SQL query, the server loads the corresponding library if it has not already been loaded.

For simplicity of managing the UDF installation, Sybase recommends that UDF developers package many UDF functions within a single library.

To facilitate the construction of UDFs, Sybase IQ includes a C-based API. The API comprises a set of predefined entry points for the UDFs, a well-defined context data structure, and a series of SQL callback functions that provide a communication mechanism from the UDF back to

the server. The Sybase IQ UDF API allows software vendors and expert end-users to develop, package, and sell their own UDFs.

## Learning Roadmap: Types of UDFs

The types of user-defined functions (UDFs) available in Sybase IQ.

UDF Type	Description	Required Li- cense	See
UDF (SQL)	A user-defined function written in SQL.	none	<i>System Administration Guide Volume 2 &gt; Using Procedures and Batches &gt; Introduction to User-Defined Functions</i>
Scalar C or C++ UDF	V3 external C or C++ procedure that operates on a single value.	IQ_UDF	<i>Learning Roadmap: Types of External C and C++ UDFs</i> on page 6
Scalar C or C++ UDF	V4 external C or C++ procedure that operates on a single value.	IQ_IDA	<i>Learning Roadmap: Types of External C and C++ UDFs</i> on page 6
Aggregate C or C++ UDF	V3 external C or C++ procedure that operates on multiple values. Aggregate UDFs are also sometimes known as UDAs or UDAFs. The context structure for coding aggregate UDFs is slightly different than the context structure used for coding scalar UDFs.	IQ_UDF	<i>Learning Roadmap: Types of External C and C++ UDFs</i> on page 6
Aggregate C or C++ UDF	V4 external C or C++ procedure that operates on multiple values. Aggregate UDFs are also sometimes known as UDAs or UDAFs. The context structure for coding aggregate UDFs is slightly different than the context structure used for coding scalar UDFs.	IQ_IDA	<i>Learning Roadmap: Types of External C and C++ UDFs</i> on page 6

UDF Type	Description	Required License	See
Table UDF	External C or C++ procedure that produces a set of rows and can be used as a table expression in the <b>FROM</b> clause of a SQL statement	IQ_IDA	<i>Learning Roadmap: Types of External C and C++ UDFs</i> on page 6
Table parameterized function	A table UDF that accepts table (non-scalar) parameters in addition to scalar parameters, and can be executed in parallel over partitions of row-sets. Also known as table parameterized user-defined functions.	IQ_IDA	<i>Learning Roadmap: Types of External C and C++ UDFs</i> on page 6
Java scalar UDF	An out-of-process (external environment) scalar user-defined function implemented in Java code.	none	<i>The Java External Environment</i> on page 339
Java table UDF	An out-of-process (external environment) table UDF implemented in Java code.	none	<i>The Java External Environment</i> on page 339

## Learning Roadmap: Types of External C and C++ UDFs

The high-performance, in-process, external C and C++ user-defined functions available with the IQ\_IDA license.

The v3 API requires either the IQ\_UDF or IQ\_IDA license. The v4 API requires the IQ\_IDA license.

UDF Type	Input Parameters	Return	API	See:
Scalar UDF	Scalar	Single scalar value	v3, v4	<i>Declaring and Defining Scalar User-Defined Functions</i> on page 33
Aggregate UDF	Scalar	Single scalar value	v3, v4	<i>Declaring and Defining Aggregate UDFs</i> on page 47

UDF Type	Input Parameters	Return	API	See:
Table UDF	Scalar	Table	v4	<i>Table UDFs and TPFs</i> on page 97
Table parameterized function (TPF)	Scalar and table	Table	v4	<i>Table Parameterized Functions</i> on page 136

These UDFs can be deterministic or nondeterministic. The result of a function can be determined by the input parameters and data (deterministic), or by some random behavior (nondeterministic). Parameters of nondeterministic UDFs typically need a random seed as one of the input parameters.

## User-Defined Functions Compliance with Sybase IQ Databases

---

Develop user-defined functions to work with Sybase IQ databases.

### Seamless Execution

UDFs must run seamlessly within the database container. Although Sybase IQ is a complex product consisting of many files, the main user interaction is through a server process (iqsrv15), using industry-standard Structured Query Language (SQL). Execution of UDFs should be accomplished entirely through SQL commands; the user does not need to understand the underlying implementation method to use the UDFs.

The `EXTFN_V3_API` and `EXTFN_V4_API` provide callback functions enabling the UDF to write to the message file (`.iqmsg`).

UDFs should manage memory and temporary results as defined by the `EXTFN_V3_API` and `EXTFN_V4_API`.

Sybase IQ is a multiuser application. Many users can simultaneously execute the same UDF. Certain OLAP queries cause a UDF to be executed multiple times within the same query, sometimes in parallel. For details on setting UDFs to run in parallel, see *Aggregate UDF calling patterns* on page 85.

### Internationalization

Sybase IQ has been internationalized for global use. Error messages are in external files, which allows you to localize error messages to new languages without having to make extensive code changes.

To support multiple languages, UDFs should also be internationalized. In general, most UDFs operate on numeric data. In some cases, a UDF may accept string keywords as one or more of

## Understanding User-Defined Functions

the parameters. Place these keywords in external files, in addition to any exception text and log messages used by the UDF.

Sybase IQ has also been localized for a few non-English languages. To support localization to the same languages that Sybase IQ supports, Sybase recommends that you internationalize UDFs, so that an independent organization can localize them at a later date.

For details about international language support in Sybase IQ, see the *System Administration Guide: Volume 1 > International Languages and Character Sets*.

See also *Debugging Using Cross-Character-Set Maps* at [www.Sybase.com](http://www.Sybase.com). This paper discusses how to debug with multi byte data, as opposed to input keywords, exception messages, and log entries.

### Platform Differences

Develop UDFs to run on a variety of platforms supported by Sybase IQ. The Sybase IQ 15.x server runs on 64-bit architectures, and is supported under several platforms of the MS Windows (64-bit) family of operating systems. Sybase IQ is also supported on versions of UNIX (64-bit), including Solaris, HP-UX, AIX, and Linux.

## Practices to Avoid

---

Learn good practices for creating user-defined functions.

- Do not hard-code library paths in SQL registration scripts. This practice makes it difficult to provide flexibility to the user to install the UDFs into the same directory as Sybase IQ.
- Do not write ambiguous code, or constructs that can unexpectedly loop forever, without providing a mechanism for the user to cancel the UDF invocation (see the function `'get_is_cancelled()'`).
- Do not perform complex, or memory-intensive operations that are repeated every invocation. When a UDF call is made against a table that contains many thousands of rows, efficient execution becomes paramount. Sybase recommends that you allocate blocks of memory for a thousand to several thousand rows at a time, rather than on a row-by-row basis.
- Do not open a database connection, or perform database operations from within a UDF. All parameters and data required for UDF execution must be passed as parameters to the UDF.
- Do not use reserved words when naming UDFs.

## Naming Conventions for User-Defined Functions

---

UDF names must follow the same restrictions as other identifiers in Sybase IQ.

Sybase IQ identifiers have a maximum length of 128 bytes. For simplicity of use, UDF names should start with an alphabetic character. Alphabetic characters as defined by Sybase IQ include the letters of the alphabet, plus underscore (\_), at sign (@), number or pound sign (#)



and dollar sign (\$). UDF names should consist entirely of these alphabetic characters as well as digits (the numbers 0 through 9). UDF names should not conflict with SQL reserved words. There is a list of SQL reserved words in *Sybase IQ Reference: Building Blocks, Tables, and Procedures > SQL Language Elements > Reserved Words*.

Although UDF names (as other identifiers) may also contain reserved words, spaces, characters other than those listed above, and may start with a non-alphabetic character, this is not recommended. If UDF names have any of these characteristics, you must enclose them in quotes or square brackets, which makes it more difficult to use them.

The UDFs reside in the same name space as other SQL functions and stored procedures. To avoid conflicts with existing stored procedures and functions, preface UDFs with a unique short (2-letter to 5-letter) acronym and underscore. Choose UDF names that do not conflict with other SQL functions or stored procedures already defined in the local environment.

These are some of the prefixes that are already in use:

- **debugger\_tutorial** – a stored procedure delivered with the native Sybase IQ installation.
- **ManageContacts** – a stored procedure delivered with the Sybase IQ demo database.
- **Show** – stored procedures used to display data from the Sybase IQ demo database.
- **sp\_Detect\_MPX\_DDL\_conflicts** – a stored procedure delivered with the native Sybase IQ installation.
- **sp\_iqevbegintxn** – a stored procedure delivered with the native Sybase IQ installation.
- **sp\_iqmpx** – functions and stored procedures provided by Sybase IQ to assist in multiplex administration.
- **ts\_** – optional financial time series and forecasting functions.

## SQL Data Types

---

UDF declarations support only certain SQL data types.

You can use the following SQL data types in a UDF declaration, either as data types for arguments to a UDF, or as return-value data types:

SQL Data Type	C or C++ Data Type Identifier	C or C++ Typedef	Description
UNSIGNED BIGINT	DT_UN-SBIGINT	a_sql_uint64	An unsigned 64-bit integer, requiring 8 bytes of storage.
BIGINT	DT_BIGINT	a_sql_int64	A signed 64-bit integer, requiring 8 bytes of storage.

## Understanding User-Defined Functions

SQL Data Type	C or C++ Data Type Identifier	C or C++ Typedef	Description
UNSIGNED INT	DT_UN-SINT	a_sql_uint32	An unsigned 32-bit integer, requiring 4 bytes of storage.
INT	DT_INT	a_sql_int32	A signed 32-bit integer, requiring 4 bytes of storage.
SMALLINT	DT_SMALLINT	short	A signed 16-bit integer, requiring 2 bytes of storage.
TINYINT	DT_TINYINT	unsigned char	An unsigned 8-bit integer, requiring 1 byte of storage.
DOUBLE	DT_DOUBLE	double	A signed 64-bit double-precision floating point number, requiring 8 bytes of storage.
REAL	DT_FLOAT	float	A signed 32-bit floating point number, requiring 4 bytes of storage.
FLOAT	DT_FLOAT	float	In SQL, depending on the associated precision, a FLOAT is either a signed 32-bit floating point number requiring 4 bytes of storage, or a signed 64-bit double-precision floating point number requiring 8 bytes of storage. You can use the SQL data type FLOAT only in a UDF declaration if the optional precision for FLOAT data types is not supplied. Without a precision, FLOAT is a synonym for REAL.
CHAR(<n>)	DT_FIXED-CHAR	char	A fixed-length blank-padded character string, in the database default character set. The maximum possible length, "<n>", is 32767. The data is not null-byte terminated.
VARCHAR(<n>)	DT_VARIABLE-CHAR	char	A varying-length character string, in the database default character set. The maximum possible length, "<n>", is 32767. The data is not null-byte terminated. For UDF input arguments, the actual length, when the value is not NULL, must be retrieved from the <i>total_len</i> field within the <i>an_extfn_value</i> structure. Similarly, for a UDF result of this type, the actual length must be set in the <i>total_len</i> field.

SQL Data Type	C or C++ Data Type Identifier	C or C++ Typedef	Description
LONG VARCHAR(<n>) or CLOB	DT_VARCHAR	char	<p>A varying-length character string, in the database default character set. Use the LONG VARCHAR data type only as an input argument, not as a return-value data type. The maximum possible length, “&lt;n&gt;”, is 4GB (gigabytes) for v3 UDFs. The data is not null-byte terminated. LONG VARCHAR data type can have a WD or TEXT index. For UDF input arguments, the actual length, when the value is not NULL, must be retrieved from the <i>total_len</i> field within the <i>an_extfn_value</i> structure.</p> <p>You need not rebuild or recompile an existing scalar or aggregate UDF to use a LOB data type as an input parameter, if the function contains a loop that reads pieces of the value via the <code>get_value()</code> and <code>get_piece()</code> methods. The loop continues until <i>remain_len</i> &gt; 0 or until 4GB is reached for v3 UDFs (there is no 4GB limit in v4).</p> <p>Table UDFs and TPFs do not use the <code>get_piece()</code> method to process and retrieve data. Table UDFs and TPFs must use the <code>Blob(a_v4_extfn_blob)</code> API instead. Use <code>blob_length</code> to determine length of input parameters.</p> <p>Large object data support requires a separately licensed Sybase IQ option.</p>
BINARY(<n>)	DT_BINARY	unsigned char	A fixed-length null-byte padded binary, value with a maximum possible binary length, “<n>”, of 32767. The data is not null-byte terminated.
VARBINARY(<n>)	DT_BINARY	unsigned char	A varying-length binary value, for which the maximum possible length, “<n>”, is 32767. The data is not null-byte terminated. For UDF input arguments, the actual length, when the value is not NULL, must be retrieved from the <i>total_len</i> field within the <i>an_extfn_value</i> structure. Similarly, for a UDF result of this type, you must set the actual length in the <i>total_len</i> field. The data is not null-byte terminated.

SQL Data Type	C or C++ Data Type Identifier	C or C++ Typedef	Description
LONG BINARY(<n>) or BLOB	DT_BINARY	unsigned char	<p>A fixed-length null-byte padded binary, value with a maximum possible binary length, “&lt;n&gt;”, of 4GB (gigabytes) for v3 UDFs. Use the LONG BINARY data type only as an input argument, not as a return-value data type.</p> <p>You need not rebuild or recompile an existing scalar or aggregate UDF to use a LOB data type as an input parameter, if the function contains a loop that reads pieces of the value via the <code>get_value()</code> and <code>get_piece()</code> methods. The loop continues until <i>remain_len</i> &gt; 0 or until 4GB is reached for v3 UDFs (there is no 4GB limit in v4).</p> <p>Table UDFs and TPFs do not use the <code>get_piece()</code> method to process and retrieve data. Table UDFs and TPFs must use the <code>Blob(a_v4_extfn_blob)</code> API instead. Use <code>blob_length</code> to determine length of input parameters.</p> <p>Large object data support requires a separately licensed Sybase IQ option.</p>

SQL Data Type	C or C++ Data Type Identifier	C or C++ Typedef	Description
DATE	DT_TIME-STAMP_STRUCT	unsigned integer	<p>A calendar date value, which is passed to or from a UDF as an unsigned integer. The value given to the UDF is guaranteed to be usable in comparison and sorting operations. A larger value indicates a later date. If the actual date components are required, the UDF must invoke the <code>convert_value</code> function in order to convert to the type <code>DT_TIME-STAMP_STRUCT</code>. This date type represents date and time with this structure:</p> <pre>typedef struct sqldatetime {     unsigned short year;          /* e.g. 1992*/     unsigned char month;         /* 0-11 */     unsigned char day_of_week;  /* 0-6  0=Sunday, 1=Monday, ... */     unsigned short day_of_year; /* 0-365 */     unsigned char day;          /* 1-31 */     unsigned char hour;         /* 0-23 */     unsigned char    mi- nute;         /* 0-59 */     unsigned char    sec- ond;         /* 0-59 */     a_sql_uint32    microsec- ond;        /* 0-999999 */ } SQLDATETIME;</pre>
	DT_TIME-STAMP_STRUCT	unsigned bigint	<p>A value that precisely describes a moment within a given day. The value given to the UDF is guaranteed to be usable in comparison and sorting operations. A larger value indicates a later time. If the actual time components are required, the UDF must invoke the <code>convert_value</code> function to convert to the type <code>DT_TIME-STAMP_STRUCT</code>.</p>

SQL Data Type	C or C++ Data Type Identifier	C or C++ Typedef	Description
DATETIME, SMALLDATE-TIME, or TIMESTAMP	DT_TIME- STAMP_S STRUCT	unsigned big- int	A calendar date and time value. The value given to the UDF is guaranteed to be usable in comparison and sorting operations. A larger value indicates a later datetime. If the actual time components are required, the UDF must invoke the <code>convert_value</code> function to convert to the type <code>DT_TIMESTAMP_STRUCT</code> .
TABLE	DT_EXTF N_TABLE	a_v4_extfn_ table	Represents an input table parameter result set. This datatype is only available on TPFs.

*Unsupported Data Types*

You cannot use the following SQL data types in a UDF declaration, either as data types for arguments to a UDF, or as return-value data types:

- BIT – should typically be handled in the UDF declaration as a TINYINT data type, and then the implicit data type conversion from BIT automatically handles the value translation.
- DECIMAL(<precision>, <scale>) or NUMERIC(<precision>, <scale>) – depending on the usage, DECIMAL is typically handled as a DOUBLE data type, but various conventions may be imposed to enable the use of INT or BIGINT data types.
- LONG VARCHAR (CLOB) – supported only as an input argument, not as a return-value data type. An exception exists for pass-through TPFs, where LONG VARCHAR is supported as a return-value data type.
- LONG BINARY (BLOB) – supported only as an input argument, not as a return-value data type. An exception exists for pass-through TPFs, where LONG BINARY is supported as a return-value data type.
- TEXT – not currently supported.

**See also**

- *Blob (a\_v4\_extfn\_blob)* on page 195
- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199
- *convert\_value* on page 295
- *Table (a\_v4\_extfn\_table)* on page 306

# Building UDFs

Design, build, and test UDFs.

## Design Basics of User-Defined Functions

---

There are some basic considerations to keep in mind while developing UDFs.

This document assumes that the UDF developer is familiar with the basics of developing software, including good program design and development and independent testing.

In addition to standard software development practices, developers of non-Java UDFs should remember that they are developing code to be executed within the Sybase IQ database container, and to understand the limitations imposed by the database container.

Developers of aggregate UDFs should also be familiar with OLAP queries, and how they translate into UDF calling patterns.

Because the UDFs may be invoked by several threads simultaneously, they must be constructed to be thread-safe.

### Sample Code

Sample UDF source code is delivered with the product. The newest version of the sample code is always delivered with the most current version of Sybase IQ.

On UNIX platforms, the sample UDF code is in `$$SYBASE/IQ-15_4/samples/udf` (where `$$SYBASE` is the installation root).

On Windows platforms, the sample UDF code is in `C:\Documents and Settings\All Users\SybaseIQ\samples\udf`.

The sample UDF code documented in the *User-Defined Functions Guide* may not be the latest version as delivered with the Sybase IQ product. Last-minute changes to the sample UDF source code are documented in the Sybase IQ Release Bulletin for your operating system platform.

## Setting the Dynamic Library Interface

---

Specify the interface style to be used in the dynamically linkable library.

Each dynamically loaded library must contain exactly one copy of this definition:

```
extern "C" a_sql_uint32 extfn_use_new_api(void )
{
```

```
    return EXTFN_V4_API;  
}
```

This definition informs the server of which interface style is being used, and therefore how to access the UDFs defined in this dynamically linkable library. For high-performance UDFs, only new interface styles `EXTFN_V3_API` and `EXTFN_V4_API` are supported.

## Upgrading to the v4 API

---

Sybase recommends that you upgrade to the v4 API included with 15.4.

### Prerequisites

Install Sybase IQ server version 15.4.

### Task

If you have existing scalar or aggregate UDFs developed for Sybase IQ server versions 15.1, 15.2, or 15.3, those UDFs use the V3 API interface style and reference the `extfnapi3.h` header file. Modify your legacy C or C++ external library files to reference the `extfnapi4.h` header file.

Existing v3 scalar and aggregate functions continue to work as designed. However, to take advantage of scalar and aggregate distribution in PlexQ, you must upgrade the header file and library version to v4. You need not change the name of the typedefs for your scalar or aggregate function.

1. Open the C or C++ external library file defining the scalar or aggregate user-defined function.
2. Locate all instances of `#include 'extfnapi3.h'` and change to `#include 'extfnapi4.h'`.
3. Set the dynamic library interface to `EXTFN_V4_API`.
4. Rebuild.

### Next

Sybase partners must ensure the library exports `extfn_get_license_info` as an entry point.

### See also

- *External Function Prototypes* on page 94
- *License Information (a\_v4\_extfn\_license\_info)* on page 301
- *Defining an Aggregate UDF* on page 54
- *Defining a Scalar UDF* on page 38
- *Developing a Table UDF* on page 103



- *Developing a TPF* on page 136

## Library Version (`extfn_get_library_version`)

---

Use the `extfn_get_library_version` method to extract the library version from the current multiplex node. The server considers partitioning a query across multiplex nodes only if the installed library is compatible with the other nodes.

### *Implementation*

A v4 library can define this optional entry point:

```
size_t extfn_get_library_version( uint8 *buff, size_t len );
```

### *Description*

Library versioning methods are at the library level, and do not have the `a_v4` prefix in their method name.

If the v4 library defines the optional entry point, the server allows query distribution to other nodes. The entry point populates the supplied buffer with the library version string (a C-style character string containing only ASCII characters, terminated with **\0**) and returns the actual size of the populated version string, which is constrained to a maximum of 256 bytes.

If an entry point is not defined, the server does not distribute the UDF to the other nodes in the multiplex.

### **See also**

- *Library Version Compatibility (`extfn_check_version_compatibility`)* on page 17
- *Setting the Dynamic Library Interface* on page 15

## Library Version Compatibility (`extfn_check_version_compatibility`)

---

Use the `extfn_check_version_compatibility` method to define compatibility criteria for library versions across nodes in a multiplex.

### *Implementation*

A v4 library can define this optional entry point:

```
a_bool extfn_check_version_compatibility( uint8 *buff, size_t len );
```

### *Description*

Library versioning methods are at the library level, and do not have the `a_v4` prefix in their method name.

## Building UDFs

This optional entry point accepts a buffer containing the version string and the version string length. It returns whether or not the library version on the target node is compatible with the version string parameter. The library developer defines the compatibility criteria.

### *Interaction with `extfn_get_library_version`*

The leader node calls `extfn_get_library_version` before checking version compatibility. If `extfn_get_library_version` is not implemented on the leader node, then there is no distribution. If `extfn_get_library_version` is implemented on the leader node, then the UDF or TPF is eligible for distribution. Being eligible for distribution is not a guarantee that distributed query processing will occur.

The `extfn_get_library_version` method can return a 0-length string; however, this does not mean that `extfn_get_library_version` is not implemented.

---

**Note:** A TPF or UDF is still eligible for distribution if `extfn_get_library_version` returns a 0-length string.

---

If `extfn_get_library_version` returns a 0-length string, whether or not the worker node accepts the distributed work depends on the `extfn_check_version_compatibility` implementation on the worker node. A worker node requires a compatible library to process distributed work.

### **See also**

- *Library Version (`extfn_get_library_version`)* on page 17
- *Setting the Dynamic Library Interface* on page 15

## **License Information (`extfn_get_license_info`)**

---

If you are a Sybase design partner, implement the `extfn_get_license_info` library-level function to enable the server to obtain licensing information from a v4 UDF.

### *Data Type*

`an_extfn_license_info`

### *Implementation*

```
(_entry an_extfn_get_license_info) ( an_extfn_license_info  
**license_info );
```

### *Parameters*

**license\_info** is an output parameter that returns the license information as received from the library. You define the license information in the `a_v4_extfn_license_info` structure.

*Description*

Sybase partners must specify the Sybase-supplied license key in the `a_v4_extfn_license_info` structure, and must ensure that the library exports `extfn_get_license_info` as an entry point.

**Adding the `extfn_get_license_info` Method**

If you are a Sybase design partner, populate strings in `a_v4_extfn_license_info` and define `extfn_get_license_info` as a v4 entry point.

1. In the `a_v4_extfn_license_info` structure, specify your company name. The maximum length is 255 characters.
2. In the `a_v4_extfn_license_info` structure, specify additional library information such as library version and build numbers. The maximum length is 255 characters.
3. In the `a_v4_extfn_license_info` structure, enter the license key provided by Sybase.
4. Ensure the library exports `extfn_get_license_info` as an entry point.

```
a_v4_extfn_license_info my_info = {
    1,
    "Company Name",
    "Library Info String",
    (void *)"KEY_STRING"
};

void SQL_CALLBACK extfn_get_license_info( an_extfn_license_info
**license_info )
/
*****
*****/
{
    *license_info = (an_extfn_license_info *)& my_info;
}
```

**Compile and Link Source Code to Build Dynamically Linkable Libraries**

Use compile and link switches when building dynamically linkable libraries for any user-defined function.

1. A UDF dynamically linkable library must include an implementation of the function `extfn_use_new_api()`. The source code for this function is in *Setting the dynamic library interface* on page 15. This function informs the server of the API style that all functions in the library adhere to. The sample source file `my_main.cxx` contains this function; you can use it without modification.

## Building UDFs

2. A UDF dynamically linkable library must also contain object code for at least one UDF function. A UDF dynamically linkable library may optionally contain multiple UDFs.
3. Link together the object code for each UDF as well as the **extfn\_use\_new\_api()** to form a single library.

For example, to build the library "libudfex:"

- Compile each source file to produce an object file:

```
my_main.cxx
my_bit_or.cxx
my_bit_xor.cxx
my_interpolate.cxx
my_plus.cxx
my_plus_counter.cxx
my_sum.cxx
my_byte_length.cxx
my_md5.cxx
my_toupper.cxx
tpf_agg.cxx
tpf_blob.cxx
tpf_dt.cxx
tpf_filt.cxx
tpf_oby.cxx
tpf_pby.cxx
tpf_rg_1.cxx
tpf_rg_2.cxx
udf_blob.cxx
udf_main.cxx
udf_rg_1.cxx
udf_rg_2.cxx
udf_rg_3.cxx
udf_utils.cxx
```

- Link together each object produced into a single library.

After the dynamically linkable library has been compiled and linked, complete one of these tasks:

- (Recommended) update the **CREATE FUNCTION ... EXTERNAL NAME** or **CREATE PROCEDURE ... EXTERNAL NAME** to include an explicit path name for the UDF library.
- Place the UDF library file into the directory where all the Sybase IQ libraries are stored.
- Start the server with a library load path that includes the location of the UDF library. On UNIX modify the LD\_LIBRARY\_PATH within the start\_iq startup script. While LD\_LIBRARY\_PATH is universal to all UNIX variants, SHLIB\_PATH is preferred on HP, and LIB\_PATH is preferred on AIX. On UNIX platforms, the external name specification can contain a fully qualified name, in which case the LD\_LIBRARY\_PATH is not used. On the Windows platform, a fully qualified name cannot be used and the library search path is defined by the PATH environment variable.

4. Run `iqdir15/samples/udf/build.bat` on Windows. Run `iqdir15/samples/udf/build.sh` on UNIX.

## Compiling and Linking the Sample UDFs for Windows

Run the `build.bat` script to compile and link the sample scalar and aggregate UDFs, table UDFs, and TPFs found in the `samples\udf` directory.

1. Navigate to `%ALLUSERSPROFILE%\samples\udf`.
2. Run `build.bat`:

Parameter	Description
<code>-clean</code>	Deletes the object and the build directory
<code>-v3</code>	Builds sample scalar and aggregate UDFs with the v3 API
<code>-v4</code>	(Default) Builds sample table UDFs and TPFs with the v4 API

## Compiling and Linking the Sample UDFs for UNIX

Run the `build.sh` script to compile and link the sample scalar and aggregate UDFs, table UDFs, and TPFs found in the `samples/udf` directory.

1. Navigate to `$IQDIR15/samples/udf`.
2. Run `build.sh`:

Parameter	Description
<code>-clean</code>	Deletes the object and the build directory
<code>-v3</code>	Builds sample scalar and aggregate UDFs with the v3 API
<code>-v4</code>	(Default) Builds sample table UDFs and TPFs with the v4 API

## AIX Switches

Use the following compile and link switches when building shared libraries on AIX.

*x1C 10.0 on a PowerPC*

---

**Important:** Include the code for `extfn_use_new_api()` in each UDF library.

---

**Note:** To compile on AIX 6.1 systems, the minimum level of the x1C compiler is 10.0.

---

**compile switches**

## Building UDFs

```
-q64 -qarch=ppc64 -qtbtable=full -qsrcmsg -qalign=natural -  
qnoansialias  
-qmaxmem=-1 -qenum=int -qhalt=e -qflag=w -qthreaded -  
qxflags=NLOOPING  
-qtplinst=none -qthreaded
```

### link switches

```
-brtl -G -lg -lpthreads_compat -lpthreads -lm_r -ldl -bnolibpath -  
v
```

## HP-UX Switches

Use the following compile and link switches when building shared libraries on HP-UX.

*aCC 6.24 on Itanium*

---

**Important:** Include the code for `extfn_use_new_api()` in each UDF library.

---

### compile switches

```
+noeh -ext +W740,749,829 +W1031 +DD64 +DSblended +FPD -Aa +ub  
-U_HP_INSTANTIATE_T_IN_LIB -Wc,-ansi_for_scope,on -mt -z
```

### link switches

```
-b -Wl,+s
```

## Linux Switches

Use the following compile and link switches when building shared libraries on Linux.

*g++ 4.1.1 on x86*

---

**Important:** Include the code for `extfn_use_new_api()` in each UDF library.

---

### compile switches

```
-fPIC -fsigned-char -fno-exceptions -pthread -fno-omit-frame-  
pointer  
-Wno-deprecated -Wno-ctor-dtor-privacy -O2 -Wall
```

---

**Note:** When compiling C++ applications for building shared libraries on Linux, adding the **-O2** and **-Wall** switches to the list of compile UDF switches decreases computation time.

---

### link switches

```
-ldl -lnsl -lm -lpthread -shared -Wl,-Bsymbolic -Wl,-shared
```

---

**Note:** You can use `gcc` on Linux as well. While linking with `gcc`, link in the C++ run time library by adding `-lstdc++` to the link switches.

---

### Examples

- *Example 1*

```
g++ -c my_interpolate.cxx -fPIC -fsigned-char -fno-exceptions -  
pthread
```

```
-fno-omit-frame-pointer -Wno-deprecated -Wno-ctor-dtor-
privacy
-I${IQDIR15}/sdk/include/
```

- *Example 2*

```
g++ -c my_main.cxx -fPIC -fsigned-char -fno-exceptions -pthread
-fno-omit-frame-pointer -Wno-deprecated -Wno-ctor-dtor-
privacy
-I${IQDIR15}/sdk/include/
```

- *Example 3*

```
ld -G my_main.o my_interpolate.o -ldl -lnsl -lm -lpthread -shared
-o my_udf_library.so
```

### *xIC 10.0 on a PowerPC*

#### **compile switches**

```
-q64 -qarch=ppc64 -qcheck=nullptr -qinfo=gen -qtbtable=full -
qsrcmsg
-qnoansialias -qminimaltoc -qmaxmem=-1 -qenum=int -qhalt=e -qflag=w
-qthreaded
-qxflags=NLOOPING -qtmplinst=none
```

#### **link switches**

```
-qmkshrobj -ldl -lg -qthreaded -lnsl -lm
```

## **Solaris Switches**

Use the following compile and link switches when building shared libraries on Solaris.

### *Sun Studio 12 on SPARC*

---

**Important:** Include the code for `extfn_use_new_api()` in each UDF library.

---

#### **compile switches**

```
-mt -noex +w -KPIC -i -instances=explicit -V -xtarget=ultra3cu -m64
-xlibmopt
-xlibmil -features=no%conststrings
-erroff=truncwarn,nokeyworddefine,diffenumtype
```

#### **link switches**

```
-z defs -G -ldl -lnsl -lsocket -ladm -lposix4 -lCrun -lCstd -lc -lm
-lefi
-liostream -lkstat
```

### *Sun Studio 12 on x86*

#### **compile switches**

```
+w2 -m64 -features=no%conststrings
-erroff=truncwarn,nokeyworddefine,diffenumtype,doubunder -errtags -
mt -noex
-KPIC -instances=explicit -xlibmopt -xlibmil
```

#### **link switches**

## Building UDFs

```
-z defs -G -ldl -lnsl -lsocket -ladm -lposix4 -lCrun -lCstd -lc -lm  
-lefi  
-liostream -lkstat -m64
```

## Windows Switches

Use the following compile and link switches when building shared libraries on Windows.

*Visual Studio 2008 on x86*

---

**Important:** Include the code for `extfn_use_new_api()` in each UDF library.

---

### compile and link switches

This example is for a DLL containing the `my_plus` function. You must include an `EXPORT` switch for the descriptor function for each UDF contained in the DLL.

```
cl /Zi /LD /I includefilepath my_main.cxx my_plus.cxx /link /  
map  
/INCREMENTAL:NO -EXPORT:extfn_use_new_api -EXPORT:my_plus /  
out:libiqudfex.dll
```

### Example

#### *Environment setup*

```
set VCBASE=c:\dev\vc9  
set MSSDK=C:\dev\mssdk6.0a  
set IQINSTALLDIR=C:\Sybase\IQ  
set OBJ_DIR=%IQINSTALLDIR%\IQ-15_4\samples\udf\objs  
set SRC_DIR=%IQINSTALLDIR%\IQ-15_4\samples\udf\src  
call %VCBASE%\VC\bin\vcvars32.bat
```

- *Example 1*

```
%VCBASE%\VC\bin\amd64\cl -c -nologo -DNDEBUG -DWINNT -D_USRDLL  
-D_WINDLL -D_WIN64 -DWIN64 -  
D_WIN32_WINNT=_WIN32_WINNT_WINXP  
-DWINVER=_WIN32_WINNT_WINXP -D_MBCS -GS -W3 -Zi -favor:AMD64  
-DSYB_LP64 -D_LARGEFILE_SOURCE -D_FILE_OFFSET_BITS=64 -  
DHMSWNT  
-D_CRT_SECURE_NO_DEPRECATED -D_CRT_NONSTDC_NO_DEPRECATED  
-DPOINTERS_ARE_64BITS -DLONG_IS_64BITS -  
D_RWSTD_NO_EXCEPTIONS  
-I"%VCBASE%\VC\include" -I"%MSSDK%\include" -I"%MSSDK%\Lib  
\AMD64"  
-I"%VCBASE%\VC\lib\amd64" -DMSDCXX -DINT64_WORKAROUND  
-DSUPPORTS_UDAF -Od -Zi -MD -I"%IQINSTALLDIR%\IQ-15_4\sdk  
\include"  
-Fo"%OBJ_DIR%\my_interpolate.o" %SRC_DIR%\my_interpolate.cxx
```

- *Example 2*

```
%VCBASE%\VC\bin\amd64\cl -c -nologo -DNDEBUG -DWINNT -D_USRDLL  
-D_WINDLL -D_WIN64 -DWIN64 -  
D_WIN32_WINNT=_WIN32_WINNT_WINXP  
-DWINVER=_WIN32_WINNT_WINXP -D_MBCS -GS -W3 -Zi -favor:AMD64  
-DSYB_LP64 -D_LARGEFILE_SOURCE -D_FILE_OFFSET_BITS=64 -
```



```
DHMSWNT
-D_CRT_SECURE_NO_DEPRECATED -D_CRT_NONSTDC_NO_DEPRECATED
-DPOINTERS_ARE_64BITS -DLONG_IS_64BITS -
D_RWSTD_NO_EXCEPTIONS
-I"%VCBASE%\VC\include" -I"%MSSDK%\include" -I"%MSSDK%\Lib
\AMD64"
-I"%VCBASE%\VC\lib\amd64" -DMSDCXX -DINT64_WORKAROUND
-DSUPPORTS_UDAF -Od -Zi -MD -I"%IQINSTALLDIR%\IQ-15_4\sdk
\include"
-Fo"%OBJ_DIR%\my_main.o" %SRC_DIR%\my_main.cxx
```

- *Example 3*

```
%VCBASE%\VC\bin\amd64\link /LIBPATH:%VCBASE%\VC\lib\amd64
/LIBPATH:%MSSDK%\lib\bin64 kernel32.lib -manifest -DLL -
nologo
-MAP:"%OBJ_DIR%\libudfex.map_deco" /OUT:"%OBJ_DIR%
\libudfex.dll"
"%OBJ_DIR%\my_interpolate.o" "%OBJ_DIR%\my_main.o" /DLL
-EXPORT:extfn_use_new_api -EXPORT:my_interpolate
```

- *Example 4*

```
%MSSDK%\bin\mt -nologo -manifest "%OBJ_DIR%
\libudfex.dll.manifest"
-outputresource:"%OBJ_DIR%\libudfex.dll;2"
```

## Testing User-Defined Functions

---

After UDF external code has been coded, compiled and linked, and the corresponding SQL functions and stored procedures have been defined, the UDFs are ready to be tested.

The reliability required by a database is extremely high. UDFs running within a database environment must maintain this high level of reliability. With the first implementation of the UDF API, UDFs run within the Sybase IQ server. If a UDF aborts prematurely or unexpectedly, the Sybase IQ server may abort. Ensure via thorough testing in a development or test environment, that UDFs do not terminate prematurely or abort unexpectedly under any circumstances.

## Enabling and Disabling User-Defined Functions

Use the `inmemory_external_procedure` security feature to enable or disable the server's ability to make use of high performance in-process UDFs.

A database should maintain data integrity. Under no circumstances should data be lost, modified, augmented, or corrupted. Since UDF execution happens within the Sybase IQ server, there is a risk of corrupting data; practice caution with memory management and any other use of pointers. Sybase strongly recommends that you install and execute UDFs within a read-only multiplex node. For additional protection, use the secured feature (**-sf**) startup option with each server to enable or disable the execution of UDF.

---

**Note:** By default, UDF execution on a multiplex writer and coordinator nodes is disabled. All other nodes are enabled by default.

---

## Building UDFs

Administrators can enable v3 and v4 UDFs for any server by specifying this in the server startup command or in the configuration file:

```
-sf -inmemory_external_procedure
```

Administrators can disable v3 and v4 UDFs for any server by specifying this in the server startup command or in the configuration file:

```
-sf inmemory_external_procedure
```

Additional information on the **-sf** flag is available in *SQL Anywhere Server - Database Administration*. The values listed in the SQL Anywhere document are not applicable to Sybase IQ and should not be used.

## Initially Executing a User-Defined Function

To ensure the safest environment possible, Sybase strongly recommends that you install and invoke UDFs from a read-only server node in a multiplex installation.

The Sybase IQ server does not load the library containing the UDF code until the first time the UDF is invoked. The first execution of a UDF residing in a library that has not yet been loaded may be unusually slow. After the library is loaded, the subsequent invocation of the same UDF or another UDF contained in the same library have the expected performance.

- **Libraries using the stored procedure SA\_EXTERNAL\_LIBRARY\_UNLOAD** – These libraries are not reloaded when the Sybase IQ server is stopped and restarted.

In environments where after-hours maintenance operations require a shutdown and restart of the IQ server, run some test queries after the server has been restarted. This ensures that the appropriate libraries are loaded in memory for optimal query performance during business hours.

## Managing External Libraries

Each external library is loaded the first time a UDF that requires it is invoked. A loaded library remains loaded for the life of the server. It is not loaded when a **CREATE FUNCTION** or **CREATE PROCEDURE** call is made, nor is it automatically unloaded when a **DROP FUNCTION** or **DROP PROCEDURE** call is made.

If the library version must be updated, the **dbo.sa\_external\_library\_unload** procedure forces the library to be unloaded without restarting the server. The call to unload the external library is successful only if the library in question is not currently in use. The procedure takes one optional parameter, a long varchar, that specifies the name of the library to be unloaded. If no parameter is specified, all external libraries not in use are unloaded.

---

**Note:** Unload existing libraries from a running Sybase IQ server before replacing the dynamic link library. The server may fail, if you do not unload the library. Before replacing a dynamically linkable library, either shut down the Sybase IQ server or use the **sa\_external\_library\_unload** function to unload the library.

---

For Windows, unload an external function library using:

```
call sa_external_library_unload('library.dll')
```

For UNIX, unload an external function library using:

```
call sa_external_library_unload('library.so')
```

If a registered function uses a complete path, for example, /abc/def/library, first unregister the function.

In Windows, use

```
call sa_external_library_unload('\abc\def\library.dll')
```

In UNIX, use

```
call sa_external_library_unload('/abc/def/library.so')
```

---

**Note:** The library path is required in the SQL function declaration only if the library is not already located within a directory in the library load path.

---

## Controlling Error Checking and Call Tracing

The **external\_UDF\_execution\_mode** option controls the amount of error checking and call tracing that is performed when statements involving v3 and v4 external user-defined functions are evaluated.

You can use **external\_UDF\_execution\_mode** during development of a UDF to aid in debugging while you are developing UDFs.

### *Allowed Values*

0, 1, 2

### *Default Value*

0

### *Scope*

Can be set as public, temporary, or user.

### *Description*

When set to 0, the default, external UDFs are evaluated in a manner that optimizes the performance of statements using UDFs.

When set to 1, external UDFs are evaluated to validate the information passed back and forth to each UDF. This setting is intended for scalar and aggregate UDFs.

When set to 2, external UDFs are evaluated to not only validate the information passed back and forth to the UDF, but also to log, in the `iqmsgg` file, every call to the functions provided by the UDFs and every callback from those functions back into the server. This setting is intended for all C or C++ external UDFs. Memory tracing is turned on for table UDFs and TPFs.

## Viewing Sybase IQ Log Files

Sybase IQ provides extensive logging and tracing capabilities. UDFs should provide the same or better level of detailed logging, in the event of problems in the UDF code.

Log files for the database are generally located with the database file and configuration file. On UNIX platforms, there are two files named after the database instance, one with a `.stderr` extension and one with a `.stdout` extension. On Windows, by default, the `stderr` file is not generated.

To capture the `stderr` messages along with the `stdout` messages under Windows, redirect the `stdout` and `stderr`:

```
iqsrv15.exe @iqdemo.cfg iqdemo.db 2>&1 > iqdemo.stdout
```

The Windows output messages are slightly different from the output messages generated on UNIX platforms.

## Using Microsoft Visual Studio Debugger for User-Defined Functions

Microsoft Visual Studio 2008 developers use Microsoft Visual Studio Debugger to step through the user-defined function code.

1. Attach the debugger to a running server:

```
devenv /debugexe "%IQDIR15%\Bin64\iqsrv15.exe"
```

2. Goto Debug | Attach to Process

3. To start the server and debugger together:

```
devenv /debugexe "%IQDIR15%\bin32\iqsrv15.exe" [commandline  
options for your server]
```

Each platform will have a debugger and each will have their own command line syntax. Sybase IQ source code is not required. The `msvs` debugger will recognize when the user-defined functions source is executed and break at the set breakpoints. When control returns from the user-defined functions to the server, you will only see machine code.

## Modifying the UDF at Runtime

Many Sybase IQ installations are in mission-critical environments, where customers require an extremely high level of availability. System Administrators must be able to install and upgrade UDFs with little or no impact to the Sybase IQ server.

An application must not attempt to access an external library while the associated library file is being moved, overwritten, or deleted. Since libraries are automatically loaded whenever an associated SQL function is invoked, it is important to follow these steps in the exact order whenever performing any type of maintenance on existing UDF libraries:

1. Ensure all users who invoke UDFs do not have any pending queries in progress

2. Revoke the execute permission from users, and drop the SQL functions and stored procedures which reference external UDF code modules
3. Unload the library from the Sybase IQ server, using the **call sa\_external\_library\_unload** command (shutting down the IQ server also automatically unloads the library).
4. Perform the desired maintenance on the external library files (copy, move, update, delete).
5. Edit SQL function and stored procedure definitions in the registration scripts to reflect external library locations, if the libraries were moved.
6. Grant the execute permission to users, and run registration scripts to re-create the SQL functions and stored procedures which reference external UDF code modules.
7. Invoke a SQL function or stored procedure that references the external UDF code to ensure the Sybase IQ server can dynamically load the external library.

## Granting and Revoking Permissions

---

A user-defined function is owned by the user who created it, by default, although the creator can specify another owner at creation time.

The owner can grant permissions to other users using the **GRANT EXECUTE** command. For example, the creator of the function *fullname* can allow *another\_user* to use *fullname* by issuing:

```
GRANT EXECUTE ON fullname TO another_user
```

Or can revoke permissions by issuing:

```
REVOKE EXECUTE ON fullname FROM another_user
```

See *System Administration Guide: Volume 1 > Managing User IDs and Permissions > Managing Individual User IDs and Permissions > Granting Permissions on Procedures in Interactive SQL*.

## Dropping User-Defined Functions

---

Once you create a user-defined function, it remains in the database until it is explicitly removed. Only the owner of the function or procedure, or a user with DBA authority, can drop a function or procedure from the database.

For example, to remove the scalar or aggregate function *fullname* from the database, enter:

```
DROP FUNCTION fullname
```

To remove a table UDF or TPF named *fullname* from the database, enter:

```
DROP PROCEDURE fullname
```



# Scalar and Aggregate UDFs

Scalar and aggregate user-defined functions return a single value to the calling environment.

---

**Note:** Scalar and aggregate UDFs are a licensable option, and require the IQ\_UDF or IQ\_IDA license. Installing the license enables user-defined functions.

---

You can install Sybase IQ in a wide variety of configurations. UDFs must be easily installed within this environment, and must be able to run within all supported configurations. The Sybase IQ installer provides a default installation directory, but allows users to select a different installation directory. UDF developers should consider providing the same flexibility when installing the UDF libraries and associated SQL function definition scripts.

## Scalar and Aggregate UDF Restrictions

---

External C/C++ scalar and aggregate user-defined functions have some restrictions.

- Write all UDFs in a manner that allows them to be called simultaneously by different users while receiving different context functions.
- If a UDF accesses a global or shared data structure, the UDF definition must implement the appropriate locking around its accesses to that data, including the releasing of that locking under all normal code paths and all error handling situations.
- UDFs implemented in C++ may provide overloaded "new" operators for their classes, but they should never overload the global "new" operator. On some platforms, the effect of doing so is not limited to the code defined within that specific library.
- Write all aggregate UDFs and all deterministic scalar UDFs such that the receipt of the same input values always produces the same output values. Any scalar function for which this is not true must be declared as NONDETERMINISTIC to avoid the potential for incorrect answers.
- Users can create a standard SQL function without a DBA authority, but they cannot create a function which will invoke an external library without having DBA permissions. Attempting to do this results in an error message "You do not have permission to use the create function statement."

## Creating a Scalar or Aggregate UDF

---

Learn how to create and configure external C or C++ scalar and aggregate user-defined functions.

For instructions on creating UDFs using Interactive SQL, see *System Administration Guide: Volume 2 > Using Procedures and Batches*.

1. Declare the UDF to the server by using the **CREATE FUNCTION** or **CREATE AGGREGATE FUNCTION** statements. Write and execute these statements as commands, or use the appropriate CREATE statement using the Sybase Central New Function wizard.

The external C/C++ form of the CREATE FUNCTION statement requires DBA or RESOURCE authority, so standard users do not have the authority to declare any UDFs of this type.

2. *Write the UDF library identification function.* on page 15.
3. Define the UDF as a set of C or C++ functions. See *Defining a scalar UDF* on page 38 or *Defining an aggregate UDF* on page 54.
4. Implement the function entry points in C/C++.
5. *Compile the UDF functions and the library identification functions.* on page 19.
6. Link the compiled file into a dynamically linkable library.

Any reference to a UDF in a SQL statement first, if necessary, links the dynamically linkable library. The *calling patterns* on page 83 are then called.

Because these high-performance external C/C++ user-defined functions involve the loading of non-server library code into the process space of the server, there are potential risks to data integrity, data security, and server robustness from poorly or maliciously written functions. To manage these risks, each Sybase IQ server can explicitly *enable or disable this functionality* on page 25.

### **Creating a UDF Using SQL Anywhere Dialects in Sybase Central**

Watcom-SQL and Transact-SQL are SQL dialects supported by SQL Anywhere, and can be used when creating user-defined functions.

1. In Sybase Central, connect to the database as a user with DBA or Resource authority.
2. Select **View > Folders**.
3. In the left pane, right-click **Procedures & Functions** and select **New > Function**.
4. Enter a name for the function and select the user who will own the function.
5. Select the SQL dialect or language for the function. Click **Next**.
6. Select the type of value to be returned in the function, and specify the size, units, and scale for the value.
7. Type a name for the return value and click **Next**.
8. Add a comment describing the purpose of the new function. Click **Finish**.
9. In the right pane, click the **SQL tab** to complete the procedure code.



## **Declaring and Defining Scalar User-Defined Functions**

Sybase IQ supports simple scalar user-defined functions (UDFs) that can be used anywhere the SQRT function can be used.

These scalar UDFs can be deterministic, which means that for a given set of argument values the function always returns the same result value, or they can be nondeterministic scalar functions, which means that the same arguments can return different results.

---

**Note:** The scalar UDF examples referenced in this chapter are installed with the IQ server, and can be found as .cxx files in `$IQDIR15/samples/udf`. You can also find them in the `$IQDIR15/lib64/libudfex` dynamically linkable library.

---

### **Declaring a Scalar UDF**

Only a DBA, or someone with DBA authority can declare an in-process external UDF. There is also a server startup option that allows an administrator to enable or disable this style of user-defined function.

After the UDF code has been written and compiled, create a SQL function that invokes the UDF from the appropriate library file, sending the input data to the UDF.

---

**Note:** You can also *create the user-defined function declaration in Sybase Central* on page 38.

---

By default, all user-defined functions use the access permissions of the owner of the UDF.

---

**Note:** Users are required to have DBA authority in order to declare UDF functions.

---

The syntax for creating a scalar UDF is:

```
scalar-udf-declaration:
CREATE FUNCTION [ owner.]function-name
    ( [ parameter , ... ] )
RETURNS data-type
    [ routine-characteristics ... ]
EXTERNAL NAME library-and-entry-point-name-string

parameter:
    param-name data-type [ DEFAULT value ]

routine-characteristics:
    [NOT] DETERMINISTIC
    | { IGNORE | RESPECT } NULL VALUES
    | SQL SECURITY { INVOKER | DEFINER }
```

The defaults for the characteristics in the above syntax are:

```
DETERMINISTIC
RESPECT NULL VALUES
SQL SECURITY DEFINER
```

To minimize potential security concerns, Sybase recommends that you use a fully qualified path name to a secure directory for the library name portion of the `EXTERNAL NAME` clause.

### SQL Security

Defines whether the function is executed as the `INVOKER`, (the user who is calling the function), or as the `DEFINER` (the user who owns the function). The default is `DEFINER`.

**SQL SECURITY INVOKER** uses additional memory, because each user that calls the procedure requires annotation. Additionally, name resolution is performed on both the user name and the `INVOKER`. Qualify all object names (tables, procedures, and so on) with their appropriate owner.

### External Name

A function using the **EXTERNAL NAME** clause is a wrapper around a call to a function in an external library. A function using **EXTERNAL NAME** can have no other clauses following the **RETURNS** clause. The library name may include the file extension, which is typically `.dll` on Windows and `.so` on UNIX. In the absence of the extension, the software appends the platform-specific default file extension for libraries.

The **EXTERNAL NAME** clause is not supported for temporary functions. See *SQL Anywhere Server - Programming > SQL Anywhere external call interface*.

---

**Note:** This reference points to SQL Anywhere documentation.

---

You can start the server with a library load path that includes the location of the UDF library. On UNIX variants, modify the `LD_LIBRARY_PATH` in the `start_iq` startup script. While `LD_LIBRARY_PATH` is universal to all UNIX variants, `SHLIB_PATH` is preferred on HP, and `LIB_PATH` is preferred on AIX.

On UNIX platforms, the external name specification can contain a fully qualified name, in which case the `LD_LIBRARY_PATH` is not used. On the Windows platform, a fully qualified name cannot be used and the library search path is defined by the `PATH` environment variable.

---

**Note:** Scalar user-defined functions and user-defined aggregate functions are not supported in updatable cursors.

---

### See also

- *Defining a Scalar UDF* on page 38

### UDF Example: my\_plus Declaration

The “my\_plus” example is a simple scalar function that returns the result of adding its two integer argument values.

### my\_plus declaration

When `my_plus` resides within the dynamically linkable library `my_shared_lib`, the declaration for this example looks like this:

```
CREATE FUNCTION my_plus (IN arg1 INT, IN arg2 INT)
  RETURNS INT
  DETERMINISTIC
  IGNORE NULL VALUES
  EXTERNAL NAME 'my_plus@libudfex'
```

This declaration says that `my_plus` is a simple scalar UDF residing in `my_shared_lib` with a descriptor routine named `describe_my_plus`. Since the behavior of a UDF may require more than one actual C/C++ entry point for its implementation, this set of entry points is not directly part of the `CREATE FUNCTION` syntax. Instead, the `CREATE FUNCTION` statement `EXTERNAL NAME` clause identifies a descriptor function for this UDF. A descriptor function, when invoked, returns a descriptor structure that is defined in detail in the next section. That descriptor structure contains the required and optional function pointers that embody the implementation of this UDF.

This declaration says that `my_plus` accepts two `INT` arguments and returns an `INT` result value. If the function is invoked with an argument that is not an `INT`, and if the argument can be implicitly converted into an `INT`, the conversion happens before the function is called. If this function is invoked with an argument that cannot be implicitly converted into an `INT`, a conversion error is generated.

Further, the declaration states that this function is deterministic. A deterministic function always returns the identical result value when supplied the same input values. This means the result cannot depend on any external information beyond the supplied argument values, or on any side effects from previous invocations. By default, functions are assumed to be deterministic, so the results are the same if this characteristic is omitted from the `CREATE` statement.

The last piece of the above declaration is the `IGNORE NULL VALUES` characteristic. Nearly all built-in scalar functions return a `NULL` result value if any of the input arguments are `NULL`. The `IGNORE NULL VALUES` states that the `my_plus` function follows that convention, and therefore this UDF routine is not actually invoked when either of its input values are `NULL`. Since `RESPECT NULL VALUES` is the default for functions, this characteristic must be specified in the declaration for this UDF to get the performance benefits. All functions that may return a non-`NULL` result given a `NULL` input value must use the default `RESPECT NULL VALUES` characteristic.

In the following example query, `my_plus` appears in the `SELECT` list along with the equivalent arithmetic expression:

```
SELECT my_plus(t.x, t.y) AS x_plus_y_one, (t.x + t.y) AS x_plus_y_two
FROM t
WHERE t.z = 2
```

In the following example, `my_plus` is used in several different places and different ways within the same query:

```
SELECT my_plus(t.x, t.y), count(*)
FROM t
WHERE t.z = 2
```

## Scalar and Aggregate UDFs

```
AND my_plus(t.x, 5) > 10
AND my_plus(t.y, 5) > 10
GROUP BY my_plus(t.x, t.y)
```

### UDF Example: my\_plus\_counter Declaration

The `my_plus_counter` example is a simple nondeterministic scalar UDF that takes a single integer argument, and returns the result of adding that argument value to an internal integer usage counter. If the input argument value is `NULL`, the result is the current value of the usage counter.

### **my\_plus\_counter declaration**

Assuming that `my_plus_counter` also resides within the dynamically linkable library `my_shared_lib`, the declaration for this example is:

```
CREATE FUNCTION my_plus_counter (IN arg1 INT DEFAULT 0)
  RETURNS INT
  NOT DETERMINISTIC
  RESPECT NULL VALUES
  EXTERNAL NAME 'describe_my_plus_counter@my_shared_lib'
```

The `RESPECT NULL VALUES` characteristic means that this function is called even if the input argument value is `NULL`. This is necessary because the semantics of `my_plus_counter` includes:

- Internally keeping a usage count that increments even if the argument is `NULL`.
- A non-null value result when passed a `NULL` argument.

Because `RESPECT NULL VALUES` is the default, the results are the same if this clause is omitted from the declaration.

Sybase IQ restricts the usage of all nondeterministic functions. They are allowed only within the `SELECT` list of the top-level query block or in the `SET` clause of an `UPDATE` statement. They cannot be used within subqueries, or within a `WHERE`, `ON`, `GROUP BY`, or `HAVING` clause. This restriction applies to nondeterministic UDFs as well as to the nondeterministic built-in functions like `GETUID` and `NUMBER`.

The last detail in the above declaration is the `DEFAULT` qualifier on the input parameter. The qualifier tells the server that this function can be called with no arguments, and that when this happens the server automatically supplies a zero for the missing argument. If a `DEFAULT` value is specified, it must be implicitly convertible into the data type of that argument.

In the following example, the first `SELECT` list item adds the running counter to the value of `t.x` for each row. The second and third `SELECT` list items each return the same value for each row as the `NUMBER` function.

```
SELECT my_plus_counter(t.x),
       my_plus_counter(),
       my_plus_counter(),
       NUMBER()
FROM t
```

UDF Example: my\_byte\_length Declaration

**my\_byte\_length** is a simple scalar user-defined function that returns the size of a column in bytes.

**my\_byte\_length declaration**

When **my\_byte\_length** resides within the dynamically linkable library `my_shared_lib`, the declaration for this example is:

```
CREATE FUNCTION my_byte_length(IN arg1 LONG BINARY)
// RETURNS UNSIGNED INT
// DETERMINISTIC
// IGNORE NULL VALUES
// EXTERNAL NAME 'my_byte_length@libudfex'
```

This declaration says that **my\_byte\_length** is a simple scalar UDF residing in `my_shared_lib` with a descriptor routine named `describe_my_byte_length`. Since the behavior of a UDF may require more than one actual C/C++ entry point for its implementation, this set of entry points is not directly part of the **CREATE FUNCTION** syntax. Instead, the **CREATE FUNCTION** statement **EXTERNAL NAME** clause identifies a descriptor function for this UDF. A descriptor function, when invoked, returns a descriptor structure. That descriptor structure contains the required and optional function pointers that embody the implementation of this UDF.

This declaration also says that **my\_byte\_length** accepts one `LONG BINARY` argument and returns an `UNSIGNED INT` result value.

---

**Note:** Large object data support requires a separately licensed Sybase IQ option.

---

The declaration states that this function is deterministic, which always returns the identical result value when supplied the same input values. This means the result cannot depend on any external information beyond the supplied argument values, or on any side effects from previous invocations. By default, functions are assumed to be deterministic, so the results are the same if this characteristic is omitted from the **CREATE** statement.

The last piece of this declaration is the `IGNORE NULL VALUES` characteristic. Nearly all built-in scalar functions return a `NULL` result value if any of the input arguments are `NULL`. The `IGNORE NULL VALUES` states that the **my\_byte\_length** function follows that convention, and therefore this UDF routine is not actually invoked when either of its input values is `NULL`. Since `RESPECT NULL VALUES` is the default for functions, this characteristic must be specified in the declaration for this UDF to get the performance benefits. All functions that may return a non-`NULL` result given a `NULL` input value must use the default `RESPECT NULL VALUES` characteristic.

This example query with **my\_byte\_length** in the **SELECT** list returns a column with one row for each row in `exTable`, with an `INT` representing the size of the binary file:

```
SELECT my_byte_length(exLOBColumn)
FROM exTable
```

### Declaring a Scalar User-Defined Function in Sybase Central

Sybase IQ supports simple scalar UDFs that can be used anywhere the SQRT function can be used. These scalar UDFs can be deterministic, which means that for a given set of argument values, the function always returns the same result value. Sybase IQ also supports nondeterministic scalar functions, which means that the same arguments can return different results.

1. In Sybase Central, connect to the database as a user with DBA or Resource authority.
2. In the left pane, right-click **Procedures & Functions** and select **New > Function**.
3. In the **Welcome dialog**, type a name for the function and select the user to be the owner of the function.
4. To create a user-defined function, select **External C/C++**. Click **Next**.
5. In the **External Function Attributes dialog**, select **Scalar**.
6. Type the name of the dynamically linkable library file, omitting the .so or .dll extension.
7. Type a name for the descriptor function. Click **Next**.
8. Select the type of value to be returned in the function, and specify the size, units, and scale for the value. Click **Next**.
9. Select whether or not the function is deterministic.
10. Specify if the function respects or ignores NULL values.
11. Select whether the privileges used for running the function are from the defining user (definer) or the calling user (invoker).
12. Add a comment describing the purpose of the new function. Click **Finish**.
13. In the right pane, click the **SQL tab** to complete the procedure code.

### Defining a Scalar UDF

The C/C++ code for defining a scalar user-defined function includes four mandatory pieces.

- **extfnapi3.h** – inclusion of the UDF interface definition header file.
- **\_evaluate\_extfn** – An evaluation function. All evaluation functions take two arguments:
  - an instance of the scalar UDF context structure that is unique to each usage of a UDF that contains a set of callback function pointers, and a pointer where a UDF can store UDF-specific data.
  - a pointer to a data structure that allows access to the argument values and to the result value through the supplied callbacks.
- **a\_v3\_extfn\_scalar** – an instance of the scalar UDF descriptor structure that contains a pointer to the evaluation function.
- **Descriptor function** – returns a pointer to the scalar UDF descriptor structure.

These parts are optional:

- **\_start\_extfn** – an initialization function generally invoked once per SQL usage. If supplied, you must also place a pointer to this function into the scalar UDF descriptor

structure. All initialization functions take one argument, a pointer to the scalar UDF context structure that is unique to each usage of a UDF. The context structure passed is the same one that is passed to the evaluation routine.

- **\_finish\_extfn** – a shutdown function generally invoked once per SQL usage. If supplied, a pointer to this function must also be placed into the scalar UDF descriptor structure. All shutdown functions take one argument, a pointer to the scalar UDF context structure that is unique to each usage of a UDF. The context structure passed is the same one that is passed to the evaluation routine.

### See also

- *Declaring a Scalar UDF* on page 33

### Scalar UDF Descriptor Structure

The scalar UDF descriptor structure, `a_v3_extfn_scalar`, is defined as:

```
typedef struct a_v3_extfn_scalar {
    // Metadata descriptor for a scalar UDF
    // supplied by the UDF library to the server
    // An optional pointer to an initialize function
    void (*_start_extfn)(a_v3_extfn_scalar_context * cntxt);
    //
    // An optional pointer to a shutdown function
    void (*_finish_extfn)(a_v3_extfn_scalar_context * cntxt);
    //
    // A required pointer to a function that will be
    // called for each invocation of the UDF on a
    // new set of argument values
    void (*_evaluate_extfn)(a_v3_extfn_scalar_context * cntxt, void
*args_handle);
    // RESERVED FIELDS MUST BE INITIALIZED TO NULL
    void *reserved1_must_be_null;
    void *reserved2_must_be_null;
    void *reserved3_must_be_null;
    void *reserved4_must_be_null;
    void *reserved5_must_be_null;
    ...
} a_v3_extfn_scalar;
```

There should always be a single instance of `a_v3_extfn_scalar` for each defined scalar UDF. If the optional initialization function is not supplied, the corresponding value in the descriptor structure should be the null pointer. Similarly, if the shutdown function is not supplied, the corresponding value in the descriptor structure should be the null pointer.

The initialization function is called at least once before any calls to the evaluation routine, and the shutdown function is called at least once after the last evaluation call. The initialization and shutdown functions are normally called only once per usage.

### Scalar UDF Context Structure

The scalar UDF context structure, `a_v3_extfn_scalar_context` that is passed to each of the functions specified within the scalar UDF descriptor structure, is defined as:

```
typedef struct a_v3_extfn_scalar_context {
//----- Callbacks available via the context -----
//
short (SQL_CALLBACK *get_value)(
    void *arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value
);
short (SQL_CALLBACK *get_piece)(
    void *arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value,
    a_sql_uint32 offset
);
short (SQL_CALLBACK *get_value_is_constant)(
    void *arg_handle,
    a_sql_uint32 arg_num,
    a_sql_uint32 *value_is_constant
);
short (SQL_CALLBACK *set_value)(
    void *arg_handle,
    an_extfn_value *value,
    short append
);
a_sql_uint32 (SQL_CALLBACK *get_is_cancelled)(
    a_v3_extfn_scalar_context *cntxt
);
short (SQL_CALLBACK *set_error)(
    a_v3_extfn_scalar_context *cntxt,
    a_sql_uint32 error_number,
    const char *error_desc_string
);
void (SQL_CALLBACK *log_message)(
    const char *msg,
    short msg_length
);
short (SQL_CALLBACK *convert_value)(
    an_extfn_value *input,
    an_extfn_value *output
//----- Data available from the context -----
void * _user_data; // read-write field
//----- For Server Internal Use Only -----
void * _for_server_internal_use;
} a_v3_extfn_scalar_context;
```

---

**Note:** The `get_piece` callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the Blob (`a_v4_extfn_blob`) and Blob Input Stream (`a_v4_extfn_blob_istream`) structures instead.

---



The `_user_data` field within the scalar UDF context structure can be populated with data the UDF requires. Usually, it is filled in with a heap allocated structure by the `_start_extfn` function, and deallocated by the `_finish_extfn` function.

The rest of the scalar UDF context structure is filled with the set of callback functions, supplied by the engine, for use within each of the user's UDF functions. Most of these callback functions return a success status through a short result value; a true return indicates success. Well-written UDF implementations should never cause a failure status, but during development (and possibly in all debug builds of a given UDF library), Sybase recommends that you check that the return status values from the callbacks. Failures can come from coding errors within the UDF implementation, such as asking for more arguments than the UDF is defined to take.

The common set of arguments used by most of the callbacks includes:

- **arg\_handle** – A pointer received by all forms of the evaluation methods, through which the values for input arguments passed to the UDF are available, and through which the UDF result value can be set.
- **arg\_num** – An integer indicating which input argument is being accessed. Input arguments are numbered left to right in ascending order starting at one.
- **cntxt** – A pointer to the context structure that the server passes to all UDF entry points.
- **value** – A pointer to an instance of the `an_extfn_value` structure that is used to either get an input argument value from the server or to set the result value of the function. The `an_extfn_value` structure has this form:

```
typedef struct an_extfn_value {
    void * data;
    a_SQL_uint32 piece_len;
    union {
        a_SQL_uint32 total_len;
        a_SQL_uint32 remain_len;
    } len;
    a_SQL_data_type type;
} an_extfn_value;
```

**Table 1. Scalar External Function Context: `a_v3_extfn_scalar_context`**

Method of <code>a_v3_extfn_scalar_context</code> structure	Description
<code>void set_cannot_be_distributed(a_v3_extfn_scalar_context * cntxt)</code>	Distribution can be disabled at the UDF level, even if distribution criteria are met at the library level. By default, the UDF is assumed to be distributable if the library is distributable. It is the responsibility of the UDF to push the decision to disable distribution to the server.

**See also**

- *Blob* (`a_v4_extfn_blob`) on page 195

- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199

### Example: my\_plus Definition

The definition for the my\_plus scalar UDF example.

#### **my\_plus definition**

Because this UDF needs no initialization or shutdown function, those values within the descriptor structure are set to 0. The descriptor function name matches the EXTERNAL NAME used in the declaration. The evaluate method does not check the data type for arguments, because they are declared as INT.

```
#include "extfnapi3.h"
#include <stdlib.h>

// A simple deterministic scalar UDF that just adds
// two integer arguments and then returns the result.
//
// Corresponding SQL declaration:
//
// CREATE FUNCTION my_plus(IN arg1 INT, IN arg2 INT)
// RETURNS INT
// DETERMINISTIC
// IGNORE NULL VALUES
// EXTERNAL NAME
'my_plus@libudfex'
//
#ifdef __cplusplus
extern "C" {
#endif

static void my_plus_evaluate(a_v3_extfn_scalar_context *cntxt,
                             void *arg_handle)
{
    an_extfn_value arg;
    an_extfn_value outval;
    a_sql_int32 arg1, arg2, result;

    // Get first argument
    (void) cntxt->get_value( arg_handle, 1, &arg );
    if (arg.data == NULL)
    {
        return;
    }
    arg1 = *((a_sql_int32 *)arg.data);

    // Get second argument
    (void) cntxt->get_value( arg_handle, 2, &arg );
    if (arg.data == NULL)
    {
        return;
    }
    arg2 = *((a_sql_int32 *)arg.data);
```

```

// Set the result value
outval.type = DT_INT;
outval.piece_len = sizeof(a_sql_int32);
result = arg1 + arg2;
outval.data = &result;
cntxt->set_value( arg_handle, &outval, 0 );
}

static a_v3_extfn_scalar my_plus_descriptor = {
    0,
    0,
    &my_plus_evaluate,
    0,          // Reserved - initialize to NULL
    0,          // Reserved - initialize to NULL
    0,          // Reserved - initialize to NULL
    0,          // Reserved - initialize to NULL
    0,          // Reserved - initialize to NULL
    NULL       // _for_server_internal_use
};

a_v3_extfn_scalar *my_plus()
{
    return &my_plus_descriptor;
}

#ifdef __cplusplus
}
#endif

```

### Example: my\_plus\_counter Definition

This scalar UDF example checks the argument value pointer data to see if the input argument value is NULL. It also has an initialization function and a shutdown function, each of which can tolerate multiple calls.

### **my\_plus\_counter definition**

```

#include "extfnapi3.h"
#include <stdlib.h>

// A simple non-deterministic scalar UDF that adds
// an internal integer usage counter to its integer
// argument and then returns the resulting integer.
//
// Here, the start function creates a little structure for
// the counter, and then the finish function deallocates it.
//
// Corresponding SQL declaration:
//
//      CREATE FUNCTION plus_counter(IN arg1 INT)
//          RETURNS INT
//          NOT DETERMINISTIC
//          RESPECT NULL VALUES

```

## Scalar and Aggregate UDFs

```
//                                     EXTERNAL NAME 'my_plus_counter@libudfex'

typedef struct my_counter {
    a_sql_int32 _counter;
} my_counter;

#if defined __cplusplus
extern "C" {
#endif

static void my_plus_counter_start(a_v3_extfn_scalar_context *cntxt)
{
    my_counter *cptr = (my_counter *)cntxt->_user_data;
    // If we have not already allocated the
    // counter structure, then do so now
    if (!cptr) {
        cptr = (my_counter *)malloc(sizeof(my_counter));
        cntxt->_user_data = cptr;
    }
    cptr->_counter = 0;
}

static void my_plus_counter_finish(a_v3_extfn_scalar_context *cntxt)
{
    // If we still have an allocated the
    // counter structure, then free it now
    if (cntxt->_user_data) {
        free(cntxt->_user_data);
        cntxt->_user_data = 0;
    }
}

static void my_plus_counter_evaluate(a_v3_extfn_scalar_context
*cntxt,
                                     void *arg_handle)
{
    an_extfn_value arg;
    an_extfn_value outval;
    a_sql_int32 arg1, result;

    // Increment the usage counter
    my_counter *cptr = (my_counter *)cntxt->_user_data;
    cptr->_counter += 1;

    // Get the one argument
    (void) cntxt->get_value( arg_handle, 1, &arg );
    if (!arg.data) {
        // argument value was NULL;
        arg1 = 0;
    } else {
        arg1 = *((a_sql_int32 *)arg.data);
    }
}
```

```

    outval.type = DT_INT;
    outval.piece_len = sizeof(a_sql_int32);
    result = arg1 + cptr->_counter;
    outval.data = &result;
    cntxt->set_value( arg_handle, &outval, 0 );
}

static a_v3_extfn_scalar my_plus_counter_descriptor =
    {
        &my_plus_counter_start,
        &my_plus_counter_finish,
        &my_plus_counter_evaluate,
        NULL, // Reserved - initialize to NULL
        NULL, // Reserved - initialize to NULL
        NULL, // Reserved - initialize to NULL
        NULL, // Reserved - initialize to NULL
        NULL, // Reserved - initialize to NULL
        NULL, // _for_server_internal_use
    };

a_v3_extfn_scalar *my_plus_counter()
{
    return &my_plus_counter_descriptor;
}

#ifdef __cplusplus
}
#endif

```

### Example: my\_byte\_length Definition

The **my\_byte\_length** scalar UDF example computes the size of a column by streaming the data in piece by piece, then returns the size of the column in bytes.

### **my\_byte\_length definition**

---

**Note:** Large object data support requires a separately licensed Sybase IQ option.

---

```

#include "extfnapi4.h"
#include <stdlib.h>
#include <string.h>
#include <math.h>
#include <assert.h>

// A simple function that returns the size of a cell value in bytes
//
// CREATE FUNCTION my_byte_length(IN arg1 LONG BINARY)
// RETURNS UNSIGNED INT
// DETERMINISTIC
// IGNORE NULL VALUES
// EXTERNAL NAME 'my_byte_length@libudfex'

#ifdef __cplusplus
extern "C" {

```

## Scalar and Aggregate UDFs

```
#endif

static void my_byte_length_evaluate(a_v3_extfn_scalar_context
*cntxt,
                                void *arg_handle)
{
    if (cntxt == NULL || arg_handle == NULL)
    {
        return;
    }

    an_extfn_value arg;
    an_extfn_value outval;

    a_sql_uint64 total_len;

    // Get first argument
    a_sql_uint32 fetchedLength = 0;
    (void) cntxt->get_value( arg_handle, 1, &arg );
    if (arg.data == NULL)
    {
        return;
    }

    fetchedLength += arg.piece_len;

    // saving total length as it loses scope inside get_piece
    total_len = arg.len.total_len;

    while (fetchedLength < total_len)
    {
        (void) cntxt->get_piece( arg_handle, 1, &arg, fetchedLength );
        fetchedLength += arg.piece_len;
    }

    //if this fails, the function did not get the full data from the
cell
    assert(fetchedLength == total_len);

    outval.type = DT_UNSYNTH;
    outval.piece_len = 4;
    outval.data = &fetchedLength;
    cntxt->set_value(arg_handle, &outval, 0);
}

static a_v3_extfn_scalar my_byte_length_descriptor = {
    0,
    0,
    &my_byte_length_evaluate,
    0, // Reserved - initialize to NULL
    0, // Reserved - initialize to NULL
    0, // Reserved - initialize to NULL
    0, // Reserved - initialize to NULL
    0, // Reserved - initialize to NULL
    0, // Reserved - initialize to NULL
    NULL // _for_server_internal_use
}
```

```
};

a_v3_extfn_scalar *my_byte_length()
{
    return &my_byte_length_descriptor;
}

#ifdef __cplusplus
}
#endif
```

**See also**

- *Example: my\_byte\_length Definition* on page 45

**Declaring and Defining Aggregate UDFs**

Sybase IQ supports aggregate UDFs. The SUM function is an example of a built-in aggregate function. A simple aggregate function produces a single result value from a set of argument values. You can write aggregate UDFs that can be used anywhere the SUM aggregate can be used.

---

**Note:** The aggregate UDF examples referenced here are installed with the server, and can be found as .cxx files in \$IQDIR15/samples/udf. You can also find them in the \$IQDIR15/lib64/libudfex dynamically linkable library.

---

An aggregate function can produce either a single result, or a set of results. The number of data points in the output result set may not necessarily match the number of data points in the input set. Multiple-output aggregate UDFs must use a temporary output file to hold the results.

**Declaring an Aggregate UDF**

Aggregate UDFs are more powerful and more complex to create than scalar UDFs.

After the UDF code has been written and compiled, create a SQL function that invokes the UDF from the appropriate library file, sending the input data to the UDF.

---

**Note:** You can also *create the user-defined function declaration in Sybase Central* on page 38.

---

When implementing an aggregate UDF, you must decide:

- Whether it will operate only across an entire data set or partition as an online analytical processing (OLAP) -style aggregate, like RANK.
- Whether it will operate as either a simple aggregate or an OLAP-style aggregate, like SUM.
- Whether it will operate only as a simple aggregate over an entire group.

The declaration and the definition of an aggregate UDF reflects these usage decisions.

The syntax for creating user-defined aggregate functions is:

```
aggregate-udf-declaration:
    CREATE AGGREGATE FUNCTION [ owner.]function-name
```

## Scalar and Aggregate UDFs

```
( [ parameter , ... ] )
RETURNS data-type
    [ aggregate-routine-characteristics ... ]
EXTERNAL NAME library-and-entry-point-name-string

parameter:
    param-name data-type [ DEFAULT value ]

aggregate-routine-characteristics:
    DUPLICATE { SENSITIVE | INSENSITIVE }
    -- is the server allowed to eliminate DISTINCT
    | SQL SECURITY { INVOKER | DEFINER }
    | OVER restrict
    | ORDER order-restrict
        | WINDOW FRAME
        | { { ALLOWED | REQUIRED }
            [ window-frame-constraints ... ]
          | NOT ALLOWED }
    | ON EMPTY INPUT RETURNS { NULL | VALUE }
    -- Call or skip function on NULL inputs

window-frame-constraints:
    VALUES { [ NOT ] ALLOWED }
    | CURRENT ROW { REQUIRED | ALLOWED }
    | [ UNBOUNDED ] { PRECEDING | FOLLOWING } restrict

restrict:    { [ NOT ] ALLOWED } | REQUIRED

order-restrict:
{ NOT ALLOWED | SENSITIVE | INSENSITIVE | REQUIRED
```

The handling of the return data type, arguments, data types, and default values are identical to that in the scalar UDF definition.

If an aggregate UDF can be used as a simple aggregate, then it can potentially be used with the `DISTINCT` qualifier. The `DUPLICATE` clause in the aggregate UDF declaration determines:

- Whether duplicate values can be considered for elimination before the aggregate UDF is called because the results are sensitive to duplicates (such as for the built-in “`COUNT(DISTINCT T.A)`”) or,
- Whether the results are insensitive to the presence of duplicates (such as for “`MAX(DISTINCT T.A)`”).

The `DUPLICATE INSENSITIVE` option allows the optimizer to consider removing the duplicates without affecting the result, giving the optimizer the choice on how to execute the query. Write the aggregate UDF to expect duplicates. If duplicate elimination is required, the server performs it before starting the set of `_next_value_extfn` calls.

Most of the remaining clauses that are not part of the scalar UDF syntax allow you to specify the usages for this function. By default, an aggregate UDF is assumed to be usable as both a simple aggregate and as an OLAP-style aggregate with any kind of window frame.

For an aggregate UDF to be used only as a simple aggregate function, declare it using:



```
OVER NOT ALLOWED
```

Any attempt to then use this aggregate as an OLAP-style aggregate generates an error.

For aggregate UDFs that allow or require an `OVER` clause, the UDF definer can specify restrictions on the presence of the `ORDER BY` clause within the `OVER` clause by specifying “`ORDER`” followed by the restriction type. Window-ordering restriction types:

- **REQUIRED** – `ORDER BY` must be specified and cannot be eliminated.
- **SENSITIVE** – `ORDER BY` may or may not be specified, but cannot be eliminated when specified.
- **INSENSITIVE** – `ORDER BY` may or may not be specified, but the server can do ordering elimination for efficiency.
- **NOT ALLOWED** – `ORDER BY` cannot be specified.

Declare an aggregate UDF that makes sense only as an OLAP-style aggregate over an entire set or partition that has been ordered, like the built-in `RANK`, with:

```
OVER REQUIRED
ORDER REQUIRED
WINDOW FRAME NOT ALLOWED
```

Declare an aggregate UDF that makes sense only as an OLAP-style aggregate using the default window frame of `UNBOUNDED PRECEDING` to `CURRENT ROW`, with:

```
OVER REQUIRED
ORDER REQUIRED
WINDOW FRAME ALLOWED
    RANGE NOT ALLOWED
    UNBOUNDED PRECEDING REQUIRED
    CURRENT ROW REQUIRED
    FOLLOWING NOT ALLOWED
```

The defaults for the all various options and restriction sets are:

```
DUPLICATE SENSITIVE
SQL SECURITY DEFINER
OVER ALLOWED
ORDER SENSITIVE
WINDOW FRAME ALLOWED
CURRENT ROW ALLOWED
PRECEDING ALLOWED
UNBOUNDED PRECEDING ALLOWED
FOLLOWING ALLOWED
UNBOUNDED FOLLOWING ALLOWED
```

## SQL Security

Defines whether the function is executed as the `INVOKER`, (the user who is calling the function), or as the `DEFINER` (the user who owns the function). The default is `DEFINER`.

When **SQL SECURITY INVOKER** is specified, more memory is used because each user that calls the procedure requires annotation. Also, when **SQL SECURITY INVOKER** is specified,

## Scalar and Aggregate UDFs

name resolution is performed on both the user name and the INVOKER. Qualify all object names (tables, procedures, and so on) with their appropriate owner.

### External Name

A function using the **EXTERNAL NAME** clause is a wrapper around a call to a function in an external library. A function using **EXTERNAL NAME** can have no other clauses following the **RETURNS** clause. The library name may include the file extension, which is typically `.dll` on Windows and `.so` on UNIX. In the absence of the extension, the software appends the platform-specific default file extension for libraries.

The **EXTERNAL NAME** clause is not supported for temporary functions. See *SQL Anywhere Server - Programming > SQL Anywhere external call interface*.

---

**Note:** This reference points to SQL Anywhere documentation.

---

The server can be started with a library load path that includes the location of the UDF library. On UNIX variants, this can be done by modifying the `LD_LIBRARY_PATH` within the `start_iq` startup script. While `LD_LIBRARY_PATH` is universal to all UNIX variants, `SHLIB_PATH` is preferred on HP, and `LIB_PATH` is preferred on AIX.

On UNIX platforms, the external name specification can contain a fully qualified name, in which case the `LD_LIBRARY_PATH` is not used. On the Windows platform, a fully qualified name cannot be used and the library search path is defined by the `PATH` environment variable.

---

**Note:** Scalar user-defined functions and user-defined aggregate functions are not supported in updatable cursors.

---

### See also

- *Defining an Aggregate UDF* on page 54
- *Context Storage of Aggregate User-Defined Functions* on page 82

### Example: my\_sum Declaration

The “my\_sum” example is similar to the built-in SUM, except it operates only on integers.

### **my\_sum declaration**

Since `my_sum`, like `SUM`, can be used in any context, it has a relatively brief declaration:

```
CREATE AGGREGATE FUNCTION my_sum(IN arg1 INT)
  RETURNS BIGINT
  ON EMPTY INPUT RETURNS NULL
  EXTERNAL NAME 'describe_my_integer_sum@my_shared_lib'
```

The various usage restrictions all default to `ALLOWED` to specify that this function can be used anywhere in a SQL statement that any aggregate function is allowed.

Without any usage restrictions, `my_sum` is usable as a simple aggregate across an entire set of rows, as shown here:

```
SELECT MIN(t.x), COUNT (*), my_sum(t.y)
FROM t
```

Without usage restrictions, `my_sum` is also usable as a simple aggregate computed for each group as specified by a `GROUP BY` clause:

```
SELECT t.x, COUNT(*), my_sum(t.y)
FROM t
GROUP BY t.x
```

Because of the lack of usage restrictions, `my_sum` is usable as an OLAP-style aggregate with an `OVER` clause, as shown in this cumulative summation example:

```
SELECT t.x,
       my_sum(t.x)
         OVER (ORDER BY t.x ROWS BETWEEN UNBOUNDED PRECEDING AND
CURRENT ROW)
         AS cumulative_x,
       COUNT(*)
FROM t
GROUP BY t.x
ORDER BY t.x
```

#### Example: my\_bit\_xor Declaration

The “`my_bit_xor`” example is analogous to the SQL Anywhere (SA) built-in `BIT_XOR`, except it operates only on unsigned integers.

#### **my\_bit\_xor declaration**

The resulting declaration is:

```
CREATE AGGREGATE FUNCTION my_bit_xor(IN arg1 UNSIGNED INT)
  RETURNS UNSIGNED INT
  ON EMPTY INPUT RETURNS NULL
  EXTERNAL NAME 'describe_my_bit_xor@my_shared_lib'
```

Like the `my_sum` example, `my_bit_xor` has no associated usage restrictions, and is therefore usable as a simple aggregate or as an OLAP-style aggregate with any kind of a window.

#### Example: my\_bit\_or Declaration

The “`my_bit_or`” example is similar to the SA built-in `BIT_OR` except it operates only on unsigned integers and can be used only as a simple aggregate.

#### **my\_bit\_or declaration**

The resulting declaration looks like:

```
CREATE AGGREGATE FUNCTION my_bit_or(IN arg1 UNSIGNED INT)
  RETURNS UNSIGNED INT
  ON EMPTY INPUT RETURNS NULL
  OVER NOT ALLOWED
  EXTERNAL NAME 'describe_my_bit_or@ my_shared_lib'
```

## Scalar and Aggregate UDFs

Unlike the `my_bit_xor` example, the `OVER NOT ALLOWED` phrase in the declaration restricts the use of this function to a simple aggregate. Because of that usage restriction, `my_bit_or` is only usable as a simple aggregate across an entire set of rows, or as a simple aggregate computed for each group as specified by a `GROUP BY` clause shown in the following example:

```
SELECT t.x, COUNT(*), my_bit_or(t.y)
FROM t
GROUP BY t.x
```

### Example: my\_interpolate Declaration

The “`my_interpolate`” example is an OLAP-style UDAF that attempts to fill in any missing values in a sequence (where missing values are denoted by NULLs) by performing linear interpolation across any set of adjacent NULL values to the nearest non-NULL value in each direction.

### **my\_interpolate declaration**

If the input at a given row is not NULL, the result for that row is the same as the input value.

**Figure 1: my\_interpolate results**

t.tran_time	t.price	my_interpolate(t.price)
4/12/08 1:40	29.50	29.50
4/12/08 1:45	29.60	29.60
4/12/08 1:50	NULL	29.70
4/12/08 1:55	29.80	29.80
4/12/08 2:00	29.65	29.65
4/12/08 2:05	NULL	29.60
4/12/08 2:10	NULL	29.55
4/12/08 2:15	29.50	29.50

To operate at a sensible cost, `my_interpolate` must run using a fixed-width, row-based window, but the user can set the width of the window based on the maximum number of adjacent NULL values he or she expects to see. This function takes a set of double-precision floating point values and produces a resulting set of doubles.

The resulting UDAF declaration looks like this:

```
CREATE AGGREGATE FUNCTION my_interpolate (IN arg1 DOUBLE)
RETURNS DOUBLE
  OVER REQUIRED
  WINDOW FRAME REQUIRED
  RANGE NOT ALLOWED
  PRECEDING REQUIRED
  UNBOUNDED PRECEDING NOT ALLOWED
  FOLLOWING REQUIRED
```

```
UNBOUNDED FOLLOWING NOT ALLOWED
EXTERNAL NAME 'describe_my_interpolate@my_shared_lib'
```

**OVER REQUIRED** means that this function cannot be used as a simple aggregate (**ON EMPTY INPUT**, if used, is irrelevant).

**WINDOW FRAME** details specify that you must use a fixed-width, row-based window that extends both forward and backward from the current row when using this function. Because of these usage restrictions, `my_interpolate` is usable as an OLAP-style aggregate with an **OVER** clause similar to:

```
SELECT t.x,
       my_interpolate(t.x)
       OVER (ORDER BY t.x ROWS BETWEEN 5 PRECEDING AND 5 FOLLOWING)
       AS x_with_gaps_filled,
       COUNT(*)
FROM t
GROUP BY t.x
ORDER BY t.x
```

Within an **OVER** clause for `my_interpolate`, the precise number of preceding and following rows may vary, and optionally, you can use a **PARTITION BY** clause; otherwise, the rows must be similar to the example above given the usage restrictions in the declaration.

### Declaring an Aggregate UDF in Sybase Central

Sybase IQ supports aggregate UDFs. The **SUM** function is an example of a built-in aggregate function. A simple aggregate function takes a set of argument values and produces a single result value from that set of inputs. User-defined aggregate functions can be written that can be used anywhere the **SUM** aggregate can be used.

1. In Sybase Central, connect to the database as a user with **DBA** or **Resource** authority.
2. In the left pane, right-click **Procedures & Functions** and select **New > Function**.
3. In the **Welcome dialog**, type a name for the function and select which user will be the owner of the function.
4. To create a user-defined function, select **External C/C++**. Click **Next**.
5. Select **Aggregate**.
6. Type the name of the dynamically linkable library file, omitting the `.so` or `.dll` extension.
7. Type a name for the descriptor function. Click **Next**.
8. Select the type of value to be returned in the function, and specify the size, units, and scale for the value. Click **Next**.
9. Select whether the privileges used for running the function are from the defining user (definer) or the calling user (invoker).
10. Specify whether the function is allowed to be, required to be, or not allowed to be, used in an **OVER** clause. Click **Next**.

If the function is not allowed to be used in an **OVER** clause, proceed with step 14.

11. Specify if the function requires the user of an **ORDER BY** clause when it is used to define a window. Click **Next**.
12. Specify if the function is allowed to be used in a **WINDOW FRAME** clause, is required to be used in an **WINDOW FRAME** clause, or is not allowed to be used in a **WINDOW FRAME** clause. Click **Next**.  
  
If the function is not allowed to be used in a **WINDOW FRAME** clause, skip to step 14.
13. Identify the constraints on the **WINDOW FRAME** clause. Click **Next**.
14. Specify if duplicate input values need to be filtered out by the database server prior to calling the function.
15. Identify if the return value of the function is **NULL** or a fixed value when it is called with no data. Click **Next**.
16. Add a comment describing the purpose of the new function. Click **Finish**.
17. In the right pane, click the **SQL tab** to complete the procedure code.

The new function appears in **Procedures & Functions**.

### Defining an Aggregate UDF

The C/C++ code for defining an aggregate user-defined function includes eight mandatory pieces.

- **extfnapi3.h** – the UDF interface definition header file. The file is `extfnapi4.h` for the v4 API.
- **\_start\_extfn** – an initialization function invoked once per SQL usage. All initialization functions take one argument: a pointer to the aggregate UDF context structure that is unique to each usage of an aggregate UDF. The context structure passed is the same one that is passed to all the supplied functions for that usage.
- **\_finish\_extfn** – a shutdown function invoked once per SQL usage. All shutdown functions take one argument: a pointer to the aggregate UDF context structure that is unique to each usage of an aggregate UDF.
- **\_reset\_extfn** – a reset function called once at the start of each new group, new partition, and if necessary, at the start of each window motion. All reset functions take one argument: a pointer to the aggregate UDF context structure that is unique to each usage of an aggregate UDF.
- **\_next\_value\_extfn** – a function called for each new set of input arguments. `_next_value_extfn` takes two arguments:

- A pointer to the aggregate UDF context, and
- An `args_handle`.

As in scalar UDFs, the `args_handle` is used with the supplied callback function pointers to access the actual argument values.

- **\_evaluate\_extfn** – an evaluation function similar to the scalar UDF evaluation function. All evaluation functions take two arguments:
  - A pointer to the aggregate UDF context structure, and

- An `args_handle`.
- **`a_v3_extfn_aggregate`** – an instance of the aggregate UDF descriptor structure that contains the pointers to all of the supplied functions for this UDF.
- **Descriptor function** – a descriptor function that returns a pointer to that aggregate UDF descriptor structure.

In addition to the mandatory pieces, there are several optional pieces that enable more optimized access for specific usage situations:

- **`_drop_value_extfn`** – an optional function pointer that is called for each input set of argument values that has fallen out of a moving window frame. This function should not set the result of the aggregation. Use the `get_value` callback function to access the input argument values, and, if necessary, through repeated calls to the `get_piece` callback function.

Set the function pointer to the null pointer if:

- This aggregate cannot be used with a window frame,
- The aggregate is not reversible in some way, or
- The user is not interested in optimal performance.

If `_drop_value_extfn` is not supplied and the user has specified a moving window, each time the window frame moves, the reset function is called and each row within the window is included by a call to the `next_value` function, and finally the evaluate function is called. If `_drop_value_extfn` is supplied, then each time the window frame moves, this drop value function is called for each row falling out of the window frame, then the `next_value` function is called for each row that has just been added into the window frame, and finally the evaluate function is called to produce the aggregate result.

- **`_evaluate_cumulative_extfn`** – an optional function pointer that may be called for each new input set of argument values. If this function is supplied, and the usage is in a row-based window frame that spans `UNBOUNDED PRECEDING` to `CURRENT ROW`, then this function is called instead of calling the next value function immediately followed by calling the evaluate function.

`_evaluate_cumulative_extfn` must set the result of the aggregation through the `set_value` callback. Access to its set of input argument values is through the usual `get_value` callback function. This function pointer should be set to the null pointer if:

- This aggregate will never be used in this manner, or
- The user is not worried about optimal performance.
- **`_next_subaggregate_extfn`** – an optional callback function pointer that works together with an `_evaluate_superaggregate_extfn` to enable some usages of this aggregate to be optimized by running in parallel.

Some aggregates, when used as simple aggregates (in other words, not OLAP-style aggregates with an `OVER` clause) can be partitioned by first producing a set of intermediate aggregate results where each intermediate result is computed from a disjointed subset of the input rows.

Examples of such partitionable aggregates include:

## Scalar and Aggregate UDFs

- SUM, where the final SUM can be computed by performing a SUM for each disjointed subset of the input rows and then performing a SUM over the sub-SUMs; and
- COUNT(\*), where the final COUNT can be computed by performing a COUNT for each disjoint subset of the input rows and then performing a SUM over the COUNTs from each partition.

When an aggregate satisfies the above conditions, the server may choose to make the computation of that aggregate parallel. For aggregate UDFs, this parallel optimization can be applied only if both the `_next_subaggregate_extfn` function pointer and the `_evaluate_superaggregate_extfn` pointer are supplied.

The `_reset_extfn` function does not set the final result of the aggregation, and by definition, has exactly one input argument value that is the same data type as the defined return value of the aggregate UDF.

Access to the subaggregate input value is through the normal `get_value` callback function. Direct communication between subaggregates and the superaggregate is impossible; the server handles all such communication. The sub-aggregates and the super-aggregate do not share a context structure. Instead, individual sub-aggregates are treated exactly the same as nonpartitioned aggregates. The independent super-aggregate sees a calling pattern that looks like this:

```
_start_extfn
_reset_extfn
_next_subaggregate_extfn (repeated 0 to N times)
_evaluate_superaggregate_extfn
_finish_extfn
```

Or like this:

```
_start_extfn
_reset_extfn
_next_subaggregate_extfn (repeated 0 to N times)
_evaluate_superaggregate_extfn
_reset_extfn
_next_subaggregate_extfn (repeated 0 to N times)
_evaluate_superaggregate_extfn
_reset_extfn
_next_subaggregate_extfn (repeated 0 to N times)
_evaluate_superaggregate_extfn
_finish_extfn
```

If neither `_evaluate_superaggregate_extfn` or `_next_subaggregate_extfn` is supplied, then the aggregate UDF is restricted, and not allowed as a simple aggregate within a query block containing `GROUP BY CUBE` or `GROUP BY ROLLUP`.

- **`_evaluate_superaggregate_extfn`** – the optional callback function pointer that works with the `_next_subaggregate_extfn` to enable some usages as a simple aggregate to be optimized through parallelization. `_evaluate_superaggregate_extfn` is called to return the result of a partitioned aggregate. The result value is sent to the server using the normal `set_value` callback function from the `a_v3_extfn_aggregate_context` structure.



**See also**

- *Declaring an Aggregate UDF* on page 47
- *Context Storage of Aggregate User-Defined Functions* on page 82
- *Blob (a\_v4\_extfn\_blob)* on page 195
- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199

**Aggregate UDF Descriptor Structure**

The aggregate UDF descriptor structure comprises several pieces.

- **typedef struct a\_v3\_extfn\_aggregate** – the metadata descriptor for an aggregate UDF function supplied by the library.
- **\_start\_extfn** – required pointer to an initialization function for which the only argument is a pointer to a\_v3\_extfn\_aggregate\_context. Typically, used to allocate some structure and store its address in the \_user\_data field within the a\_v3\_extfn\_aggregate\_context. \_start\_extfn is only ever called once per a\_v3\_extfn\_aggregate\_context.

```
void (*_start_extfn)(a_v3_extfn_aggregate_context *);
```

- **\_finish\_extfn** – required pointer to a shutdown function for which the only argument is a pointer to a\_v3\_extfn\_aggregate\_context. Typically, used to deallocate some structure with the address stored within the \_user\_data field in the a\_v3\_extfn\_aggregate\_context. \_finish\_extfn is only ever called once per a\_v3\_extfn\_aggregate\_context.

```
void (*_finish_extfn)(a_v3_extfn_aggregate_context *);
```

- **\_reset\_extfn** – required pointer to a start-of-new-group function, for which the only argument is a pointer to a\_v3\_extfn\_aggregate\_context. Typically, used to reset some values in the structure for which the address was stashed within the \_user\_data field in the a\_v3\_extfn\_aggregate\_context. \_reset\_extfn is called repeatedly.

```
void (*_reset_extfn)(a_v3_extfn_aggregate_context *);
```

- **\_next\_value\_extfn** – required function pointer to be called for each new input set of argument values. The function does not set the result of the aggregation. Access to input argument values are through the get\_value callback function and, if necessary, through repeated calls to the get\_piece callback function, which is required only if piece\_len is less than total\_len.

```
void (*_next_value_extfn)(a_v3_extfn_aggregate_context *cntxt, void *args_handle);
```

---

**Note:** The **get\_piece** callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the Blob (a\_v4\_extfn\_blob) and Blob Input Stream (a\_v4\_extfn\_blob\_istream) structures instead.

---

- **\_evaluate\_extfn** – required function pointer to be called to return the resulting aggregate result value. \_evaluate\_extfn is sent to the server using the set\_value callback function.

```
void (*_evaluate_extfn)(a_v3_extfn_aggregate_context *cntxt, void *args_handle);
```

- **\_drop\_value\_extfn** – Optional function pointer that is called for each input set of argument values that has fallen out of a moving window frame. Do not use this function to set the result of the aggregation. Access to input argument values are through the get\_value

callback function and, if necessary, through repeated calls to the `get_piece` callback function; however, access is required only if `piece_len` is less than `total_len`. Set `_drop_value_extfn` to the null pointer if:

- The aggregate cannot be used with a window frame.
- The aggregate is not reversible in some way.
- The user is not interested in optimal performance.

---

**Note:** The `get_piece` callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the Blob (`a_v4_extfn_blob`) and Blob Input Stream (`a_v4_extfn_blob_istream`) structures instead.

---

If this function is not supplied, and the user has specified a moving window, then each time the window frame moves, the reset function is called and each row now within the window is included by a call to the `next_value` function. Finally, the evaluate function is called.

However, if this function is supplied, each time the window frame moves, this `drop_value` function is called for each row falling out of the window frame, then the `next_value` function is called for each row that has just been added into the window frame. Finally, the evaluate function is called to produce the aggregate result.

```
void (*_drop_value_extfn)(a_v3_extfn_aggregate_context *cntxt, void *args_handle);
```

- **`_evaluate_cumulative_extfn`** – optional function pointer to be called for each new input set of argument values. If this function is supplied, and the usage is in a row-based window frame that spans UNBOUNDED PRECEDING to CURRENT ROW, then this function is called instead of `next_value`, immediately followed by calling `evaluate`. `_evaluate_cumulative_extfn` must set the result of the aggregation through the `set_value` callback. Access to input argument values are through the `get_value` callback function and, if necessary, through repeated calls to the `get_piece` callback function, which is only required if `piece_len` is less than `total_len`.

```
void (*_evaluate_cumulative_extfn)(a_v3_extfn_aggregate_context *cntxt, void *args_handle);
```

---

**Note:** The `get_piece` callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the Blob (`a_v4_extfn_blob`) and Blob Input Stream (`a_v4_extfn_blob_istream`) structures instead.

---

- **`_next_subaggregate_extfn`** – optional callback function pointer that, with the `_evaluate_superaggregate_extfn` function (and in some usages also with the `_drop_subaggregate_extfn` function), enables some usages of the aggregate to be optimized through parallel and partial results aggregation.

Some aggregates, when used as simple aggregates (in other words, not OLAP-style aggregates with an OVER clause) can be partitioned by first producing a set of intermediate aggregate results where each of the intermediate results is computed from a disjoint subset of the input rows. Examples of such partitionable aggregates include:

- SUM, where the final SUM can be computed by performing a SUM for each disjoint subset of the input rows and then performing a SUM over the sub-SUMs; and

- `COUNT(*)`, where the final `COUNT` can be computed by performing a `COUNT` for each disjoint subset of the input rows and then performing a `SUM` over the `COUNT`s from each partition.

When an aggregate satisfies the above conditions, the server may choose to make the computation of that aggregate parallel. For aggregate UDFs, this optimization can be applied only if both the `_next_subaggregate_extfn` callback and the `_evaluate_superaggregate_extfn` callback are supplied. This usage pattern does not require `_drop_subaggregate_extfn`.

Similarly, if an aggregate can be used with a `RANGE`-based `OVER` clause, an optimization can be applied if `_next_subaggregate_extfn`, `_drop_subaggregate_extfn`, and `_evaluate_superaggregate_extfn` functions are all supplied by the Aggregate UDF implementation.

`_next_subaggregate_extfn` does not set the final result of the aggregation, and by definition, has exactly one input argument value that is the same data type as the return value of the aggregate UDF. Access to the sub-aggregate input value is through the `get_value` callback function and, if necessary, through repeated calls to the `get_piece` callback function, which is required only if `piece_len` is less than `total_len`.

---

**Note:** The `get_piece` callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the `Blob` (`a_v4_extfn_blob`) and `Blob Input Stream` (`a_v4_extfn_blob_istream`) structures instead.

---

Direct communication between sub-aggregates and the super-aggregate is impossible; the server handles all such communication. The sub-aggregates and the super-aggregate do not share the context structure. Individual sub-aggregates are treated exactly the same as nonpartitioned aggregates. The independent super-aggregate sees a calling pattern that looks like this:

```

_start_extfn
_reset_extfn
_next_subaggregate_extfn (repeated 0 to N times)
_evaluate_superaggregate_extfn
_finish_extfn

```

```

void (*_next_subaggregate_extfn)(a_v3_extfn_aggregate_context
*cntxt, void *args_handle);

```

- **`_drop_subaggregate_extfn`** – optional callback function pointer that, together with `_next_subaggregate_extfn` and `_evaluate_superaggregate_extfn`, enables some usages involving `RANGE`-based `OVER` clauses to be optimized through a partial aggregation. `_drop_subaggregate_extfn` is called whenever a set of rows sharing a common ordering key value have collectively fallen out of a moving window. This optimization is applied only if all three functions are provided by the UDF.

```

void (*_drop_subaggregate_extfn)(a_v3_extfn_aggregate_context
*cntxt, void *args_handle);

```

- **`_evaluate_superaggregate_extfn`** – optional callback function pointer that, together with `_next_subaggregate_extfn` (and in some cases also with `_drop_subaggregate_extfn`), enables some usages to be optimized by running in parallel.

## Scalar and Aggregate UDFs

`_evaluate_superaggregate_extfn` is called, as described above, when it is time to return the result of a partitioned aggregate. The result value is sent to the server using the `set_value` callback function from the `a_v3_extfn_aggregate_context` structure:

```
void (*_evaluate_superaggregate_extfn)
(a_v3_extfn_aggregate_context *cntxt, void *args_handle);
```

- **NULL fields** – initialize these fields to NULL:

```
void * reserved1_must_be_null;
void * reserved2_must_be_null;
void * reserved3_must_be_null;
void * reserved4_must_be_null;
void * reserved5_must_be_null;
```

- **Status indicator bit field** – a bit field containing indicators that allow the engine to optimize the algorithm used to process the aggregate.

```
a_sql_uint32 indicators;
```

- **\_calculation\_context\_size** – the number of bytes for the server to allocate for each UDF calculation context. The server may allocate multiple calculation contexts during query processing. The currently active group context is available in `a_v3_extfn_aggregate_context_user_calculation_context`.

```
short _calculation_context_size;
```

- **\_calculation\_context\_alignment** – specifies the alignment requirement for the user's calculation context. Valid values include 1, 2, 4, or 8.

```
short _calculation_context_alignment;
```

- **External memory requirements** – the following fields allow the optimizer to consider the cost of externally allocated memory. With these values, the optimizer can consider the degree to which multiple simultaneous calculations can be made. These counters should be estimates based on a typical row or group, and should not be maximum values. If no memory is allocated by the UDF, set these fields to zero.

- `external_bytes_per_group` – The amount of memory allocated to a group at the start of each aggregate. Typically, any memory allocated during the `reset()` call.
- `external_bytes_per_row` – The amount of memory allocated by the UDF for each row of a group. Typically, the amount of memory allocated during `next_value()`.

```
double external_bytes_per_group;
double external_bytes_per_row;
```

- **Reserved fields for future use** – initialize these fields:

```
a_sql_uint64 reserved6_must_be_null;
a_sql_uint64 reserved7_must_be_null;
a_sql_uint64 reserved8_must_be_null;
a_sql_uint64 reserved9_must_be_null;
a_sql_uint64 reserved10_must_be_null;
```

- **Closing syntax** – Complete the descriptor with this syntax:

```
//----- For Server Internal Use Only -----
void * _for_server_internal_use;
} a_extfn_aggregate;
```

**See also**

- *Blob (a\_v4\_extfn\_blob)* on page 195
- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199

**Calculation Context**

The `_user_calculation_context` field allows the server to concurrently execute calculations on multiple groups of data.

An Aggregate UDF must keep intermediate counters for calculations as it is processing rows. The simple model for managing these counters is to allocate memory at the start API function, store a pointer to it in the aggregate context's `_user_data` field, then release the memory at the aggregate's finish API. An alternative method, based on the `_user_calculation_context` field, allows the server to concurrently execute calculations on multiple groups of data.

The `_user_calculation_context` field is a server-allocated memory pointer, created by the server for each concurrent processing group. The server ensures that the `_user_calculation_context` always points to the correct calculation context for the group of rows currently being processed. Between UDF API calls, depending on the data, the server may allocate new `_user_calculation_context` values. The server may save and restore calculation context areas to disk while processing a query.

The UDF stores all intermediate calculation values in this field. This illustrates a typical usage:

```
struct my_average_context
{
    int    sum;
    int    count;
};

reset(a_v3_aggregate_context *context)
{
    mycontext = (my_average_context *) context->_user_calculation_context;
    mycontext->count = 0;
    mycontext->sum = 0;
}

next_value(a_v3_aggregate_context *context, void *args_handle)
{
    mycontext = (my_average_context *) context->_user_calculation_context;
    mycontext->count++;
    ..
}
```

In this model, the `_user_data` field can still be used, but no values relating to intermediate result calculations can be stored there. The `_user_calculation_context` is NULL at both the start and finish entry points.

## Scalar and Aggregate UDFs

To use the `_user_calculation_context` to enable concurrent processing, the UDF must specify the size and alignment requirements for its calculation context, and define a structure to hold its values and set `a_v3_extfn_aggregate` and `_calculation_context_size` to the `sizeof()` of that structure.

The UDF must also specify the data alignment requirements of `_user_calculation_context` through `_calculation_context_alignment`. If `_user_calculation_context` memory contains only a character byte array, no particular alignment is necessary, and you can specify an alignment of 1. Likewise, double floating point values might require an 8-byte alignment. Alignment requirements vary by platform and data type. Specifying a larger alignment than necessary always works; however, using the smallest alignment uses memory more efficiently.

### Aggregate UDF Context Structure

The aggregate UDF context structure, `a_v3_extfn_aggregate_context`, has exactly the same set of callback function pointers as the scalar UDF context structure.

In addition, it has a read/write `_user_data` pointer just like the scalar UDF context, and a set of read-only data fields that describe the current usage and location. Each unique instance of the UDF within a statement has one aggregate UDF context instance that is passed to each of the functions specified within the aggregate UDF descriptor structure when they are called. The aggregate context structure is defined as:

- **typedef struct a\_v3\_extfn\_aggregate\_context** – One created for each instance of an external function referenced within a query. If used within a parallelized subtree within a query, there is a separate context for parallel subtree.
- **Callbacks available via the context** – Common arguments to the callback routines include:
  - **arg\_handle** – A handle to function instance and arguments provided by the server.
  - **arg\_num** – The argument number. Return values are 0..N.
  - **data** – The pointer to argument data.

The context must call `get_value` before `get_piece`, but needs to call `get_piece` only if `piece_len` is less than `total_len`.

```
short (SQL_CALLBACK *get_value)(
    void *      arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value
);
short (SQL_CALLBACK *get_piece)(
    void *      arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value,
    a_sql_uint32 offset
);
```

- **Determining whether an argument is a constant** – The UDF can ask whether a given argument is a constant. This can be useful, for example, to allow work to be done once at the first call to the `_next_value` function rather than for every call to the `_next_value` function.

```
short (SQL_CALLBACK *get_value_is_constant)(
    void *      arg_handle,
    a_sql_uint32 arg_num,
    a_sql_uint32 * value_is_constant
);
```

- **Returning a null value** – To return a null value, set "data" to NULL in an\_extfn\_value. The total\_len field is ignored on calls to set\_value, the data supplied becomes the value of the argument if append is FALSE; otherwise, the data is appended to the current value of the argument. It is expected that set\_value is called with append=FALSE for an argument before being called with append=TRUE for the same argument. The append field is ignored for fixed-length data types (in other words, all numeric data types).

```
short (SQL_CALLBACK *set_value)(
    void *      arg_handle,
    an_extfn_value *value,
    short      append
);
```

- **Determining whether the statement was interrupted** – If a UDF entry point performs work for an extended period of time (many seconds), then it should, if possible, call the get\_is\_cancelled callback every second or two to see if the user has interrupted the current statement. If the statement has been interrupted, a nonzero value is returned, and the UDF entry point should then immediately perform. Eventually, the \_finish\_extfn function is called to do any necessary cleanup, but no other UDF entry points are subsequently called.

```
a_sql_uint32 (SQL_CALLBACK *get_is_cancelled)
(a_v3_extfn_aggregate_context * cntxt);
```

- **Sending error messages** – If a UDF entry point encounters some error that should result in an error message being sent back to the user and the current statement being shut down, the set\_error callback routine should be called. set\_error causes the current statement to roll back; the user sees Error from external UDF: <error\_desc\_string>, and the SQLCODE is the negated form of <error\_number>. After a call to set\_error, the UDF entry point immediately performs a return. Eventually, \_finish\_extfn is called to perform any necessary cleanup, but no other UDF entry points are subsequently called.

```
void (SQL_CALLBACK *set_error)(
    a_v3_extfn_aggregate_context * cntxt,
    a_sql_uint32 error_number,
    // use error_number values >17000 & <100000
    const char * error_desc_string
);
```

- **Writing messages to the message log** – Messages longer than 255 bytes may be truncated.

```
void (SQL_CALLBACK *log_message)(
    const char *msg,
    short msg_length
);
```

- **Converting one data type to another** – for input:
  - an\_extfn\_value.data – input data pointer.
  - an\_extfn\_value.total\_len – length of input data.

## Scalar and Aggregate UDFs

- **an\_extfn\_value.type** – DT\_ datatype of input.

For output:

- **an\_extfn\_value.data** – UDF-supplied output data pointer.
- **an\_extfn\_value.piece\_len** – maximum length of output data.
- **an\_extfn\_value.total\_len** – server set length of converted output.
- **an\_extfn\_value.type** – DT\_ datatype of desired output.

```
short (SQL_CALLBACK *convert_value)(  
    an_extfn_value *input,  
    an_extfn_value *output  
);
```

- **Fields reserved for future use** – These are reserved for future use:

```
void * reserved1;  
void * reserved2;  
void * reserved3;  
void * reserved4;  
void * reserved5;
```

- **Data available from the context** – This data pointer can be filled in by any usage with any context data the external routine requires. The UDF allocates and deallocates this memory. A single instance of `_user_data` is active for each statement. Do not use this memory for intermediate result values.

```
void * _user_data;
```

- **Currently active calculation context** – UDFs should use this memory location to store intermediate values that calculate the aggregate. This memory is allocated by the server based on the size requested in the `a_v3_extfn_aggregate`. Intermediate calculations must be stored in this memory, since the engine may perform simultaneous calculations over more than one group. Before each UDF entry point, the server ensures that the correct context data is active.

```
void * _user_calculation_context;
```

- **Other available aggregate information** – Available at all external function entry points, including `start_extfn`. Zero indicates an unknown or not-applicable value. Estimated average number of rows per partition or group.
  - **a\_sql\_uint64\_max\_rows\_in\_frame**; – Calculates the maximum number of rows defined in the window frame. For range-based windows, this indicates unique values. Zero indicates an unknown or not-applicable value.
  - **a\_sql\_uint64\_estimated\_rows\_per\_partition**; – Displays the estimated average number of rows per partition or group. 0 indicates an unknown or not-applicable value.
  - **a\_sql\_uint32\_is\_used\_as\_a\_superaggregate**; – Identifies whether this instance is a normal aggregate or a superaggregate. Returns a result of 0 if the instance is a normal aggregate.
- **Determining window specifications** – Window specifications if a window is present on the query:
  - **a\_sql\_uint32\_is\_window\_used**; – Determines if the statement is windowed.



- **a\_sql\_uint32\_window\_has\_unbounded\_preceding;** – A return value of 0 indicates the window does not have unbounded preceding.
- **a\_sql\_uint32\_window\_contains\_current\_row;** – A return value of 0 indicates the window does not contain the current row.
- **a\_sql\_uint32\_window\_is\_range\_based;** – If the return code is 1, the window is range-based. If the return code is 0, the window is row-based.
- **Available at reset\_extfn() calls** – Returns the actual number of rows in current partition, or 0 for nonwindowed aggregate.

```
a_sql_uint64 _num_rows_in_partition;
```

- **Available only at evaluate\_extfn() calls for windowed aggregates** – Currently evaluated row number in partition (starting with 1). This is useful during the evaluation phase of unbounded windows.

```
a_sql_uint64 _result_row_from_start_of_partition;
```

- **Closing syntax** – Complete the context with:

```
//----- For Server Internal Use Only -----
void *_for_server_internal_use;
} a_v3_extfn_aggregate_context;
```

**Aggregate external function context: a\_v3\_extfn\_aggregate\_context :**

Method of a_v3_extfn_aggregate_context structure	Description
void set_cannot_be_distributed(a_v3_extfn_aggregate_context * cntxt )	Distribution can be disabled at the UDF level, even if distribution criteria are met at the library level. By default, the UDF is assumed to be distributable if the library is distributable. It is the responsibility of the UDF to push the decision to disable distribution to the server.

**See also**

- *Blob (a\_v4\_extfn\_blob)* on page 195
- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199

Example: my\_sum Definition

The aggregate UDF **my\_sum** example operates only on integers.

**my\_sum definition**

Since **my\_sum**, like SUM, can be used in any context, all the optimized optional entry points have been supplied. In this example, the normal `_evaluate_extfn` function can also be used as the `_evaluate_superaggregate_extfn` function.

```
#include "extfnapi4.h"
#include <stdlib.h>
#include <assert.h>
```

## Scalar and Aggregate UDFs

```
// Simple aggregate UDF that adds up a set of
// integer arguments, and whenever asked returns
// the resulting big integer total. For int
// arguments, the only difference between this
// UDF and the SUM built-in aggregate is that this
// UDF will return NULL if there are no input rows.
//
// The start function creates a little structure for
// the running total, and the finish function then
// deallocates it.
//
// Since there are no aggregate usage restrictions
// for this aggregate UDF, the corresponding SQL declaration
// will look like:
//
//          CREATE AGGREGATE FUNCTION my_sum(IN arg1 INT)
//          RETURNS BIGINT
//          ON EMPTY INPUT RETURNS NULL
//          EXTERNAL NAME 'my_integer_sum@libudfex'

typedef struct my_total {
    a_sql_int64    _total;
    a_sql_uint64  _num_nonnulls_seen;
} my_total;

extern "C"
void my_integer_sum_start(a_v3_extfn_aggregate_context *cntxt)
{
}

extern "C"
void my_integer_sum_finish(a_v3_extfn_aggregate_context *cntxt)
{
}

extern "C"
void my_integer_sum_reset(a_v3_extfn_aggregate_context *cntxt)
{
    my_total *cptr = (my_total *)cntxt->_user_calculation_context;
    cptr->_total = 0;
    cptr->_num_nonnulls_seen = 0;
}

extern "C"
void my_integer_sum_next_value(a_v3_extfn_aggregate_context *cntxt,
                              void *arg_handle)
{
    an_extfn_value arg;
    a_sql_int32 arg1;

    my_total *cptr = (my_total *)cntxt->_user_calculation_context;
```

```

// Get the one argument, and if non-NULL then add it to the total
//
if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
    arg1 = *((a_sql_int32 *)arg.data);
    cptr->_total += arg1;
    cptr->_num_nonnulls_seen++;
}
}

extern "C"
void my_integer_sum_drop_value(a_v3_extfn_aggregate_context *cntxt,
                              void *arg_handle)
{
    an_extfn_value arg;
    a_sql_int32 arg1;
    my_total *cptr = (my_total *)cntxt->_user_calculation_context;

    // Get the one argument, and if non-NULL then subtract it from the
    // total
    if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
        arg1 = *((a_sql_int32 *)arg.data);
        cptr->_total -= arg1;
        cptr->_num_nonnulls_seen--;
    }
}

extern "C"
void my_integer_sum_evaluate(a_v3_extfn_aggregate_context *cntxt,
                             void *arg_handle)
{
    an_extfn_value outval;
    my_total *cptr = (my_total *)cntxt->_user_calculation_context;

    // Set the output result value.  If the inputs
    // were all NULL, then set the result as NULL.
    //
    outval.type = DT_BIGINT;
    outval.piece_len = sizeof(a_sql_int64);
    if (cptr->_num_nonnulls_seen > 0) {
        outval.data = &cptr->_total;
    } else {
        outval.data = 0;
    }
    cntxt->set_value( arg_handle, &outval, 0 );
}

extern "C"
void my_integer_sum_cum_evaluate(
    a_v3_extfn_aggregate_context *cntxt,
    void *arg_handle)
{
    an_extfn_value outval;

```

## Scalar and Aggregate UDFs

```
an_extfn_value arg;
int arg1;
my_total *cptr = (my_total *)cntxt->_user_calculation_context;

// Get the one argument, and if non-NULL then add it into the
total.
//
if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
    arg1 = *((a_sql_int32 *)arg.data);
    cptr->_total += arg1;
    cptr->_num_nonnulls_seen++;
}

// Then set the output result value. If the inputs
// were all NULL, then set the result as NULL.
//
outval.type = DT_BIGINT;
outval.piece_len = sizeof(a_sql_int64);
if (cptr->_num_nonnulls_seen > 0) {
    outval.data = &cptr->_total;
} else {
    outval.data = 0;
}
cntxt->set_value( arg_handle, &outval, 0 );
}

extern "C"
void my_integer_sum_next_subagg_value(
    a_v3_extfn_aggregate_context *cntxt,
    void *arg_handle)
{
    an_extfn_value arg;
    a_sql_int64 arg1;

    my_total *cptr = (my_total *)cntxt->_user_calculation_context;

    // Get the one argument, and if non-NULL then add it to the total
    //
    if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
        arg1 = *((a_sql_int64 *)arg.data);
        cptr->_total += arg1;
        cptr->_num_nonnulls_seen++;
    }
}

extern "C"
void my_integer_sum_drop_subagg_value(
    a_v3_extfn_aggregate_context *cntxt,
    void *arg_handle)
{
    an_extfn_value arg;
    a_sql_int64 arg1;

    my_total *cptr = (my_total *)cntxt->_user_calculation_context;
```

```

    // Get the one argument, and if non-NULL then subtract it from the
    total
    //
    if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
        arg1 = *((a_sql_int64 *)arg.data);
        cptr->_total -= arg1;
        cptr->_num_nonnulls_seen--;
    }
}

a_v3_extfn_aggregate my_integer_sum_descriptor =
{
    &my_integer_sum_start,
    &my_integer_sum_finish,
    &my_integer_sum_reset,
    &my_integer_sum_next_value,
    &my_integer_sum_evaluate,
    &my_integer_sum_drop_value,
    &my_integer_sum_cum_evaluate,
    &my_integer_sum_next_subagg_value,
    &my_integer_sum_drop_subagg_value,
    &my_integer_sum_evaluate,
    NULL, // reserved1_must_be_null
    NULL, // reserved2_must_be_null
    NULL, // reserved3_must_be_null
    NULL, // reserved4_must_be_null
    NULL, // reserved5_must_be_null
    0, // indicators
    ( short )sizeof( my_total ), // context size
    8, // context alignment
    0.0, //external_bytes_per_group
    0.0, // external bytes per row
    0, // reserved6_must_be_null
    0, // reserved7_must_be_null
    0, // reserved8_must_be_null
    0, // reserved9_must_be_null
    0, // reserved10_must_be_null
    NULL // _for_server_internal_use
};

extern "C"
a_v3_extfn_aggregate *my_integer_sum()
{
    return &my_integer_sum_descriptor;
}

```

### Example: my\_bit\_xor Definition

The aggregate UDF **my\_bit\_xor** example is similar to the SA built-in BIT\_XOR, except **my\_bit\_xor** operates only on unsigned integers.

### **my\_bit\_xor definition**

Because the input and the output data types are identical, use the normal `_next_value_extfn` and `_evaluate_extfn` functions to accumulate subaggregate values and produce the superaggregate result.

```
#include "extfnapi4.h"
#include <stdlib.h>
#include <assert.h>

// Generic aggregate UDF that exclusive-ORs a set of
// unsigned integer arguments, and whenever asked
// returns the resulting unsigned integer result.
//
// The start function creates a little structure for
// the running result, and the finish function then
// deallocates it.
//
// Since there are no aggregate usage restrictions
// for this aggregate UDF, the corresponding SQL declaration
// will look like:
//
//          CREATE AGGREGATE FUNCTION my_bit_xor(IN arg1 UNSIGNED
INT)
//          RETURNS UNSIGNED INT
//          ON EMPTY INPUT RETURNS NULL
//          EXTERNAL NAME 'my_bit_xor@libudfex'

typedef struct my_xor_result {
    a_sql_uint64 _num_nonnulls_seen;
    a_sql_uint32 _xor_result;
} my_xor_result;

#if defined __cplusplus
extern "C" {
#endif

static void my_xor_start(a_v3_extfn_aggregate_context *cntxt)
{
}

static void my_xor_finish(a_v3_extfn_aggregate_context *cntxt)
{
}

static void my_xor_reset(a_v3_extfn_aggregate_context *cntxt)
```

```

{
    my_xor_result *cptr = (my_xor_result *)cntxt-
>_user_calculation_context;
    cptr->_xor_result = 0;
    cptr->_num_nonnulls_seen = 0;
}

static void my_xor_next_value(a_v3_extfn_aggregate_context *cntxt,
                             void *arg_handle)
{
    an_extfn_value arg;
    a_sql_uint32 arg1;

    my_xor_result *cptr = (my_xor_result *)cntxt-
>_user_calculation_context;

    // Get the one argument, and add it to the total
    if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
        arg1 = *((a_sql_uint32 *)arg.data);
        cptr->_xor_result ^= arg1;
        cptr->_num_nonnulls_seen++;
    }
}

static void my_xor_drop_value(a_v3_extfn_aggregate_context *cntxt,
                              void *arg_handle)
{
    an_extfn_value arg;
    a_sql_uint32 arg1;
    my_xor_result *cptr = (my_xor_result *)cntxt-
>_user_calculation_context;

    // Get the one argument, and remove it from the total
    if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
        arg1 = *((a_sql_uint32 *)arg.data);
        cptr->_xor_result ^= arg1;
        cptr->_num_nonnulls_seen--;
    }
}

static void my_xor_evaluate(a_v3_extfn_aggregate_context *cntxt,
                            void *arg_handle)
{
    an_extfn_value outval;
    my_xor_result *cptr = (my_xor_result *)cntxt-
>_user_calculation_context;

    outval.type = DT_UNSENT;
    outval.piece_len = sizeof(a_sql_uint32);
    if (cptr->_num_nonnulls_seen > 0) {
        outval.data = &cptr->_xor_result;
    } else {
        outval.data = 0;
    }
}

```

## Scalar and Aggregate UDFs

```
    }
    cntxt->set_value( arg_handle, &outval, 0 );
}

static void my_xor_cum_evaluate(a_v3_extfn_aggregate_context
*cntxt,
                               void *arg_handle)
{
    an_extfn_value  outval;
    an_extfn_value  arg;
    a_sql_uint32  arg1;
    my_xor_result *cptr = (my_xor_result *)cntxt->
    _user_calculation_context;

    // Get the one argument, and include it in the result,
    // unless that input value is null.
    //
    if (cntxt->get_value( arg_handle, 1, &arg) && arg.data) {
        arg1 = *((a_sql_uint32 *)arg.data);
        cptr->_xor_result ^= arg1;
        cptr->_num_nonnulls_seen++;
    }

    // Then set the output result value
    outval.type = DT_UNSENT;
    outval.piece_len = sizeof(a_sql_uint32);
    if (cptr->_num_nonnulls_seen > 0) {
        outval.data = &cptr->_xor_result;
    } else {
        outval.data = 0;
    }
    cntxt->set_value( arg_handle, &outval, 0 );
}

static a_v3_extfn_aggregate my_xor_descriptor =
{
    &my_xor_start,
    &my_xor_finish,
    &my_xor_reset,
    &my_xor_next_value,
    &my_xor_evaluate,
    &my_xor_drop_value,
    &my_xor_cum_evaluate,
    &my_xor_next_value,
    &my_xor_drop_value,
    &my_xor_evaluate,
    NULL, // reserved1_must_be_null
    NULL, // reserved2_must_be_null
    NULL, // reserved3_must_be_null
    NULL, // reserved4_must_be_null
    NULL, // reserved5_must_be_null
    0,    // indicators
    ( short )sizeof( my_xor_result ), // context size
    8,    // context alignment
}
```



```

        0.0, // external_bytes_per_group
        0.0, // external bytes per row
        0,   // reserved6_must_be_null
        0,   // reserved7_must_be_null
        0,   // reserved8_must_be_null
        0,   // reserved9_must_be_null
        0,   // reserved10_must_be_null
        NULL // _for_server_internal_use
    };

a_v3_extfn_aggregate *my_bit_xor()
{
    return &my_xor_descriptor;
}

#ifdef __cplusplus
}
#endif

```

### Example: my\_bit\_or Definition

The aggregate UDF **my\_bit\_or** example is similar to the SA built-in BIT\_OR, except **my\_bit\_or** operates only on unsigned integers, and can be used only as a simple aggregate.

### **my\_bit\_or** definition

The **my\_bit\_or** definition is somewhat simpler than the **my\_bit\_xor** example.

```

#include "extfnapi4.h"
#include <stdlib.h>
#include <assert.h>

// A simple (non-OLAP) aggregate UDF that ORs a set
// of unsigned integer arguments, and whenever asked
// returns the resulting unsigned integer result.
//
// The start function creates a little structure for
// the running result, and the finish function then
// deallocates it.
//
// The aggregate usage restrictions for this aggregate UDF
// only allow its use as a simple aggregate, so the
// corresponding SQL declaration will look like:
//
//      CREATE AGGREGATE FUNCTION my_bit_or(IN arg1 UNSIGNED INT)
//          RETURNS UNSIGNED INT
//          ON EMPTY INPUT RETURNS NULL
//          OVER NOT ALLOWED
//          EXTERNAL NAME 'my_bit_or@libudfex'

typedef struct my_or_result {
    a_sql_uint32 _or_result;
    a_sql_uint32 _non_null_seen;
} my_or_result;

```

```

#if defined __cplusplus
extern "C" {
#endif

static void my_or_start(a_v3_extfn_aggregate_context *cntxt)
{
}

static void my_or_finish(a_v3_extfn_aggregate_context *cntxt)
{
}

static void my_or_reset(a_v3_extfn_aggregate_context *cntxt)
{
    my_or_result *cptr = (my_or_result *)cntxt->_user_calculation_context;
    cptr->_or_result = 0;
    cptr->_non_null_seen = 0;
}

static void my_or_next_value(a_v3_extfn_aggregate_context *cntxt,
                             void *arg_handle)
{
    an_extfn_value arg;
    a_sql_uint32 arg1;

    my_or_result *cptr = (my_or_result *)cntxt->_user_calculation_context;

    // Get the one argument, and add it to the total
    if (cntxt->get_value( arg_handle, 1, &arg ) && arg.data)
    {
        arg1 = *((a_sql_uint32 *)arg.data);
        cptr->_or_result |= arg1;
        cptr->_non_null_seen = 1;
    }
}

static void my_or_evaluate(a_v3_extfn_aggregate_context *cntxt,
                           void *arg_handle)
{
    an_extfn_value outval;
    my_or_result *cptr = (my_or_result *)cntxt->_user_calculation_context;

    outval.type = DT_UNSINT;
    outval.piece_len = sizeof(a_sql_uint32);
    if (cptr->_non_null_seen)
    {
        outval.data = &cptr->_or_result;
    }
}

```

```

else
{
    // Return null if no values seen
    outval.data = 0;
}
cntxt->set_value( arg_handle, &outval, 0 );
}

static a_v3_extfn_aggregate my_or_descriptor =
{
    &my_or_start,
    &my_or_finish,
    &my_or_reset,
    &my_or_next_value,
    &my_or_evaluate,
    NULL, // drop_val_extfn
    NULL, // cume_eval,
    NULL, // next_subaggregate_extfn
    NULL, // drop_subaggregate_extfn
    NULL, // evaluate_superaggregate_extfn
    NULL, // reserved1_must_be_null
    NULL, // reserved2_must_be_null
    NULL, // reserved3_must_be_null
    NULL, // reserved4_must_be_null
    NULL, // reserved5_must_be_null
    0, // indicators
    ( short )sizeof( my_or_result ), // context size
    8, // context alignment
    0.0, //external_bytes_per_group
    0.0, // external bytes per row
    0, // reserved6_must_be_null
    0, // reserved7_must_be_null
    0, // reserved8_must_be_null
    0, // reserved9_must_be_null
    0, // reserved10_must_be_null
    NULL // _for_server_internal_use
};

extern "C"
a_v3_extfn_aggregate *my_bit_or()
{
    return &my_or_descriptor;
}

#ifdef __cplusplus
}
#endif

```

Example: my\_interpolate definition

The aggregate UDF **my\_interpolate** example is an OLAP-style aggregate UDF that attempts to fill in NULL values within a sequence by performing linear interpolation across any set of adjacent NULL values to the nearest non-NULL value in each direction.

**my\_interpolate definition**

To operate at a sensible cost, **my\_interpolate** must run using a fixed-width, row-based window, but the user can set the width of the window based on the maximum number of adjacent NULL values expected. If the input at a given row is not NULL, the result for that row is the same as the input value. This function takes a set of double-precision floating-point values and produces a resulting set of doubles.

```
#include "extfnapi4.h"
#include <stdlib.h>
#include <assert.h>

// MY_INTERPOLATE
//
// OLAP-style aggregate UDF that accepts a double precision
// floating point argument. If the current argument value is
// not NULL, then the result value is the same as the
// argument value. On the other hand, if the current row's
// argument value is NULL, then the result, where possible,
// will be the arithmetic interpolation across the nearest
// preceding and nearest following values that are not NULL.
// In all cases the result is also a double precision value.
//
// The start function creates a structure for maintaining the
// argument values within the window including their NULLness.
// The finish function then deallocates this structure.
//
// Since there are some strict aggregate usage restrictions
// for this aggregate (must be used with a row-based window
// frame that includes the current row), the corresponding
// SQL declaration will look like:
//
// CREATE AGGREGATE FUNCTION my_interpolate(IN arg1 DOUBLE)
// RETURNS DOUBLE
// OVER REQUIRED
// WINDOW FRAME REQUIRED
// RANGE NOT ALLOWED
// PRECEDING REQUIRED
// UNBOUNDED PRECEDING NOT ALLOWED
// FOLLOWING REQUIRED
// UNBOUNDED FOLLOWING NOT ALLOWED
// EXTERNAL NAME 'my_interpolate@libudfex'

typedef struct my_window {
    int    _allocated_elem;
    int    _first_used;
    int    _next_insert_loc;
```

```

int      *_is_null;
double   *_dbl_val;
int      _num_rows_in_frame;
} my_window;

#if defined __cplusplus
extern "C" {
#endif

static void my_interpolate_reset(a_v3_extfn_aggregate_context
*cntxt)
{
    assert(cntxt->_user_data);
    my_window *cptr = (my_window *)cntxt->_user_data;

    cptr->_first_used = 0;
    cptr->_next_insert_loc = 0;
    cptr->_num_rows_in_frame = 0;
    for (int i=0; i<cptr->_allocated_elem; i++) {
        cptr->_is_null[i] = 1;
    }
}

static void my_interpolate_start(a_v3_extfn_aggregate_context
*cntxt)
{
    my_window *cptr = (my_window *)cntxt->_user_data;

    // Make sure function was defined correctly
    if (!cntxt->_is_window_used)
    {
        cntxt->set_error(cntxt, 20001, "Function requires window");
        return;
    }
    if (cntxt->_window_has_unbounded_preceding ||
        cntxt->_window_has_unbounded_following)
    {
        cntxt->set_error(cntxt, 20002, "Window cannot be unbounded");
        return;
    }
    if (cntxt->_window_is_range_based)
    {
        cntxt->set_error(cntxt, 20003, "Window must be row based");
        return;
    }

    if (!cptr) {
        //
        cptr = (my_window *)malloc(sizeof(my_window));
        if (cptr) {
            cptr->_is_null = 0;
            cptr->_dbl_val = 0;

```

## Scalar and Aggregate UDFs

```
    cptr->_num_rows_in_frame = 0;
    cptr->_allocated_elem = ( int )cntxt->_max_rows_in_frame;
    cptr->_is_null = (int *)malloc(cptr->_allocated_elem
                                   * sizeof(int));
    cptr->_dbl_val = (double *)malloc(cptr->_allocated_elem
                                       * sizeof(double));
    cntxt->_user_data = cptr;
}
}
if (!cptr || !cptr->_is_null || !cptr->_dbl_val) {
    // Terminate this query
    cntxt->set_error(cntxt, 20000, "Unable to allocate memory");
    return;
}
my_interpolate_reset(cntxt);
}

static void my_interpolate_finish(a_v3_extfn_aggregate_context
*cntxt)
{
    if (cntxt->_user_data) {
        my_window *cptr = (my_window *)cntxt->_user_data;
        if (cptr->_is_null) {
            free(cptr->_is_null);
            cptr->_is_null = 0;
        }
        if (cptr->_dbl_val) {
            free(cptr->_dbl_val);
            cptr->_dbl_val = 0;
        }
        free(cntxt->_user_data);
        cntxt->_user_data = 0;
    }
}

static void my_interpolate_next_value(a_v3_extfn_aggregate_context
*cntxt,
                                     void *arg_handle)
{
    an_extfn_value arg;
    double arg1;
    my_window *cptr = (my_window *)cntxt->_user_data;

    // Get the one argument, and stash its value
    // within the rotating window arrays
    //
    int curr_cell_num = cptr->_next_insert_loc % cptr-
>_allocated_elem;
    if (cntxt->get_value( arg_handle, 1, &arg ) && arg.data != NULL ) {
        arg1 = *((double *)arg.data);
        cptr->_dbl_val[curr_cell_num] = arg1;
        cptr->_is_null[curr_cell_num] = 0;
    } else {
        cptr->_is_null[curr_cell_num] = 1;
    }
}
```

```

    }

    // Then increment the insertion location and number of rows in
frame
    cptr->_next_insert_loc = ((cptr->_next_insert_loc + 1)
                             % cptr->_allocated_elem);
    cptr->_num_rows_in_frame++;
}

static void my_interpolate_drop_value(a_v3_extfn_aggregate_context
*cntxt,
                                     void * /*arg_handle*/)
{
    my_window *cptr = (my_window *)cntxt->_user_data;

    // Drop one value from the window by incrementing past it and
// decrement the number of rows in the frame
    cptr->_first_used = ((cptr->_first_used + 1) % cptr-
>_allocated_elem);
    cptr->_num_rows_in_frame--;
}

static void my_interpolate_evaluate(a_v3_extfn_aggregate_context
*cntxt,
                                     void *arg_handle)
{
    an_extfn_value  outval;
    my_window *cptr = (my_window *)cntxt->_user_data;
    double  result;
    int     result_is_null = 1;
    double  preceding_value;
    int     preceding_value_is_null = 1;
    double  preceding_distance = 0;
    double  following_value;
    int     following_value_is_null = 1;
    double  following_distance = 0;
    int  j;

    // Determine which cell is the current cell
    int curr_cell_num =
        ((int)(cntxt->_result_row_from_start_of_partition-1))%cptr-
>_allocated_elem;
    int tmp_cell_num;

    int result_row_offset_from_start_of_frame = cptr->_first_used <=
curr_cell_num ?
        ( curr_cell_num - cptr->_first_used ) :
        ( curr_cell_num + cptr->_allocated_elem - cptr-
>_first_used );

```

## Scalar and Aggregate UDFs

```
// Compute the result value
if (cptr->_is_null[curr_cell_num] == 0) {
    //
    // If the current rows input value is not NULL, then there is
    // no need to interpolate, just use that input value.
    //
    result = cptr->_dbl_val[curr_cell_num];
    result_is_null = 0;
    //
} else {
    //
    // If the current rows input value is NULL, then we do
    // need to interpolate to find the correct result value.
    // First, find the nearest following non-NULL argument
    // value after the current row.
    //
    int rows_following = cptr->_num_rows_in_frame -
        result_row_offset_from_start_of_frame - 1;
    for (j=0; j<rows_following; j++) {
        tmp_cell_num = ((curr_cell_num + j + 1) % cptr-
            >_allocated_elem);
        if (cptr->_is_null[tmp_cell_num] == 0) {
            following_value = cptr->_dbl_val[tmp_cell_num];
            following_value_is_null = 0;
            following_distance = j + 1;
            break;
        }
    }
    // Second, find the nearest preceding non-NULL
    // argument value before the current row.
    //
    int rows_before = result_row_offset_from_start_of_frame;
    for (j=0; j<rows_before; j++) {
        tmp_cell_num = ((curr_cell_num + cptr->_allocated_elem - j - 1)
            % cptr->_allocated_elem);
        if (cptr->_is_null[tmp_cell_num] == 0) {
            preceding_value = cptr->_dbl_val[tmp_cell_num];
            preceding_value_is_null = 0;
            preceding_distance = j + 1;
            break;
        }
    }
    // Finally, see what we can come up with for a result value
    //
    if (preceding_value_is_null && !following_value_is_null) {
        //
        // No choice but to mirror the nearest following non-NULL value
        // Example:
        //
        // Inputs:  NULL      Result of my_interpolate:  40.0
        //           NULL      Result of my_interpolate:  40.0
        //           40.0      Result of my_interpolate:  40.0
        //
        result = following_value;
        result_is_null = 0;
    }
    //

```



```

    } else if (!preceding_value_is_null && following_value_is_null) {
        //
        // No choice but to mirror the nearest preceding non-NULL value
        // Example:
        //
        //     Inputs:  10.0      Result of my_interpolate:  10.0
        //              NULL      Result of my_interpolate:  10.0
        //
        result = preceding_value;
        result_is_null = 0;
        //
    } else if (!preceding_value_is_null && !following_value_is_null)
    {
        //
        // Here we get to do real interpolation based on the
        // nearest preceding non-NULL value, the nearest following
        // non-NULL value, and the relative distances to each.
        // Examples:
        //
        //     Inputs:  10.0      Result of my_interpolate:  10.0
        //              NULL      Result of my_interpolate:  20.0
        //              NULL      Result of my_interpolate:  30.0
        //              40.0      Result of my_interpolate:  40.0
        //
        //     Inputs:  10.0      Result of my_interpolate:  10.0
        //              NULL      Result of my_interpolate:  25.0
        //              40.0      Result of my_interpolate:  40.0
        //
        result = ( preceding_value
                   + ( (following_value - preceding_value)
                       * ( preceding_distance
                           / (preceding_distance +
following_distance))));
        result_is_null = 0;
    }
}

// And last, pass the result value out
outval.type = DT_DOUBLE;
outval.piece_len = sizeof(double);
if (result_is_null) {
    outval.data = 0;
} else {
    outval.data = &result;
}
cntxt->set_value( arg_handle, &outval, 0 );
}

static a_v3_extfn_aggregate my_interpolate_descriptor =
{
    &my_interpolate_start,
    &my_interpolate_finish,
    &my_interpolate_reset,
    &my_interpolate_next_value, //( timeseries_expression )
    &my_interpolate_evaluate,

```

```

    &my_interpolate_drop_value,
    NULL, // cume_eval,
    NULL, // next_subaggregate_extfn
    NULL, // drop_subaggregate_extfn
    NULL, // evaluate_superaggregate_extfn
    NULL, // reserved1_must_be_null
    NULL, // reserved2_must_be_null
    NULL, // reserved3_must_be_null
    NULL, // reserved4_must_be_null
    NULL, // reserved5_must_be_null
    0, // indicators
    0, // context size
    0, // context alignment
    0.0, //external_bytes_per_group
    ( double )sizeof( double ), // external bytes per row
    0, // reserved6_must_be_null
    0, // reserved7_must_be_null
    0, // reserved8_must_be_null
    0, // reserved9_must_be_null
    0, // reserved10_must_be_null
    NULL // _for_server_internal_use
};

a_v3_extfn_aggregate *my_interpolate()
{ return &my_interpolate_descriptor; }

#ifdef __cplusplus
}
#endif

```

### **Context Storage of Aggregate User-Defined Functions**

The context area is used to transfer or communicate data between multiple invocations of the UDF within the same query (particularly within OLAP-style queries).

Context variables control whether the intermediate results of aggregate functions are to be managed by the UDF itself (forcing the Sybase IQ server to run the UDFs serially), or whether the memory is to be managed by the Sybase IQ server.

If the `_calculation_context_size` is set to 0, then the UDF is required to manage all interim results in memory, (forcing the Sybase IQ server to invoke the UDF sequentially over the data (instead of being able to invoke many instances of the UDF in parallel during an OLAP query)).

If the `_calculation_context_size` is set to a nonzero value, the Sybase IQ server manages a separate context area for each invocation of the UDF, allowing multiple instances of the UDF to be invoked in parallel. To make the most efficient use of memory, consider setting the `_calculation_context_alignment` a value smaller than the default (depending on the size of the context storage needed).

For details on context storage, refer to the description of `_calculation_context_size` and `_calculation_context_alignment` in the section *Aggregate UDF descriptor structure* on page 57. These variables are near the end of the descriptor structure.

For a detailed discussion about the use of context storage, see *Calculation context* on page 61.

---

**Important:** To store intermediate results in memory within an aggregate UDF, initialize the memory with the `_start_extfn` function, and clean up and de-allocate any memory with the `_finish_extfn` function.

---

### See also

- *Declaring an Aggregate UDF* on page 47
- *Defining an Aggregate UDF* on page 54

## Calling Scalar and Aggregate UDFs

---

You can use a user-defined function, subject to permissions, any place you use a built-in nonaggregate function.

This Interactive SQL statement returns a full name from two columns containing a first and last name:

```
SELECT fullname (GivenName, LastName)
FROM Employees;
```

<b>fullname (Employees.GivenName,Employees.SurName)</b>
Fran Whitney
Matthew Cobb
Philip Chin
...

The following statement returns a full name from a supplied first and last name:

```
SELECT fullname ('Jane', 'Smith');
```

<b>fullname ('Jane','Smith')</b>
Jane Smith

Any user who has been granted Execute permissions for the function can use the *fullname* function.

## Scalar and Aggregate UDF Calling Patterns

---

Calling patterns are steps the functions perform as results are gathered.

## **Scalar and Aggregate UDF Callback Functions**

The set of callback functions are supplied by the engine through the `a_v3_extfn_scalar_context` structure and used within the user's UDF functions.

- **get\_value** – The function used within an evaluation method to retrieve the value of each input argument. For narrow argument data types (smaller than 256 bytes), a call to `get_value` is sufficient to retrieve the entire argument value. For wider argument data types, if the `piece_len` field within the `an_extfn_value` structure passed to this callback comes back with a value smaller than the value in the `total_len` field, use the `get_piece` callback to retrieve the rest of the input value.
- **get\_piece** – The function used to retrieve subsequent fragments of a long argument input value.

---

**Note:** The **get\_piece** callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the `Blob (a_v4_extfn_blob)` and `Blob Input Stream (a_v4_extfn_blob_istream)` structures instead.

---

- **get\_is\_constant** – A function that determines whether the specified input argument value is a constant. This can be useful for optimizing a UDF, for example, where work can be performed once during the first call to the `_evaluate_extfn` function, rather than for every evaluation call.
- **set\_value** – The function used within an evaluation function to tell the server the result value of the UDF for this call. If the result data type is narrow, one call to `set_value` is sufficient. However, if the result data value is wide, then multiple calls to `set_value` are required to pass the entire value, and the `append` argument to the callback should be true for each fragment except the last. To return a NULL result, the UDF should set the data field within the result value's `an_extfn_value` structure to the null pointer.
- **get\_is\_cancelled** – A function to determine whether the statement has been cancelled. If a UDF entry point is performing work for an extended period of time (many seconds), then it should, if possible, call the `get_is_cancelled` callback every second or two to see if the user has interrupted the current statement. The return value is 0 if the statement has not been interrupted.

Sybase IQ can handle extremely large data sets, and some queries can run for long periods of time. Occasionally, a query takes an unusually long time to execute. The SQL client lets the user cancel a query if it is taking too long to complete. Native functions track when a user has canceled a query. UDFs must also be written in a manner that tracks whether a query has been canceled by the user. In other words, UDFs should support the ability for users to cancel long-running queries that invoke UDFs.

- **set\_error** – A function that can be used to communicate an error back to the server, and eventually to the user. Call this callback routine if a UDF entry point encounters an error that should result in an error message being sent back to the user. When called, `set_error` rolls back the current statement, and the user receives `Error from external UDF: error_desc_string`, and the `SQLCODE` is the negated form of the supplied `error_number`. To avoid collisions with existing errors, UDFs should use `error_number`

values between 17000 and 99999. The maximum length of “error\_desc\_string” is 140 characters.

- **log\_message** – The function used to send a message to the server's message log. The string must be a printable text string no longer than 255 bytes.
- **convert\_value** – The function allows data conversion between data types. The primary use is the conversion between DT\_DATE, DT\_TIME, and DT\_TIMESTAMP, and DT\_TIMESTAMP\_STRUCT. An input and output an\_extfn\_value is passed to the function.

### See also

- *Scalar UDF Calling Pattern* on page 85
- *Aggregate UDF Calling Patterns* on page 85
- *Blob (a\_v4\_extfn\_blob)* on page 195
- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199

## Scalar UDF Calling Pattern

Expected calling pattern for supplied function pointers for a scalar UDF calling pattern.

```
_start_extfn(if supplied)
_evaluate_extfn (repeated 0 to numerous times)
_finish_extfn(if supplied)
```

### See also

- *Scalar and Aggregate UDF Callback Functions* on page 84
- *Aggregate UDF Calling Patterns* on page 85

## Aggregate UDF Calling Patterns

The calling patterns for the user-supplied aggregate UDF functions are more complex and varied than the scalar calling patterns.

The examples that follow this table definition:

```
create table t (a int, b int, c int)
insert into t values (1, 1, 1)
insert into t values (2, 1, 1)
insert into t values (3, 1, 1)
insert into t values (4, 2, 1)
insert into t values (5, 2, 1)
insert into t values (6, 2, 1)
```

The following abbreviation is used:

**RR = a\_v3\_extfn\_aggregate\_context.\_result\_row\_offset\_from\_start\_of\_partition** – This value indicates the current row number inside the current partition for which a value is calculated. The value is set during windowed aggregates and is intended to be used during the evaluation step of unbounded windows; it is available at all evaluate calls.

## Scalar and Aggregate UDFs

Sybase IQ is a multi user application. Many users can simultaneously execute the same UDF. Certain OLAP queries execute UDFs multiple times within the same query, sometimes in parallel.

### See also

- *Scalar and Aggregate UDF Callback Functions* on page 84
- *Scalar UDF Calling Pattern* on page 85

### **Simple Aggregate Ungrouped**

The simple aggregate ungrouped calling pattern totals the input values of all rows and produces a result.

#### Query

```
select my_sum(a) from t
```

#### Calling pattern

```
_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args) -- input a=1
_next_value_extfn(cntxt, args) -- input a=2
_next_value_extfn(cntxt, args) -- input a=3
_next_value_extfn(cntxt, args) -- input a=4
_next_value_extfn(cntxt, args) -- input a=5
_next_value_extfn(cntxt, args) -- input a=6
_evaluate_extfn(cntxt, args) -- returns 21
_finish_extfn(cntxt)
```

#### Result

```
my_sum(a)
21
```

### **Simple Aggregate Grouped**

The simple aggregate grouped calling pattern totals the input values of all rows in the group and produces a result. `_reset_extfn` identifies the beginning of a group.

#### Query

```
select b, my_sum(a) from t group by b order by b
```

#### Calling pattern

```
_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args) -- input a=1
_next_value_extfn(cntxt, args) -- input a=2
_next_value_extfn(cntxt, args) -- input a=3
_evaluate_extfn(cntxt, args) -- returns 6
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args) -- input a=4
```

```

_next_value_extfn(cntxt, args) -- input a=5
_next_value_extfn(cntxt, args) -- input a=6
_evaluate_extfn(cntxt, args)  -- returns 15
_finish_extfn(cntxt)

```

**Result**

```

b,    my_sum(a)
1,    6
2,    15

```

**OLAP-Style Aggregate Calling Pattern with Unbounded Window**

Partitioning on “b” creates the same partitions as grouping on “b”. An unbounded window causes the “a” value to be evaluated for each row of the partition. Because this is an unbounded query, all values are fed to the UDF first, followed by an evaluation cycle. Context indicators are set to 1 for `_window_has_unbounded_preceding` and `_window_has_unbounded_following`

**Query**

```

select b, my_sum(a) over (partition by b rows between
unbounded preceding and
unbounded following)
from t

```

**Calling pattern**

```

_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)    input a=1
_next_value_extfn(cntxt, args)    input a=2
_next_value_extfn(cntxt, args)    input a=3
_evaluate_extfn(cntxt, args)      rr=1  returns 6
_evaluate_extfn(cntxt, args)      rr=2  returns 6
_evaluate_extfn(cntxt, args)      rr=3  returns 6
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)    input a=4
_next_value_extfn(cntxt, args)    input a=5
_next_value_extfn(cntxt, args)    input a=6
_evaluate_extfn(cntxt, args)      rr=1  returns 15
_evaluate_extfn(cntxt, args)      rr=2  returns 15
_evaluate_extfn(cntxt, args)      rr=3  returns 15
_finish_extfn(cntxt)

```

**Result**

```

b,    my_sum(a)
1,    6
1,    6
1,    6
2,    15
2,    15
2,    15

```

### **OLAP-Style Unoptimized Cumulative Window Aggregate**

If `_evaluate_cumulative_extfn` is not supplied, this cumulative sum is evaluated through this calling pattern, which is less efficient than `_evaluate_cumulative_extfn`.

#### *Query*

```
select b, my_sum(a) over (partition by b
rows between unbounded preceding and current row)
from t
order by b
```

#### *Calling pattern*

```
_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)    -- input a=1
_evaluate_extfn(cntxt, args)     -- returns 1
_next_value_extfn(cntxt, args)    -- input a=2
_evaluate_extfn(cntxt, args)     -- returns 3
_next_value_extfn(cntxt, args)    -- input a=3
_evaluate_extfn(cntxt, args)     -- returns 6
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)    -- input a=4
_evaluate_extfn(cntxt, args)     -- returns 4
_next_value_extfn(cntxt, args)    -- input a=5
_evaluate_extfn(cntxt, args)     -- returns 9
_next_value_extfn(cntxt, args)    -- input a=6
_evaluate_extfn(cntxt, args)     -- returns 15
_finish_extfn(cntxt)
```

#### *Result*

```
b,  my_sum(a)
1,  1
1,  3
1,  6
2,  4
2,  9
2,  15
```

### **OLAP-Style Optimized Cumulative Window Aggregate**

If `_evaluate_cumulative_extfn` is supplied, this cumulative sum is evaluated where the `next_value/evaluate` sequence is combined into a single `_evaluate_cumulative_extfn` call for each row within each partition.

#### *Query*

```
select b, my_sum(a) over (partition by b rows between unbounded
preceding and current row)
from t
order by b
```



*Calling pattern*

```

_start_extfn(cntxt)
_reset_extfn(cntxt)
_evaluate_cumulative_extfn(cntxt, args) -- input a=1 returns 1
_evaluate_cumulative_extfn(cntxt, args) -- input a=2 returns 3
_evaluate_cumulative_extfn(cntxt, args) -- input a=3 returns 6
_reset_extfn(cntxt)
_evaluate_cumulative_extfn(cntxt, args) -- input a=4 returns 4
_evaluate_cumulative_extfn(cntxt, args) -- input a=5 returns 9
_evaluate_cumulative_extfn(cntxt, args) -- input a=6 returns 15
_finish_extfn(cntxt)

```

*Result*

```

b,  my_sum(a)
1,  1
1,  3
1,  6
2,  4
2,  9
2,  15

```

**OLAP-Style Unoptimized Moving Window Aggregate**

If `_drop_value_extfn` function is not supplied, this moving window sum is evaluated through this significantly less efficient than using `_drop_value_extfn`.

*Query*

```

select b, my_sum(a) over (partition by b rows between 1 preceding and
current row)
from t

```

*Calling pattern*

```

_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=1
_evaluate_extfn(cntxt, args)            returns 1
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=1
_next_value_extfn(cntxt, args)           input a=2
_evaluate_extfn(cntxt, args)            returns 3
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=2
_next_value_extfn(cntxt, args)           input a=3
_evaluate_extfn(cntxt, args)            returns 5
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=4
_evaluate_extfn(cntxt, args)            returns 4
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=4
_next_value_extfn(cntxt, args)           input a=5

```

## Scalar and Aggregate UDFs

```
_evaluate_extfn(cntxt, args)           returns 9
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)        input a=5
_next_value_extfn(cntxt, args)        input a=6
_evaluate_extfn(cntxt, args)          returns 11
_finish_extfn(cntxt)
```

### Result

```
b, my_sum(a)
1, 1
1, 3
1, 5
2, 4
2, 9
2, 11
```

### **OLAP-Style Optimized Moving Window Aggregate**

If the `_drop_value_extfn` function is supplied, this moving window sum is evaluated using this calling pattern, which is more efficient than using `_drop_value_extfn`.

### Query

```
select b, my_sum(a) over (partition by b rows between 1 preceding and
current row)
from t
```

### Calling pattern

```
_start_extfn(cntxt)
_reset_extfn(cntxt)
_evaluate_aggregate_extfn(cntxt, args)    -- returns 1
_evaluate_aggregate_extfn(cntxt, args)    -- returns 3
_drop_value_extfn(cntxt)                  -- input a=1
_next_value_extfn(cntxt, args)            -- input a=3
_evaluate_aggregate_extfn(cntxt, args)    -- returns 5
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)            -- input a=4
_evaluate_aggregate_extfn(cntxt, args)    -- returns 4
_next_value_extfn(cntxt, args)            -- input a=5
_evaluate_aggregate_extfn(cntxt, args)    -- returns 9
_drop_value_extfn(cntxt)                  -- input a=4
_next_value_extfn(cntxt, args)            -- input a=6
_evaluate_aggregate_extfn(cntxt, args)    -- returns 11
_finish_extfn(cntxt)
```

### Result

```
b, my_sum(a)
1, 1
1, 3
1, 5
2, 4
2, 9
2, 11
```

**OLAP-Style Unoptimized Moving Window Following Aggregate**

If `_drop_value_extfn` function is not supplied, this moving window sum is evaluated through the following calling pattern. This case is similar to the previous moving window example, but the row being evaluated is not the last row given by next value function.

**Query**

```
select b, my_sum(a) over (partition by b rows between 1 preceding and
1 following)
from t
```

**Calling pattern**

```
_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)          input a=1
_next_value_extfn(cntxt, args)          input a=2
_evaluate_extfn(cntxt, args)             returns 3
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=1
_next_value_extfn(cntxt, args)           input a=2
_next_value_extfn(cntxt, args)           input a=3
_evaluate_extfn(cntxt, args)             returns 6
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=2
_next_value_extfn(cntxt, args)           input a=3
_evaluate_extfn(cntxt, args)             returns 5
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=4
_next_value_extfn(cntxt, args)           input a=5
_evaluate_extfn(cntxt, args)             returns 9
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=4
_next_value_extfn(cntxt, args)           input a=5
_next_value_extfn(cntxt, args)           input a=6
_evaluate_extfn(cntxt, args)             returns 15
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)           input a=5
_next_value_extfn(cntxt, args)           input a=6
_evaluate_extfn(cntxt, args)             returns 11
_finish_extfn(cntxt)
```

**Result**

```
b, my_sum(a)
1, 3
1, 6
1, 5
2, 9
2, 15
2, 11
```

**OLAP-Style Optimized Moving Window Following Aggregate**

If `_drop_value_extfn` function is supplied, this moving window sum is evaluated through the following calling pattern. Again, this case is similar to the previous moving window example, but the row being evaluated is not the last row given by next value function.

*Query*

```
select b, my_sum(a) over (partition by b rows between 1 preceding and
1 following)
from t
```

*Calling pattern*

```
_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)          input a=1
_next_value_extfn(cntxt, args)          input a=2
_evaluate_extfn(cntxt, args)            returns 3
_next_value_extfn(cntxt, args)          input a=3
_evaluate_extfn(cntxt, args)            returns 6
_dropvalue_extfn(cntxt)                  input a=1
_evaluate_extfn(cntxt, args)            returns 5
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)          input a=4
_next_value_extfn(cntxt, args)          input a=5
_evaluate_extfn(cntxt, args)            returns 9
_next_value_extfn(cntxt, args)          input a=6
_evaluate_extfn(cntxt, args)            returns 15
_dropvalue_extfn(cntxt)                  input a=4
_evaluate_extfn(cntxt, args)            returns 11
_finish_extfn(cntxt)
```

*Result*

```
b, my_sum(a)
1, 3
1, 6
1, 5
2, 9
2, 15
2, 11
```

**OLAP-Style Unoptimized Moving Window without Current**

Assume the UDF `my_sum` works like the built-in `SUM`. If `_drop_value_extfn` function is not supplied, this moving window count is evaluated through the following calling pattern. This case is similar to the previous moving window examples, but the current row is not part of the window frame.

*Query*

```
select b, my_sum(a) over (rows between 3 preceding and 1 preceding)
from t
```

*Calling pattern*

```

_start_extfn(cntxt)
_reset_extfn(cntxt)
_evaluate_extfn(cntxt, args)           returns NULL
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)         input a=1
_evaluate_extfn(cntxt, args)           returns 1
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)         input a=1
_next_value_extfn(cntxt, args)         input a=2
_evaluate_extfn(cntxt, args)           returns 3
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)         input a=1
_next_value_extfn(cntxt, args)         input a=2
_next_value_extfn(cntxt, args)         input a=3
_evaluate_extfn(cntxt, args)           returns 6
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)         input a=2
_next_value_extfn(cntxt, args)         input a=3
_next_value_extfn(cntxt, args)         input a=4
_evaluate_extfn(cntxt, args)           returns 9
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args)         input a=3
_next_value_extfn(cntxt, args)         input a=4
_next_value_extfn(cntxt, args)         input a=5
_evaluate_extfn(cntxt, args)           returns 12
_finish_extfn(cntxt)

```

*Result*

b	my_sum(a)
1	NULL
1	1
1	3
2	6
2	9
2	12

**OLAP-Style Optimized Moving Window without Current**

If `_drop_value_extfn` function is supplied, this moving window count is evaluated through the following calling pattern. This case is similar to the previous moving window examples, but the current row is not part of the window frame.

*Query*

```

select b, my_sum(a) over (rows between 3 preceding and 1 preceding)
from t

```

*Calling pattern*

```

_start_extfn(cntxt)
_reset_extfn(cntxt)

```

## Scalar and Aggregate UDFs

```
_evaluate_extfn(cntxt, args)           returns NULL
_next_value_extfn(cntxt, args)         input a=1
_evaluate_extfn(cntxt, args)           returns 1
_next_value_extfn(cntxt, args)         input a=2
_evaluate_extfn(cntxt, args)           returns 3
_next_value_extfn(cntxt, args)         input a=3
_evaluate_extfn(cntxt, args)           returns 6
_dropvalue_extfn(cntxt)                 input a=1
_next_value_extfn(cntxt, args)         input a=4
_evaluate_extfn(cntxt, args)           returns 9
_dropvalue_extfn(cntxt)                 input a=2
_next_value_extfn(cntxt, args)         input a=5
_evaluate_extfn(cntxt, args)           returns 12
_finish_extfn(cntxt)
```

### Result

b	my_sum(a)
1	NULL
1	1
1	3
2	6
2	9
2	12

### External Function Prototypes

Define the API by a header file named `extfnapi.v3.h` (`extfnapi.v4.h` for the v4 API) in the subdirectory of your Sybase IQ installation directory. This header file handles the platform-dependent features of external function prototypes.

To notify the database server that the library is not written using the old API, provide a function as follows:

```
uint32 extfn_use_new_api( )
```

This function returns an unsigned 32-bit integer. If the return value is nonzero, the database server assumes that you are using the new API.

If the DLL does not export this function, the database server assumes that the old API is in use. When using the new API, the returned value must be the API version number defined in `extfnapi.v4h`.

Each library should implement and export this function as:

```
unsigned int extfn_use_new_api(void)
{
    return EXTFN_V4_API;
}
```

The presence of this function, and that it returns `EXTFN_V4_API` informs the Sybase IQ engine that the library contains UDFs written to the new API documented in this book.

*Function prototypes*

The name of the function must match that referenced in the **CREATE PROCEDURE** or **CREATE FUNCTION** statement. Declare the function as:

```
void function-name ( an_extfn_api *api, void *argument-handle )
```

The function must return void, and must take as arguments a structure used to pass the arguments, and a handle to the arguments provided by the SQL procedure.

The `an_extfn_api` structure has this form:

```
typedef struct an_extfn_api {
short (SQL_CALLBACK *get_value)(
    void *      arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value
);
short (SQL_CALLBACK *get_piece)(
    void *      arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value,
    a_sql_uint32 offset
);
short (SQL_CALLBACK *set_value)(
    void *      arg_handle,
    a_sql_uint32 arg_num,
    an_extfn_value *value
    short      append
);
void (SQL_CALLBACK *set_cancel)(
    void *      arg_handle,
    void *      cancel_handle
);
} an_extfn_api;
```

---

**Note:** The **get\_piece** callback is valid in v3 and v4 scalar and aggregate UDFs. For v4 table UDFs and TPFs, use the Blob (`a_v4_extfn_blob`) and Blob Input Stream (`a_v4_extfn_blob_istream`) structures instead.

---

The `an_extfn_value` structure has this form:

```
typedef struct an_extfn_value {
void *      data;
a_sql_uint32 piece_len;
union {
a_sql_uint32 total_len;
a_sql_uint32 remain_len;
} len;
a_sql_data_type type;
} an_extfn_value;
```

### *Notes*

Calling `get_value` on an OUT parameter returns the data type of the argument, and returns data as NULL.

The `get_piece` function for any given argument can be called only immediately after the `get_value` function for the same argument.

To return NULL, set data to NULL in `an_extfn_value`.

The `append` field of `set_value` determines whether the supplied data replaces (false) or appends to (true) the existing data. You must call `set_value` with `append=FALSE` before calling it with `append=TRUE` for the same argument. The `append` field is ignored for fixed-length data types.

The header file itself contains additional notes.

### **See also**

- *Blob (a\_v4\_extfn\_blob)* on page 195
- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199



# Table UDFs and TPFs

Table UDFs are external user-defined C, C++, or Java table functions. Unlike scalar and aggregate UDFs, table UDFs produce row sets as output. SQL queries consume a table UDF's output sets as table expressions.

Scalar and aggregate UDFs can use either the v3 or v4 API, but table UDFs can use only v4.

Declare a table UDF SQL function using the **CREATE PROCEDURE** statement. Scalar and aggregate UDFs use the **CREATE FUNCTION** statement.

Table parameterized functions (TPFs) are enhanced table UDFs that accept either scalar values or row sets as input.

## See also

- *Table Parameterized Functions* on page 136
- *Declaring and Defining Scalar User-Defined Functions* on page 33
- *Declaring and Defining Aggregate UDFs* on page 47
- *Learning Roadmap: Types of External C and C++ UDFs* on page 6
- *Creating a Java Table UDF* on page 352

## User Roles

---

Two types of users work with table UDFs: UDF developers, and SQL analysts.

- **UDF developer** – develops the table UDF in C or C++.
- **SQL analyst** – develops and analyzes the SQL queries that reference the table expression in the **FROM** clause. The table expression is the set of rows produced by the table UDF.

## See also

- *Learning Roadmap for Table UDF Developers* on page 97
- *Learning Roadmap for SQL Analysts* on page 98

## Learning Roadmap for Table UDF Developers

---

Use annotated examples to learn how to develop a C or C++ table UDF. After completing the development tasks, the SQL analyst can then reference your UDF in a SQL query.

This roadmap assumes:

- You have a C or C++ development environment on your machine.

## Table UDFs and TPFs

- You are familiar with standard programming practices.

Task	See
Become familiar with table UDF and TPF restrictions.	<i>Table UDF Restrictions</i> on page 99
Create a table UDF.	<i>Developing a Table UDF</i> on page 103
(Optional) Define the library version validators for distributed query processing (DQP).	<i>Library Version (extfn_get_library_version)</i> on page 17 <i>Library Version Compatibility (extfn_check_version_compatibility)</i> on page 17
Compile and link source code.	<i>Compile and Link Source Code to Build Dynamically Linkable Libraries</i> on page 19
Declare the UDF to the server using the <b>CREATE PROCEDURE</b> statement. Write and execute these statements as commands, or issue the <b>CREATE</b> statement using Sybase Central or Sybase Control Center.	<i>Learning Roadmap for SQL Analysts</i> on page 98

## Learning Roadmap for SQL Analysts

Reference a C or C++ table UDF in your SQL query.

Task	See:
Obtain the .dll or .so file (for example, myudf.dll) from the UDF developer.  Place the .dll file in the bin64 directory; place the .so file in the lib64 or LD_LIBRARY_PATH directory.	Not applicable.
Define the <b>CREATE PROCEDURE</b> statement, referencing the .dll file and the callback function.  For example: <pre>CREATE PROCEDURE my_udf( IN num_row INT ) RESULT( id INT ) EXTERNAL NAME 'udf_rg_proc@myudf.dll'</pre>	<i>CREATE PROCEDURE Statement (Table UDF)</i> on page 169

Task	See:
Select rows from the UDF. For example: <pre>SELECT * FROM my_udf ( 5 )</pre>	<i>SELECT Statement</i> on page 185 <i>FROM Clause</i> on page 179

**See also**

- *SQL Reference for Table UDF and TPF Queries* on page 166

## Table UDF Restrictions

---

Table UDFs and TPFs have some restrictions.

- IQ\_IDA license users can run table UDFs and TPFs only on reader nodes of a multiplex.
- The **TEMPORARY PROCEDURE** clause is not allowed for any external procedures. Attempting to create a temporary external procedure results in an error at creation time.
- The **NO RESULT SET** clause is not allowed. Table UDFs and TPFs must explicitly declare the contents of their results.
- If the optional **DYNAMIC RESULT SETS integer-expression** clause is specified, the value must be set to 1. Table UDFs and TPFs do not return multiple result sets.
- A table UDF or TPF cannot be referenced in a **CALL SQL** statement or **EXEC** embedded SQL statement. A table UDF or TPF can be referenced only in a **FROM** clause of a SQL statement.
- The **LANGUAGE** clause cannot be used for table UDFs or TPFs. If the **LANGUAGE** clause is present, syntax errors are reported at execution time.
- The **parameter** clause is limited to keyword **IN**; **INOUT** and **OUT** keywords are not supported for table UDFs or TPFs.
- The **EXTERNAL NAME** clause has the same syntax as scalar and aggregate UDFs.

## Get Started

---

Familiarize yourself with sample files, concepts, and restrictions before developing table UDFs and TPFs.

## Sample Files

Sample table UDF files are installed with the server. Use the samples as models when defining your own table UDFs.

Sample files are located in:

- %ALLUSERSPROFILE%\samples\udf (Windows)

## Table UDFs and TPFs

- `$IQDIR15/samples/udf` (UNIX)

File	Description
<code>apache_log_reader.cxx</code>	Implementation of a table UDF that reads an Apache log file and presents the rows from the file in table format. This UDF illustrates a real-world example of how you can use a UDF to make computer-generated data available to a SQL query writer in real time.
<code>build.sh / build.bat</code>	Script that compiles and links the sample scalar and aggregate UDFs, table UDFs, and TPFs found in the <code>samples/udf</code> directory.
<code>my_md5.cxx</code>	A simple deterministic scalar UDF that calculates the MD5 hash value of an input file (a LOB binary argument).
<code>tpf_agg.cxx</code>	Consumes rows from an input table, performs an aggregation on the input data, and returns rows back to the server.
<code>tpf_blob.cxx</code>	Implementation of a TPF that reads LOB data from an input table and passes the data to the result set, if an even number of the specified character or digit is present. This TPF illustrates how to read LOB data and how a user can pass LOB datatypes through to the result set.
<code>tpf_dt.cxx</code>	
<code>tpf_filt.cxx</code>	Illustrates how a TPF can be used to filter rows. The example uses the row block provided by caller and passes it to the input table parameter. The input table schema must match the output result set of this function.
<code>tpf_oby.cxx</code>	Illustrates how a TPF can generate ordered output and pass it along.
<code>tpf_pby.cxx</code>	Illustrates how a TPF can generate partitioned output and pass it along.
<code>tpf_rg_1.cxx</code>	Similar to the table UDF sample <code>udf_rg_2.cxx</code> . It produces rows of data based on an input parameter.
<code>tpf_rg_2.cxx</code>	Builds upon the sample in <code>tpf_rg_1.cxx</code> , but uses <code>fetch_into</code> instead of <code>fetch_block</code> to read rows from the input table.

File	Description
<code>udf_main.cxx</code>	This file is linked into all of the examples and includes a common set of required entry points for the v4 API. This allows you to reuse the code rather than including it in each example.
<code>udf_rg_1.cxx</code>	A simple table UDF that generates rows of integer data.
<code>udf_rg_2.cxx</code>	A simple table UDF that generates rows of integer data that uses <code>describes</code> to ensure the schema defined in SQL matches the UDF's implementation. It also describes some optimizer attributes.
<code>udf_rg_3.cxx</code>	A simple table UDF that generates integer data in blocks of 100 using the <code>_fetch_block</code> fetch method.
<code>udf_utils.cxx</code>	A set of utility functions and macros that are useful to UDF/TPF developers. The examples rely on items in this file.
<code>udf_utils.h</code>	A set of utility functions and macros that are useful to UDF/TPF developers. The examples rely on items in this file.

## Understanding Producers Versus Consumers

The server and UDF form a producer and consumer relationship when exchanging rows of data.

Production and consumption refer to table row data. The producer produces table rows; the consumer consumes table rows.

The server executes scalar and aggregate UDFs once for each matching row of a query. These UDFs consume input scalar parameters and produce, and return, a single scalar parameter. This data exchange occurs during the `evaluate` method using the `get_value()` and `set_value()` APIs.

However, scalar production and consumption is an inefficient method of data exchange if your UDF must produce or consume a table. Table UDFs that produce a table, and TPFs that consume a table, use the `row_block` data structure of the v4 API. Row blocks allow for bulk row and column data exchange. The row block is populated by a producer, and read from by a consumer.

In this example, the table UDF `my_table_udf()` is a producer of data. Sybase IQ, the server, is the consumer of the data:

```
SELECT * FROM my_table_udf()
```

## Table UDFs and TPFs

In general, a table UDF is always a producer of data. The server, however, may not always be the consumer:

```
SELECT * FROM my_tpf( TABLE( SELECT * FROM my_table_udf() ) )
```

The outer TPF, `my_tpf()`, is the consumer for the table input parameter specified by **SELECT \* from my\_table\_udf()**. Sybase IQ is the consumer of the table produced by the `my_tpf()` TPF. A TPF, therefore, can be both a consumer and a producer.

The TPF does not have to consume from a table UDF. In this example, the TPF consumes the table data produced by the inner query, which is produced by the Sybase IQ server:

```
SELECT * FROM my_tpf( TABLE( SELECT * FROM my_table where my_table.c1 < 10 ) )
```

In a TPF, therefore, Sybase IQ can be both the consumer and producer of table data.

In the v4 API, a row block defines a memory area where data is produced to, and consumed from. In general, the layout of a row block conceptually matches the row and column format of the table; a row block consists of a number of rows, and each row consists of a number of columns. Either the producer or consumer must allocate the row block, and must also deallocate it when the time comes.

Rows and column have their own specific attributes that only apply to them. For example, rows have a status flag which indicates if the row is present or not. This flag lets a TPF change the row status without having to move the column data. Columns have a null mask that indicates if the data value is null or not. Row blocks also have some additional attributes: maximum number of rows, and current number of rows, for example. These row block attributes are useful when a UDF wants to create a row block to handle a large set of rows, but produce a smaller number of rows as required.

The process of consuming a row is handled via one of the two fetch APIs:

- `fetch_into`
- `fetch_block`

The `fetch_into` is called when the consumer allocates the row block and passes it to the producer. The producer is then requested to populate as many rows as possible, up to the maximum number of rows. The `fetch_block` is called when the consumer wants the producer to allocate the row block. `Fetch_block` is efficient if you are developing a TPF that filters rows of data. The server (consumer) allocates the row block and fetches from the TPF using the `fetch_into` API. The TPF can then pass the same row block to the input parameter using the `fetch_block` API.

### See also

- *Row Block Data Exchange* on page 128

## Developing a Table UDF

---

The general steps for developing a table UDF include determining input and output, declaring the v4 library, defining the `a_v4_extfn_proc` descriptor, defining a library entry point function, defining how the server gets row information, implementing the `a_v4_extfn_proc` structure functions, and implement the `a_v4_extfn_table_func` structure functions.

### 1. Determine the input and output for the table UDF.

The input is defined by the parameters the procedure accepts, and the output is defined by how the **RESULT** clause for the procedure is declared. The declaration of the table UDF in SQL is separate from the implementation of the table UDF. This means that a particular implementation of a table UDF may be bound to a specific declaration. When developing a table UDF, ensure that the implementation and declaration match.

### 2. Declare the library as a v4 Library.

For Sybase IQ to recognize the library as a v4 library, the library must include the `extfnapi4.h` header file located in the subdirectory of your Sybase IQ installation directory.

This header defines the v4 API features and functions and is a superset of the v3 API; `extfnapi4.h` includes `extfnapi3.h`.

To create table UDF or TPFs, the library must provide the `extfn_use_new_api()` entry point. For v4 libraries, `extfn_use_new_api()` must return `EXTFN_V4_API`.

### 3. Define the `a_v4_extfn_proc` descriptor.

When developing a v4 table UDF or TPF, the library must declare what functions are available for the server to call.

Create a variable of type `a_v4_extfn_proc` and set each member of this structure to the address of the function within the table UDF that implements the function. The information in this variable is made available to the server via a library entry pointer. Not all members of `a_v4_extfn_proc` are required and there are two reserved fields which you must set to `NULL`.

Use this descriptor function as a model when developing your own function:

```
static a_v4_extfn_proc udf_proc_descriptor =
{
    udf_proc_start,          // optional
    udf_proc_finish,        // optional
    udf_proc_evaluate,      // required
    udf_proc_describe,      // required
    udf_proc_enter_state,   // optional
    udf_proc_leave_state,   // optional
}
```

## Table UDFs and TPFs

```
    NULL,          // Reserved: must be NULL
    NULL          // Reserved: must be NULL
};
```

### 4. Define a library entry-point function.

The table UDF library must provide a function entry point that returns an `a_v4_extfn_proc` descriptor pointer. This is the same descriptor as described in step 3.

This callback function is the main required entry point for the library.

Use this function as a model when developing your own library entry point:

```
extern "C"
a_v4_extfn_proc * SQL_CALLBACK udf_rg_proc()
/*****/
{
    return &udf_proc_descriptor;
}
```

### 5. Define how the server gets row information from the table UDF.

When developing a v4 table UDF or TPF, the library must declare how row information is transferred to the server.

Create a variable of type `a_v4_extfn_table_func` and set each member of this structure to the address of the function within the table UDF that implements the function. The information in this variable is made available to the server at runtime.

Not all members of `a_v4_extfn_table_func` are required and there are two reserved fields which must be set to `NULL`.

Use this descriptor as a model when developing your own table UDF:

```
{
    udf_table_func_open,          // required
    udf_table_func_fetch_into,   // one of fetch_into or
fetch_block required
    udf_table_func_fetch_block,  // one of fetch_into or
fetch_block required
    udf_table_func_rewind,       // optional
    udf_table_func_close,       // required
    NULL,                        // Reserved: must be NULL
    NULL,                        // Reserved: must be NULL
};
```

At the start of execution, the server calls the `a_v4_extfn_proc` function `_evaluate_extfn` to give the table UDF an opportunity to tell the server what table functions it is implementing. To do this, the table UDF must create an instance of `a_v4_extfn_table` that is given to the server. This structure contains a pointer to the `a_v4_extfn_table_func` descriptor and the number of columns in the result set.

Use this descriptor as a model when developing your own table UDF:



```
static a_v4_extfn_table_udf_rg_table = {
    &udf_table_funcs, // Table function descriptor
    1                // number_of_columns
};
```

#### 6. Implement the `a_v4_extfn_proc` structure functions.

The table UDF must provide an implementation for each of the `a_v4_extfn_proc` functions that it declares in the `a_v4_extfn_proc` descriptor in step 3.

#### 7. Implement the `a_v4_extfn_table_func` structure functions.

The table UDF must provide an implementation for each of the `a_v4_extfn_table_func` functions that it declares in the `a_v4_extfn_table_func` descriptor in step 5.

#### See also

- *Scalar and Aggregate UDF Calling Patterns* on page 83
- *udf\_rg\_2* on page 111
- *udf\_rg\_3* on page 115
- *Implementing Sample Table UDF `udf_rg_1`* on page 106
- *Table UDF Implementation Examples* on page 105
- *External Function (`a_v4_extfn_proc`)* on page 284
- *Table Functions (`a_v4_extfn_table_func`)* on page 314
- *\_evaluate\_extfn* on page 286

## Table UDF Implementation Examples

Implementation examples start with a simple table UDF and increase in complexity.

The table UDF implementation examples are included in the samples directory. These examples start with a simple table UDF and build upon its complexity and functionality as the examples progress.

The examples are available in a precompiled dynamic library called `libv4apiex`. (The extension of this library name is platform dependent.) This library has linked in the functions defined in `udf_main.cxx`, which contains the library level functions, such as `extfn_use_new_api`. Put `libv4apiex` in a directory the server can read.

#### See also

- *Running the Sample Table UDF in `udf_rg_1.cxx`* on page 111
- *Running the Sample Table UDF in `udf_rg_2.cxx`* on page 114
- *Running the Sample Table UDF in `udf_rg_3.cxx`* on page 118

**Implementing Sample Table UDF `udf_rg_1`**

The sample table UDF called `udf_rg_1` illustrates how a v4 Table UDF can generate  $n$  rows of data. The implementation of the table UDF is in the samples directory in `udf_rg_1.cxx`.

**1. Determine the input and output for the table UDF.**

This example produces  $n$  rows of data based on the value of an input parameter. The input is a single integer parameter and the output is rows that consist of a single column of type integer.

The **CREATE PROCEDURE** statement required to define this procedure is:

```
CREATE OR REPLACE PROCEDURE udf_rg_1( IN num INT )
RESULT( c1 INT )
EXTERNAL NAME 'udf_rg_1@libv4apiex'
```

**2. Declare the library as a v4 library.**

In this example, `udf_rg_1.cxx` includes the `extfnapi4.h` header file:

```
#include "extfnapi4.h"
```

To inform the server that this library contains v4 table UDFs, this function export is defined in `udf_main.cxx`:

```
a_sql_uint32 SQL_CALLBACK extfn_use_new_api( void )
/*****/
{
    return EXTFN_V4_API;
}
```

**3. Define the `a_v4_extfn_proc` descriptor.**

This declares the necessary descriptor in `udf_rg_1.cxx`:

```
static a_v4_extfn_proc udf_rg_descriptor =
{
    NULL,          // _start_extfn
    NULL,          // _finish_extfn
    udf_rg_evaluate, // _evaluate_extfn
    udf_rg_describe, // _describe_extfn
    NULL,          // _leave_state_extfn
    NULL,          // _enter_state_extfn
    NULL,          // Reserved: must be NULL
    NULL,          // Reserved: must be NULL
};
```

**4. Define a library entry point function.**

This callback function declares the main entry point function. It simply returns a pointer to the `a_v4_proc_descriptor` variable `udf_rg_descriptor`.

```
extern "C"
a_v4_extfn_proc * SQL_CALLBACK udf_rg_1_proc()
/*****/
{
```

```

    return &udf_rg_descriptor;
}

```

5. Define how the server gets row information from the table UDF.

This declares the `a_v4_extfn_table_func` descriptor that is used to tell the server how to retrieve row data from the table UDF:

```

static a_v4_extfn_table_func udf_rg_table_funcs =
{
    udf_rg_open,           // _open_extfn
    udf_rg_fetch_into,    // _fetch_into_extfn
    NULL,                 // _fetch_block_extfn
    NULL,                 // _rewind_extfn
    udf_rg_close,        // _close_extfn
    NULL,                 // Reserved: must be NULL
    NULL                  // Reserved: must be NULL
};

```

In this example, the `_fetch_into_extfn` function transfers row data to the server. This is the easiest data transfer method to understand and implement. This document refers to data transfer methods as *rowblock data exchange*. There are two rowblock data exchange functions: `_fetch_into_extfn` and `_fetch_block_extfn`.

At runtime, when the `_evaluate_extfn` function is called, the UDF publishes the table functions descriptor by setting the result set parameter. To do this, the UDF must create an instance of `a_v4_extfn_table`:

```

static a_v4_extfn_table udf_rg_table = {
    &udf_rg_table_funcs,    // Table function descriptor
    1                      // number_of_columns
};

```

This structure contains a pointer to the `udf_rg_table_funcs` structure and the number of columns in the result set. This table UDF produces a single column in its result set.

6. Implement the `a_v4_extfn_proc` structure functions.

In this example, the required function `_describe_extfn` function does not do anything. Other examples demonstrate how a table UDF can use the `describe` function:

```

static void UDF_CALLBACK udf_rg_describe(
    a_v4_extfn_proc_context *ctx )
/
*****
/
{
    // This required function is not needed in this simple example.
}

```

The `_evaluate_extfn` method sends the server information about getting the result set from the UDF. This is done by calling the `a_v4_extfn_proc_context` method `set_value` on argument 0. Argument 0 represents the return value, which for a table

UDF is a `DT_EXTFN_TABLE`. This method constructs an `_extfn_value` structure, setting the data type to `DT_EXTFN_TABLE` and setting the value pointer of this to point to the `a_v4_extfn_table` object created in step 5. For table UDFs, the type must always be `DT_EXTFN_TABLE`.

```
static void UDF_CALLBACK udf_rg_evaluate(
    a_v4_extfn_proc_context *ctx,
    void *args_handle )
/*****/
{
    an_extfn_value    result_table = { &udf_rg_table,
                                      sizeof( udf_rg_table ),
                                      sizeof( udf_rg_table ),
                                      DT_EXTFN_TABLE };

    // Tell the server what functions table functions are being
    // implemented and how many columns are in our result set.
    ctx->set_value( args_handle, 0, &result_table );
}
```

#### 7. Implement the `a_v4_extfn_table_func` structure functions.

In this example, the table UDF needs to read in the parameter passed in that contains the number of rows to generate, and cache this information to be used later. Because the `_open_extfn` method is called for each new value that the parameter has, this is an appropriate place to get this information.

In addition to the total number of rows to generate, the table UDF must also remember the next row to generate. When the server begins fetching rows from the table UDF, it may need to repeatedly call the `_fetch_into_extfn` method. This means that the table UDF must remember the last row that was generated.

This structure is created in `udf_rg_1.cxx` to contain the state information between calls:

```
struct udf_rg_state {
    a_sql_int32    next_row;    // The next row to produce
    a_sql_int32    max_row;    // The number of rows to generate.
};
```

The open method first reads in the value of argument 1 using the `a_v4_proc_context` method `get_value`. An instance of `udf_rg_state` is allocated using the `a_v4_proc_context` function `alloc`. Table UDFs should use the memory management functions (`alloc` and `free`) on the `a_v4_proc_context` structure whenever possible to manage their memory. The state object is then saved in the `user_data` field of `a_v4_proc_context`. Memory stored in this field is available to the table UDF until execution finishes.

```
static short UDF_CALLBACK udf_rg_open(
    a_v4_extfn_table_context *tctx )
/*****/
{
    an_extfn_value    value;
```

```

udf_rg_state *    state = NULL;

// Read in the value of the input parameter and store it away in a
// state object.  Save the state object in the context.
if( !tctx->proc_context->get_value( tctx->args_handle,
    1,
    &value ) ) {

// Send an error to the client if we could not get the value.
tctx->proc_context->set_error(
    tctx->proc_context,
    17001,
    "Error: Could not get the value of parameter 1" );

return 0;
}

// Allocate memory for the state using the a_v4_extfn_proc_context
// function alloc.
state = (udf_rg_state *)
tctx->proc_context->alloc( tctx->proc_context,
    sizeof( udf_rg_state ) );

// Start generating at row zero.
state->next_row = 0;

// Save the value of parameter 1
state->max_row = *(a_sql_int32 *)value.data;

// Save the state on the context
tctx->user_data = state;

return 1;
}

```

The `_fetch_info_extfn` method returns row data to the server. This method is called repeatedly until it returns false. For this example, the table UDF retrieves the state information from the `user_data` field of the `a_v4_extfn_proc_context` object to determine the next row to generate and the total number of rows to generate. This method is free to generate up to the maximum number of rows indicated in the rowblock structure passed in.

For this example, the table UDF generates a single column of type INT. It copies the data for the `next_row` saved in the state into the data pointer of the first column. Each time through the loop, the table UDF copies a new value into the data pointer and stops when either the maximum number of rows to produce is reached or the row block is full.

```

static short UDF_CALLBACK udf_rg_fetch_into(
    a_v4_extfn_table_context *tctx,
    a_v4_extfn_row_block *rb)
/*****/
{
    udf_rg_state *state = (udf_rg_state *)tctx->user_data;

// Because we are implementing fetch_into, the server has provided
// us with a row block. We need to inform the server how many rows
// this call to _fetch_into has produced.
rb->num_rows = 0;

```

## Table UDFs and TPFs

```
// The server provided row block structure contains a max_rows
// field. This field is the maximum number of rows that this row
// block can handle. We can not exceed this number. We will also
// stop producing rows when we have produced the number of rows
// required as per the max_row in the state.
while( rb->num_rows < rb->max_rows && state->next_row < state->max_row ) {

    // Get the current row from the row block data.
    a_v4_extfn_row      &row = rb->row_data[ rb->num_rows ];

    // Get the column data for the current row.
    a_v4_extfn_column_data  &col0 = row.column_data[ 0 ];

    // Copy the integer value for the next row to generate
    // into the column data for the current row.
    memcpy( col0.data, &state->next_row, col0.max_piece_len );

    state->next_row++;
    rb->num_rows++;
}

// If we produced any rows, return true.
return( rb->num_rows > 0 );
}
```

The table UDF calls the `_close_extfn` method once per new value for the parameters, after all the rows have been fetched. In other words, for each `_open_extfn` call, there is a subsequent `_close_extfn` call. In this example, the table UDF must free the memory allocated during the `_open_extfn` call which it does by retrieving the state from the `user_data` field of a `a_v4_extfn_proc_context` object and calling the `free` method.

```
static short UDF_CALLBACK udf_rg_close(
    a_v4_extfn_table_context *tctx)
/*****/
{
    udf_rg_state * state = NULL;

    // Retrieve the state that was saved in user_data
    state = (udf_rg_state *)tctx->user_data;

    // Free the memory for the state using the
    a_v4_extfn_proc_context
    // function free.
    tctx->proc_context->free( tctx->proc_context, state );
    tctx->user_data = NULL;

    return 1;
}
```

### See also

- [udf\\_rg\\_2](#) on page 111
- [udf\\_rg\\_3](#) on page 115
- [Row Block Data Exchange](#) on page 128
- [Describe API](#) on page 204
- [\\_evaluate\\_extfn](#) on page 286

- *fetch\_into* on page 308
- *Table (a\_v4\_extfn\_table)* on page 306
- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287
- *\_open\_extfn* on page 316
- *\_close\_extfn* on page 319

### Running the Sample Table UDF in *udf\_rg\_1.cxx*

The sample *udf\_rg\_1* is included in a precompiled dynamic library called *libv4apiex* (extension is platform-dependent). Its implementation is in the *samples* directory in *udf\_rg\_1.cxx*.

1. Put the library *libv4apiex* in a directory that can be read by the server.
2. To declare the table UDF to the server, issue:

```
CREATE PROCEDURE udf_rg_1( IN num INT )
RESULT( c1 INT )
EXTERNAL NAME 'udf_rg_1@libv4apiex'
```

3. Select rows from the table UDF:

```
SELECT * FROM udf_rg_1( 5 );
```

### udf\_rg\_2

The sample table UDF *udf\_rg\_2* builds on the sample in *udf\_rg\_1.cxx* and has the same behavior. The procedure is called ***udf\_rg\_2*** and its implementation is in the *samples* directory in *udf\_rg\_2.cxx*.

The table UDF ***udf\_rg\_2*** provides an alternate implementation of the *\_describe\_extfn* method in the *a\_v4\_extfn\_proc* descriptor.

```
static void UDF_CALLBACK udf_rg_describe(
    a_v4_extfn_proc_context *ctx )
/*****/
{
    a_sql_int32      desc_rc;

    // The following describes will ensure that the schema defined
    // by the user matches the schema supported by this table udf.
    // This is achieved by telling the server what our schema is
    // using describe_xxxx_set methods.
    if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {

        a_sql_data_type      type      = DT_NOTYPE;
        a_sql_uint32         num_cols  = 0;
        a_sql_uint32         num_parms = 0;

        // Inform the server that we support a single input
        // parameter.
        num_parms = 1;
        desc_rc = ctx->describe_udf_set
            ( ctx,
```

## Table UDFs and TPFs

```
EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS,  
    &num_parms,  
    sizeof( num_parms ) );  
  
// Checks the return code and sets an error if the  
// describe was unsuccessful for any reason.  
UDF_CHECK_DESCRIBE( ctx, desc_rc );  
  
// Inform the server that the type of parameter 1 is int.  
type = DT_INT;  
desc_rc = ctx->describe_parameter_set  
    ( ctx,  
      1,  
      EXTFNAPIV4_DESCRIBE_PARM_TYPE,  
      &type,  
      sizeof( type ) );  
  
UDF_CHECK_DESCRIBE( ctx, desc_rc );  
  
// Inform the server that the number of columns in our  
// result set is 1.  
num_cols = 1;  
desc_rc = ctx->describe_parameter_set  
    ( ctx,  
      0,  
      EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS,  
      &num_cols,  
      sizeof( num_cols ) );  
  
UDF_CHECK_DESCRIBE( ctx, desc_rc );  
  
// Inform the server that the type of column 1 in our  
// result set is int.  
type = DT_INT;  
desc_rc = ctx->describe_column_set  
    ( ctx,  
      0,  
      1,  
      EXTFNAPIV4_DESCRIBE_COL_TYPE,  
      &type,  
      sizeof( type ) );  
  
UDF_CHECK_DESCRIBE( ctx, desc_rc );  
  
}  
  
// The following describes will inform the server of various  
// optimizer related characteristics.  
if( ctx->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {  
  
    an_extfn_value          p1_value;  
    a_v4_extfn_estimate     num_rows;  
  
    // If the value of parameter 1 was constant, then we can  
    // inform the server how many distinct values will be.  
    desc_rc = ctx->describe_parameter_get
```



```

    ( ctx,
      1,
      EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE,
      &p1_value,
      sizeof( p1_value ) );

UDF_CHECK_DESCRIBE( ctx, desc_rc );

if( desc_rc != EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE ) {

    // Inform the server that this UDF will produce n rows.
    num_rows.value = *(a_sql_int32 *)p1_value.data;
    num_rows.confidence = 1;
    desc_rc = ctx->describe_parameter_set
    ( ctx,
      0,
      EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS,
      &num_rows,
      sizeof( num_rows ) );

    UDF_CHECK_DESCRIBE( ctx, desc_rc );

    // Inform the server that this UDF will produce n distinct
    // values for column 1 of its result set.
    desc_rc = ctx->describe_column_set
    ( ctx,
      0,
      1,
      EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES,
      &num_rows,
      sizeof( num_rows ) );

    UDF_CHECK_DESCRIBE( ctx, desc_rc );

}

}
}
}

```

This describe method has two primary functions:

- Inform the server what schema it supports.
- Inform the server of some known optimization attributes.

The describe function is called during several states. However, not all describe attributes are usable in every state. The describe method determines the state in which it is executing by checking the *current\_state* variable on the *a\_v4\_extfn\_proc* structure.

During the Annotation state, the **udf\_rg\_2** table udf informs the server that it has one parameter of type INTEGER and its result set contains a single column of type INTEGER. This is accomplished by setting these attributes:

- EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS

## Table UDFs and TPFs

- EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS
- EXTFNAPIV4\_DESCRIBE\_COL\_TYPE

If the information set in these `describe` methods does not match the procedure definition from the **CREATE PROCEDURE** statement, the `describe_parameter_set` and `describe_column_set` methods return `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE`. The `describe` method then sets an error to indicate to the client there is a mismatch.

This example uses the macro `UDF_CHECK_DESCRIBE` defined in `udf_utils.h` to check the return value from a `describe` and set an error, if it is not successful.

During optimization, the `udf_rg_2` table udf informs the server that it returns the same number of rows indicated in parameter one. Since the generated rows increment, the values are also unique. During optimization, only parameters that have a constant value are available. Use the `describe` attribute `EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE` to obtain the value of a constant parameter. Once the table udf determines that the attribute value is available, `udf_rg_2` sets `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS` and `EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES` to the value obtained.

### See also

- *udf\_rg\_3* on page 115
- *Implementing Sample Table UDF udf\_rg\_1* on page 106

### Running the Sample Table UDF in udf\_rg\_2.cxx

The sample `udf_rg_2` is included in a pre-compiled dynamic library called `libv4apiex` (extension is platform-dependent). Its implementation is in the `samples` directory in `udf_rg_2.cxx`.

1. To declare the table UDF to the server, issue:

```
CREATE OR REPLACE PROCEDURE udf_rg_2( IN num INT )
RESULT( c1 INT )
EXTERNAL NAME 'udf_rg_2@libv4apiex'
```

2. Select rows from the table UDF:

```
SELECT * FROM udf_rg_2( 5 );
```

3. To see how the `describe` affects behavior, issue a **CREATE PROCEDURE** statement that has a different schema than the one published by the table UDF. For example:

```
CREATE OR REPLACE PROCEDURE udf_rg_2( IN num INT, IN extra INT )
RESULT( c1 INT )
EXTERNAL NAME 'udf_rg_2@libv4apiex'
```

4. Select rows from the table UDF:

```
SELECT * FROM udf_rg_2( 5 );
```

IQ returns an error.

### udf\_rg\_3

The sample table UDF **udf\_rg\_3** builds upon **udf\_rg\_2** and has similar behavior. The procedure is called **udf\_rg\_3** and its implementation is in the `samples` directory in `udf_rg_3.cxx`.

The difference between the behavior of table UDFs **udf\_rg\_3** and **udf\_rg\_2** is that **udf\_rg\_3** generates only 100 unique values from 0 to 99, then repeats the sequence as necessary. This table UDF provides `_start_extfn` and `_finish_extfn` methods and has a modified version of `_describe_extfn` to account for the different semantics of the function.

Using `fetch_block` instead of `fetch_into` allows the table UDF to own the row block structure and use its own data layout. To illustrate this, the numbers generated are pre-allocated in an array. When a fetch is performed, rather than copying data into the server provided row block, the table UDF points the row block data pointers directly to the memory containing the data, thus preventing additional copies.

The following ancillary structure stores the numbers array. This structure also keeps a pointer to the allocated row block, which deallocates the row block.

```
#define MAX_ROWS 100
struct RowData {

    a_sql_int32          numbers[MAX_ROWS];
    a_sql_uint32         piece_len;
    a_v4_extfn_row_block * rows;

    void Init()
    {
        rows = NULL;
        piece_len = sizeof( a_sql_int32 );
        for( int i = 0; i < MAX_ROWS; i++ ) {
            numbers[i] = i;
        }
    }
};
```

This structure is allocated when execution of the table UDF starts, and deallocated when execution finishes, by providing `_start_extfn` and `_finish_extfn` methods in the `a_v4_extfn_proc_context`.

```
static void UDF_CALLBACK udf_rg_start(
    a_v4_extfn_proc_context *ctx )
/*****/
{
    // The start_extfn method is a good place to allocate our row
    // data. This method is called only once at the beginning of
    // execution.
    RowData *row_data = (RowData *)
        ctx->alloc( ctx, sizeof( RowData ) );
    row_data->Init();
}
```

## Table UDFs and TPFs

```
    ctx->_user_data = row_data;
}
```

The finish method performs two functions:

- Deallocates the RowData structure.
- Destroys the row block, if the table UDF encounters an error during fetch and cannot destroy the row block.

```
static void UDF_CALLBACK udf_rg_finish(
    a_v4_extfn_proc_context *ctx )
/*****/
{
    if( ctx->_user_data != NULL ) {

        RowData *row_data = (RowData *)ctx->_user_data;

        // If rows is non-null here, it means an error occurred and
        // fetch_block did not complete.
        if( row_data->rows != NULL ) {
            DestroyRowBlock( ctx, row_data->rows, 0, false );
        }

        ctx->free( ctx, ctx->_user_data );
        ctx->_user_data = NULL;
    }
}
```

The fetch\_block method is:

```
static short UDF_CALLBACK udf_rg_fetch_block(
    a_v4_extfn_table_context *tctx,
    a_v4_extfn_row_block **rows )
/*****/
{
    udf_rg_state * state = (udf_rg_state*)tctx->user_data;
    RowData * row_data = (RowData *)tctx->proc_context->_user_data;

    // First call, we need to build the row block
    if( *rows == NULL ) {

        // This function will build a row block structure that holds
        // MAX_ROWS rows of data. See udf_utils.cxx for details.
        *rows = BuildRowBlock( tctx->proc_context, 0, MAX_ROWS, false );

        // This pointer gets saved here because in some circumstances
        // when an error occurs, its possible we may have allocated
        // the rowblock structure but then never called back into
        // fetch_block to deallocate it. In this case, when the finish
        // method is called, we will end up deallocating it there.
        row_data->rows = *rows;
    }

    (*rows)->num_rows = 0;

    // The row block we allocated contains a max_rows member that was
    // set to the macro MAX_ROWS (100 in this case). This field is the
    // maximum number of rows that this row block can handle. We can
    // not exceed this number. We will also stop producing rows when
```

```

// we have produced the number of rows required as per the max_row
// in the state.
while( (*rows)->num_rows < (*rows)->max_rows &&
       state->next_row < state->max_row ) {

    a_v4_extfn_row      &row = (*rows)->row_data[ (*rows)->num_rows ];
    a_v4_extfn_column_data &col0 = row.column_data[ 0 ];

    // Row generation here is a matter of pointing the data
    // pointer in the rowblock to our pre-allocated array of
    // integers that was stored in the proc_context.
    col0.data = &row_data->numbers[(*rows)->num_rows % MAX_ROWS];
    col0.max_piece_len = sizeof( a_sql_int32 );
    col0.piece_len = &row_data->piece_len;
    state->next_row++;
    (*rows)->num_rows++;
}
if( (*rows)->num_rows > 0 ) {
    return 1;
} else {
    // When we are finished generating data, we can destroy the
    // row block structure.
    DestroyRowBlock( tctx->proc_context, *rows, 0, false );
    row_data->rows = NULL;
    return 0;
}
}

```

The first time this method is called, a row block is allocated using the helper function **BuildRowBlock**, which is in `udf_utils.cxx`. A pointer to this row block is saved in the `RowData` structure for later use.

Row generation is achieved by setting the data pointer for the column data to the address of the next number in sequence in the previously allocated numbers array. The `piece_len` pointer for the column data must also be initialized, by setting it to the address of the `piece_len` member of `RowData`. Since the rows are a fixed data length, this number is the same for all rows.

When `fetch` is called the last time and there is no more data to produce, the row block structure is destroyed using the **DestroyRowBlock** helper function in `udf_utils.cxx`.

To accommodate this table UDF generating only 100 unique values, `EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES` is set to a value of 100. This code excerpt from the describe method demonstrates this:

```

static void UDF_CALLBACK udf_rg_describe(
    a_v4_extfn_proc_context *ctx )
/*****
{
...
...
...

    a_v4_extfn_estimate distinct = {
        MAX_ROWS, 1.0
    };
}

```

## Table UDFs and TPFs

```
// Inform the server that this UDF will produce MAX_ROWS
// distinct values for column 1 of its result set.
desc_rc = ctx->describe_column_set
    ( ctx,
      0,
      1,
      EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES,
      &distinct,
      sizeof( distinct ) );

UDF_CHECK_DESCRIBE( ctx, desc_rc );
...
...
...
}
```

### See also

- *udf\_rg\_2* on page 111
- *Implementing Sample Table UDF udf\_rg\_1* on page 106

### Running the Sample Table UDF in udf\_rg\_3.cxx

The sample `udf_rg_3` is included in a precompiled dynamic library called `libv4apiex` (extension is platform-dependent). Its implementation is in the `samples` directory in `udf_rg_3.cxx`.

1. To declare the Table UDF to the server, issue:

```
CREATE OR REPLACE PROCEDURE udf_rg_3( IN num INT )
RESULT( c1 INT )
EXTERNAL NAME 'udf_rg_3@libv4apiex'
```

2. Select rows from the table UDF:

```
SELECT * FROM udf_rg_3( 200 );
```

This query produces values for `c1` from 0...99 followed by 0...99.

### apache\_log\_reader

The sample table UDF `apache_log_reader` reads the contents of an Apache access log or an Apache error log into table data. It is implemented in the file `apache_log_reader.cxx` in the `samples` directory.

A sample access log (`apache_access.log`) and sample error log (`apache_error.log`) are included in the `samples` directory.

The `apache_log_reader` sample opens the log file in the `_open_extfn` method. It reads in the data and parses it into the schema supported by the procedure in the `_fetch_into_extfn` method. It then closes the log file using the `_close_extfn` method.

**See also**

- `_open_extfn` on page 316
- `_fetch_into_extfn` on page 317
- `_close_extfn` on page 319

**Running the Sample Table UDF in `apache_log_reader.cxx`**

The sample `apache_log_reader` is included in a precompiled dynamic library called `libv4apiex` (extension is platform-dependent). Its implementation is in the `samples` directory in `apache_log_reader.cxx`.

1. To declare the table UDF to the server, issue:

```
create procedure apache_log_reader
(
    in file_name varchar(4000),
    in log_format varchar(32),
    in ip_padding varchar(1)
)
result
(
    ip_address varchar(15),
    log_name varchar(4000),
    user_name varchar(4000),
    access_time datetime,
    time_zone int,
    request varchar(4000),
    response int,
    bytes_sent int,
    referer varchar(4000),
    browser varchar(4000),
    error_type varchar(4000),
    error_msg varchar(4000)
)
external name 'apache_log_reader@libv4apiex'
```

2. Select rows from the table UDF. Use the full path to the access log when executing the SQL query.

```
SELECT * FROM apache_log_reader( 'apache_access.log', 'access',
null );
```

**udf\_blob**

The sample table UDF `udf_blob` illustrates how a table UDF or TPF can read LOB input parameters using the `blob` API.

`udf_blob` counts the number of occurrences of a letter in the first input parameter. The datatype of parameter **1** can be `LONG VARCHAR` or `VARCHAR(64)`. If the type is `LONG VARCHAR`, the table UDF uses the `blob` API to read in the value. If the type is `VARCHAR(64)`, the entire value is available using `get_value`.

This code snippet from the `_open_extfn` method illustrates how parameter **1** is read using the `blob` API:

## Table UDFs and TPFs

```
static short UDF_CALLBACK udf_blob_open(
    a_v4_extfn_table_context * tctx )
/*****/
{
...
...
    a_v4_extfn_blob      *blob      = NULL;

    ret = tctx->proc_context->get_value( tctx->args_handle, 2,
&value );
    UDF_SQLERROR_RT( tctx->proc_context,
        "get_value for argument 2 failed",
        ret == 1,
        0 );

    letter_to_find = *(char *)value.data;

    ret = tctx->proc_context->get_value( tctx->args_handle, 1,
&value );
    UDF_SQLERROR_RT( tctx->proc_context,
        "get_value for argument 1 failed",
        ret == 1,
        0 );

    if( EXTFN_IS_NULL(value) || EXTFN_IS_EMPTY(value) ) {
        state->return_value = 0;
        return 1;
    }

    if( EXTFN_IS_INCOMPLETE(value) ) {
        // If the value is incomplete, then that means we
// are dealing with a blob.
        tctx->proc_context->get_blob( tctx->args_handle, 1, &blob );
        return_value = ProcessBlob( tctx->proc_context,
            blob,
            letter_to_find );
        blob->release( blob );
    } else {
        // The entire value was put into the value pointer.
        return_value = CountNum( (char *)value.data,
value.piece_len,
letter_to_find );
    }
...
...
}
```

Parameter 1 is retrieved using `get_value`. If the value is empty or NULL, then no further processing is required. If the value is determined to be a blob using the macro `EXTFN_IS_INCOMPLETE`, then the Table UDF gets an instance of `a_v4_extfn_blob` using the `get_blob` method of `a_v4_extfn_proc_context`. The `ProcessBlob`



method reads from the blob to determine how many occurrences of the specified letter are present.

### See also

- *Blob (a\_v4\_extfn\_blob)* on page 195
- *\_open\_extfn* on page 316
- *get\_blob* on page 300
- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287

### Running the Sample Table UDF udf\_blob.cxx

The sample udf\_blob is included in a precompiled dynamic library called libv4apiex (extension is platform-dependent). Its implementation is in the samples directory in udf\_blob.cxx.

1. To declare the table UDF to the server, issue:

```
CREATE PROCEDURE udf_blob( IN data long varchar, letter char(1) )
RESULT ( c1 BIGINT )
EXTERNAL NAME 'udf_blob@libv4apiex'
```

2. Select rows from the table UDF:

```
set temporary option Enable_LOB_Variables = 'On';
create variable testblob long varchar;
set testblob = 'aaaaaaaaabbbbbbbbbbbb';
select * from udf_blob(testblob, 'a');
```

The supplied string contains the letter "a" 10 times.

## Query Processing States

---

The SQL statement that references a UDF goes through query processing states in the Sybase IQ server. In each of these states, the server uses the v4 API to communicate and negotiate with the UDF.

### See also

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207

### Initial State

Initial state on the server. The only UDF method called during the Initial state is `_start_extfn`.

The server calls the start method for each instance of the UDF created. If a query is executed by a single server thread, then the start method is called once. If a query is handled by several

## Table UDFs and TPFs

threads, or distributed across several nodes, the server creates different UDF instances and, as a result, the start method is called several times.

UDFs can set function instance level data within the `_user_data` field of the `a_v4_extfn_proc_context` structure, which is the argument to the start method.

## **Annotation State**

During the annotation state the server updates the parse tree with the metadata necessary for efficient and correct query optimization.

The `[_enter_state]`, `_describe_extfn`, and `[_leave_state]` methods are called. The `_enter_state` and `_leave_state` methods are optional and called if provided by the UDF.

The annotation state is represented in the v4 API by `EXTFNAPIV4_STATE_ANNOTATION` from the `a_v4_extfn_state` enumeration:

```
typedef enum a_v4_extfn_state {  
    ... EXTFNAPIV4_STATE_ANNOTATION, ...  
} a_v4_extfn_state;
```

As a UDF developer, you can perform some initial schema negotiation in this phase. Schema negotiation can occur either through the UDF describing to the server what it supports, or the UDF asking the server how it was declared.

When the UDF describes itself to the server, the server detects mismatches and returns SQL errors back to the client. For example, if a UDF describes that it requires four parameters and the SQL writer only declared the UDF with two, the server detects this and returns a SQL error back to the client.

When the UDF itself performs the validation by asking the server how it was declared, it adjusts its runtime execution accordingly: it either matches the declaratio, or it returns an error via the `set_error` v4 API. For example, assume you build a UDF that returns the maximum value of up to five input scalar integers. At runtime, the UDF determines how many input parameters were provided and adjusts its internal logic accordingly. SQL analysts could then create the procedure as:

```
CREATE PROCEDURE my_sum_2( IN a INT, IN b INT ) EXTERNAL  
"my_sum@my_lib"  
CREATE PROCEDURE my_sum_3( IN a INT, IN b INT, IN c INT ) EXTERNAL  
"my_sum@my_lib"
```

Both functions use the same underling implementation of `my_sum`. The UDF recognizes that there are only two parameters for `my_sum_2`, and attempts to sum parameters 1 and 2. For `my_sum_3`, the UDF sums parameters 1, 2 and 3.

As a UDF developer, you can obtain values for constant literal parameters only in the Annotation state. No other values are available until the Execution state. To get parameter

values during the annotation state use the `describe_parameter_get` method with the `PARAM_CONSTANT_VALUE` and `PARAM_IS_CONSTANT` attributes.

In the Annotation state, UDFs have access to schema `describe` attributes:

- `EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS`
- `EXTFNAPIV4_DESCRIBE_PARAM_NAME`
- `EXTFNAPIV4_DESCRIBE_PARAM_TYPE`
- `EXTFNAPIV4_DESCRIBE_PARAM_WIDTH`
- `EXTFNAPIV4_DESCRIBE_PARAM_SCALE`
- `EXTFNAPIV4_DESCRIBE_PARAM_IS_CONSTANT`
- `EXTFNAPIV4_DESCRIBE_PARAM_CONSTANT_VALUE`
- `EXTFNAPIV4_DESCRIBE_PARAM_TABLE_NUM_COLUMNS`
- `EXTFNAPIV4_DESCRIBE_COL_NAME`
- `EXTFNAPIV4_DESCRIBE_COL_TYPE`
- `EXTFNAPIV4_DESCRIBE_COL_WIDTH`
- `EXTFNAPIV4_DESCRIBE_COL_SCALE`
- `EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT`
- `EXTFNAPIV4_DESCRIBE_COL_CONSTANT_VALUE`

During the Annotation phase the UDF can set the above values to define its schema to the server. If the server detects a mismatch between what the UDF describes and the SQL procedure declaration, it returns an error. This technique is referred to as *self-describing*.

An alternative technique, schema validation, can be employed by the UDF. This involves the UDF getting the values for the schema `describe` types, and then setting an error if a mismatch is detected. With this approach, validation is left to the UDF, but the UDF can choose to support multiple schemas with a single implementation (for example, the ability to support multiple datatypes for a given parameter or being able to support varying number of parameters).

### See also

- *EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS Attribute (Get)* on page 274
- *EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS Attribute (Set)* on page 276
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_NAME Attribute (Get)* on page 239
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_NAME Attribute (Set)* on page 258
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_TYPE Attribute (Get)* on page 240
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_TYPE Attribute (Set)* on page 259
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_WIDTH Attribute (Get)* on page 241
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_WIDTH Attribute (Set)* on page 260
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_SCALE Attribute (Get)* on page 242
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_SCALE Attribute (Set)* on page 261
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_IS\_CONSTANT Attribute (Get)* on page 247
- *EXTFNAPIV4\_DESCRIBE\_PARAM\_IS\_CONSTANT Attribute (Set)* on page 263

## Table UDFs and TPFs

- *EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Get)* on page 248
- *EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Set)* on page 264
- *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Get)* on page 206
- *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Set)* on page 222
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Get)* on page 208
- *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Set)* on page 226
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Get)* on page 213
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Set)* on page 230
- *EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Get)* on page 214
- *EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Set)* on page 230

## Query Optimization State

During the Optimization state, the server is in the initial process of constructing a query plan. The server collects schema information and some preliminary statistical information.

The [`_enter_state`], `_describe_extfn`, and [`_leave_state`] methods are called. The `_enter_state` and `_leave_state` methods are optional, and called if provided by the UDF.

The query optimization state is represented in the v4 API by `EXTFNAPIV4_STATE_OPTIMIZATION` from the `a_v4_extfn_state` enumeration:

```
typedef enum a_v4_extfn_state {  
    ... EXTFNAPIV4_STATE_OPTIMIZATION, ...  
} a_v4_extfn_state;
```

Negotiations during the query optimization state include:

- The server and UDF determine the partitioning/ordering/clustering already specified for input tables.
- The server and UDF determine the partitioning/ordering required for input tables.
- The UDF declares physical properties (such as an ordering property) for the result table.
- The UDF describes any properties and statistics (for example, cost estimates) which can be used during the query optimization process.
  - Table scope estimates include:
    - **Number of rows** – the total number of rows present in the UDF during the execution state. This value is available for both the input table parameter and the returned table.

- **Row size** – an estimate of the average number of bytes in each row.
- Column scope estimates include:
  - **Distinct count** – the number of distinct values in a column over the total number of rows in a table. This value is available for both the input table parameter and the returned table.

In the Optimization state, UDFs have access to describe attributes:

- EXTFNAPIV4\_DESCRIBE\_PARM\_NAME
- EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE
- EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH
- EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE
- EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT
- EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND
- EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND
- EXTFNAPIV4\_DESCRIBE\_COL\_NAME
- EXTFNAPIV4\_DESCRIBE\_COL\_TYPE
- EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH
- EXTFNAPIV4\_DESCRIBE\_COL\_SCALE
- EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL
- EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT
- EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE
- EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER
- EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT

### See also

- *DEFAULT\_TABLE\_UDF\_ROW\_COUNT* Option on page 178
- *EXTFNAPIV4\_DESCRIBE\_PARM\_NAME* Attribute (Get) on page 239
- *EXTFNAPIV4\_DESCRIBE\_PARM\_NAME* Attribute (Set) on page 258
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE* Attribute (Get) on page 240
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE* Attribute (Set) on page 259
- *EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH* Attribute (Get) on page 241
- *EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH* Attribute (Set) on page 260
- *EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE* Attribute (Get) on page 242
- *EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE* Attribute (Set) on page 261
- *EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT* Attribute (Get) on page 247

## Table UDFs and TPFs

- *EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT Attribute (Set)* on page 263
- *EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Get)* on page 248
- *EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Set)* on page 264
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS Attribute (Get)* on page 249
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS Attribute (Set)* on page 264
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS Attribute (Get)* on page 250
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS Attribute (Set)* on page 265
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Get)* on page 251
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Set)* on page 266
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Get)* on page 252
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Set)* on page 268
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Get)* on page 254
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)* on page 269
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Get)* on page 255
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Set)* on page 271
- *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Get)* on page 206
- *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Set)* on page 222
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Get)* on page 208
- *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Set)* on page 226
- *EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL (Get)* on page 209
- *EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL (Set)* on page 227
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Get)* on page 213
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Set)* on page 230
- *EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Get)* on page 214
- *EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Set)* on page 230
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER (Get)* on page 215

- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER (Set)* on page 231
- *EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Get)* on page 220
- *EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Set)* on page 236

## **Plan Building State**

.During the plan building state, the server builds the query execution plan based on the best plan found during the query optimization state.

The [*\_enter\_state*], [*\_describe\_extfn*], and [*\_leave\_state*] methods are called. The *\_enter\_state* and *\_leave\_state* methods are optional and called if provided by the UDF.

The plan building state is represented in the v4 API by *EXTFNAPIV4\_STATE\_PLAN\_BUILDING* from the *a\_v4\_extfn\_state* enumeration:

```
typedef enum a_v4_extfn_state {
    ... EXTFNAPIV4_STATE_PLAN_BUILDING, ...
} a_v4_extfn_state;
```

At this point in query processing, the server determines what columns are needed from the UDF, and requests information about the columns needed from the table parameters.

If the UDF supports parallel processing, and if the server agrees that the query is eligible for parallelism, the server creates multiple instances of the UDF for distributed query processing.

In the Plan Building state, UDFs have access to all describe attributes.

As an example, the following code fragment queries the server to determine which columns are used:

```
a_sql_int32    rc;
rg_udf        *rgUdf = (rg_udf *)ctx->_user_data;
rg_table      *rgTable = rgUdf->rgTable;
a_sql_uint32  buffer_size = 0;

buffer_size = sizeof(a_v4_extfn_column_list) ( rgTable->
number_of_columns - 1 ) * sizeof(a_sql_uint32);
a_v4_extfn_column_list *ulist = (a_v4_extfn_column_list *)ctx->
alloc(
                                ctx,
                                buffer_size );
memset(ulist, 0, buffer_size);

rc = ctx->describe_parameter_get( ctx,
                                0,
                                EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS,
                                ulist,
                                buffer_size );

if( rc != buffer_size ) {
```

## Table UDFs and TPFs

```
    ctx->free( ctx, ulist );
    UDF_SQLERROR( PC(ctx), "Describe parameter type get failure.",
rc == buffer_size );
} else {
    rgTable->unused_col_list = ulist;
}
```

Assuming the above code fragment is from a Table UDF that produces 4 result set columns, and assuming the SQL statement was

```
SELECT c1, c2 FROM my_table_proc();
```

then the `describe` API returns only `c1` and `c2`. This lets the UDF optimize the production of the result set values.

### See also

- *Describe API* on page 204

## Execution State

During the execution state, the server makes an execution call into the UDF.

The execution plan, created in the plan building state, is used in the execution state to compute the result set of the SQL query.

These methods can be called: `[_enter_state]`, `_describe_extfn`, `evaluate_extfn`, `_open_extfn`, `_fetch_into_extfn`, `_fetch_block_extfn`, `_close_extfn`, `[_leave_state]`, and `_finish_extfn`.

The execution state is represented in the `a_v4_extfn_state` API by this enumeration:

```
typedef enum a_v4_extfn_state {
    ... EXTFNAPIV4_STATE_EXECUTING, ...
} a_v4_extfn_state;
```

In the execution state:

- Input table parameter rows and nonconstant scalar input parameter values are available.
- The UDF can open a result set on input table parameters, and fetch rows.

### *Executing Partition State*

If an input table parameter exists, and if a **PARTITION BY** clause exists in the SQL query, then the server invokes the UDF once per available partition.

## Row Block Data Exchange

A row block is the data transfer area between the producer and the consumer.

A table UDF can only produce rows. It can use an existing row block, or it can build its own row block.



A TPF can both produce and consume rows. A TPF produces rows in the same way a table UDF produces rows and can use an existing row block or build its own row block. A TPF can consume rows from an input table and can provide the producer with a row block, or request the producer to create its own row block.

### See also

- *Row Block (a\_v4\_extfn\_row\_block)* on page 305
- *Table (a\_v4\_extfn\_table)* on page 306
- *Table Functions (a\_v4\_extfn\_table\_func)* on page 314
- *\_open\_extfn* on page 316
- *\_fetch\_into\_extfn* on page 317
- *\_fetch\_block\_extfn* on page 318
- *\_rewind\_extfn* on page 318
- *\_close\_extfn* on page 319

## Fetch Methods for Row Blocks

The fetch methods for row blocks are `_fetch_into_extfn` and `_fetch_block_extfn`. These methods are part of the `a_v4_extfn_table_func` structure.

When producing data, if the table UDF or TPF builds its own row block, the UDF must provide the `fetch_block` API method. If the UDF does not build its own row block, the UDF must provide the `fetch_into` API method.

When consuming data, if the TPF builds its own row block, the UDF calls the `fetch_into` method on the producer. If the TPF does not build its own row block, the TPF must call the `fetch_block` method on the producer.

The UDF can select which fetch method to use for data production and consumption. In general, these guidelines apply:

- **fetch\_into** – Use this API when the consumer owns the memory for the data transfer area and requests that the producer use this area. In this scenario, the consumer cares about how the data transfer area is set up, and the producer performs the necessary data copies into this area.
- **fetch\_block** – Use this API when the consumer does not care about the format of the data transfer area. `fetch_block` requests the producer to create a data transfer area and provides a pointer to that area. The consumer owns the memory and the consumer is responsible for copying data from this area.

### See also

- *Table Parameterized Functions* on page 136
- *fetch\_into* on page 308
- *fetch\_block* on page 311

### **The fetch\_block Method**

Use the `fetch_block` method for underlying data storage.

The `fetch_block` method is used as an entry point when the consumer does not need the data in a particular format. `fetch_block` requests that the producer create a data transfer area and provide a pointer to that area. The consumer owns the memory and takes responsibility for copying data from this area.

The `fetch_block` method is more efficient than `fetch_into`, if the consumer does not need a specific layout. The `fetch_block` call provides a row block that can be populated, and this block is passed on the next `fetch_block` call. This method is part of the `a_v4_extfn_table_context` structure.

If the underlying data storage does not map easily to the row block structure, the UDF can simply point the row block to addresses in its memory. This prevents unnecessary data copies to satisfy another memory layout scheme.

The API uses a data transfer area that is defined by the structure `a_v4_extfn_row_block`, which is defined as a set of rows, where each row is defined as a set of columns. The row block creator can allocate enough storage to hold a single row or a set of rows. The producer can fill the rows, but cannot exceed the maximum number of rows allocated for the row block. If the producer has additional rows, the producer informs the consumer by returning the numeric value 1 from the `fetch` method.

Fetch is executed against a table object, which is either the object produced as the result set of a table UDF or the object consumed as a result set of an input table parameter.

### **See also**

- *Using a Row Block to Produce Data* on page 131
- *fetch\_block* on page 311

### **The fetch\_into Method**

Use the `fetch_into` API when the consumer owns the memory for the data transfer area and requests that the producer use this area.

The `fetch_into` method is useful when the producer does not know how data should be arranged in memory. This method is used as an entry point when the consumer has a transfer area with a specific format. The `fetch_into()` function writes the fetched rows into the provided row block. This method is part of the `a_v4_extfn_table_context` structure.

The API uses a data transfer area that is defined by the structure `a_v4_extfn_row_block`, which is defined as a set of rows, where each row is defined as a set of columns. The creator of the row block can allocate enough storage to hold a single row or a set of rows. The producer can fill the rows, but cannot exceed the maximum number of rows allocated for the row block. If the producer has additional rows, the producer informs the consumer by returning the numeric value 1 from the `fetch` method.

This API enables consumers to optionally construct the row block, such that the data pointers refer to its own data structures. This allows the producer to directly populate memory within the consumer. A consumer may not want to do this, if data cleansing or validation checks are required first.

Fetch is executed against a table object, which is either the object produced as the result set of a table UDF or the object consumed as a result set of an input table parameter.

### See also

- *Using a Row Block to Produce Data* on page 131
- *fetch\_into* on page 308

## Using a Row Block to Produce Data

A table UDF or TPF can use row block structures to produce data.

The `a_v4_extfn_row_block` row block has three fields:

- **max\_rows** – How many table rows the row block can store in a piece of memory.
- **num\_rows** – The number of rows actually produced or available for consumption. Cannot be larger than `max_rows`.
- **row\_data** – The array of rows produced or available for consumption. Each row is an `a_v4_extfn_row` structure.

### See also

- *Table UDF Implementation Examples* on page 105
- *fetch\_into* on page 308
- *fetch\_block* on page 311
- *Row Block (a\_v4\_extfn\_row\_block)* on page 305
- *Row (a\_v4\_extfn\_row)* on page 304
- *Column Data (a\_v4\_extfn\_column\_data)* on page 200

## Producing Data Using fetch\_into

Use the `fetch_into` API method to produce data.

1. Set `num_rows` to a value based on the number of rows produced in the fetch call.
2. For each row produced, set the `row_status` flag of `a_v4_extfn_row` to 1 (available) or 0 (not available). The default value is 1.
3. For each column (`a_v4_extfn_column_data`) in the row set:

Options	Description
<code>is_null</code>	Set to true, if the value returned is NULL. The default is false.

Options	Description
<b>data</b>	The data returned must be copied into this pointer
<b>piece_len</b>	The actual length of data returned. For fixed-length data types, this cannot exceed <i>max_piece_len</i> . Defaults to <i>max_piece_len</i> for fixed data types.

- For each column, return 1 to indicate rows produced, and return 0 to indicate otherwise.

### **Producing Data Using fetch\_block**

Use the `fetch_block` API method to produce data.

- Set `max_rows` to the number of rows the producer-allocated row block structure can hold.
- On the first fetch call, allocate a row block structure that can hold `max_rows`.
- Set `num_rows` to a value based on the number of rows produced in the fetch call.
- For each row produced, set the `row_status` flag of `a_v4_extfn_row` to 1 (available) or 0 (not available). The default value is 1.
- For each column (`a_v4_extfn_column_data`) in the row set:

Options	Description
<b>null_value</b>	Indicates the value is used to indicate NULL
<b>null_mask</b>	Identifies the bits that represent the NULL value.
<b>is_null</b>	If the value is NULL, set <code>is_null</code> to a value such that <code>(*(cd-&gt;is_null) &amp; cd-&gt;&gt;null_mask) == cd-&gt;&gt;null_value</code> .
<b>data</b>	Set this pointer to the area in the producer's memory containing the data to be returned.
<b>piece_len</b>	The actual length of data being returned. For fixed-length data types, this cannot exceed <i>max_piece_len</i> . This value defaults to <i>max_piece_len</i> for fixed data types.

- Return 1 from `fetch_into` to indicate rows were produced, and return 0 to indicate otherwise. On the last fetch call, deallocate any memory that is allocated for the row block structure.

## **Row Block Allocation**

Row block allocation is required when a producer produces data using the `fetch_block` method or when the consumer uses the `fetch_into` method to retrieve data.

`udf_utils.cxx` contains sample code that illustrates how to allocate and deallocate a row block.

These relevant data structures in the `extfnapi_v4.h` header file are used when allocating a row block:

```

typedef struct a_v4_extfn_column_data {
    a_sql_byte      *is_null;
    a_sql_byte      null_mask;
    a_sql_byte      null_value;

    void            *data;
    a_sql_uint32    *piece_len;
    size_t          max_piece_len;

    void            *blob_handle;
} a_v4_extfn_column_data;

typedef struct a_v4_extfn_row {
    a_sql_uint32    *row_status;
    a_v4_extfn_column_data *column_data;
} a_v4_extfn_row;

typedef struct a_v4_extfn_row_block {
    a_sql_uint32    max_rows;
    a_sql_uint32    num_rows;
    a_v4_extfn_row *row_data;
} a_v4_extfn_row_block;

```

When allocating a row block, the developer must decide how many rows the row block is capable of holding, how many columns each row has, and the number of bytes required for each of those columns.

For a row block of size  $m$ , where each row has  $n$  columns, the developer must allocate an array of  $m$  `a_v4_extfn_row` structures. For each row in this array, the developer must allocate  $n$  `a_v4_extfn_column_data` structures.

These tables outline allocation requirements for each member of the row block structures:

**Table 2. a\_v4\_extfn\_row\_block Structure**

Field	Requirement
max_rows	Set to the number of rows this row block can hold
num_rows	Initialize to zero. Is set to the number of actual rows a row block contains during usage
*row_data	Allocate an array containing max_rows of <code>a_v4_extfn_row</code> structures

**Table 3. a\_v4\_extfn\_row Structure**

Field	Requirement
*row_status	Allocate enough memory to hold an <code>a_sql_uint32</code>

Field	Requirement
*column_data	Allocate an array containing the number of columns in the result set of a_v4_extfn_column_data structures

Table 4. a\_v4\_extfn\_column\_data Structure

Field	Requirement
*is_null	Allocate enough memory to hold an <b>a_sql_byte</b>
null_mask	Set to a value such that the formula (*is_null & null_mask) == null_value indicates a column value is NULL
null_value	Set to a value such that the formula (*is_null & null_mask) == null_value indicates a column value is NULL
*data	Allocate an array of bytes large enough to hold the data for the data type of the column
*piece_len	Allocate enough memory to hold an <b>a_sql_uint32</b>
max_piece_len	Set to the maximum width for the column
*blob_handle	Always owned by the server. Initialize to NULL.

**See also**

- *SQL Data Types* on page 9
- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287

## Table UDF Query Plan Objects

---

The table UDF values and TPF values visible in the query plan.

- **Blocks Fetched** – shows the number of chunks used to transfer all the data produced by the UDF. This value equates to the number of times the server called the fetch method of the UDF.
- **Maximum rows per \_fetch\_into\_extfn** – (visible only if the UDF is using \_fetch\_into\_extfn.) Displays the maximum number of rows a UDF can produce on each call to \_fetch\_into\_extfn as determined by the server.
- **Minimum/Maximum values for an output column** – displays minimum/maximum values per column if the UDF has set them via extfnapi\_v4\_describe\_col\_maximum\_value. Minimum/maximum appear only for arithmetic data type columns.

- **ORDER BY node (TPF only)** – for a TPF, the query plan shows an ORDER BY node as a child of the TPF SubQuery node. The ORDER BY node indicates that the data is ordered as it flows into the table parameter.
- **Output Row Width** – (visible only if the UDF is using `_fetch_into_extfn`.) Shows the width of an output column in bytes. This value is used in calculating the maximum number of rows.
- **TableUDF node** – represents an instance of a table UDF in the query. The TableUDF node is a leaf node.
- **TPF node (TPF only)** – same as the TableUDF node except that TPF node permits use of an input table parameter. Unlike a TableUDF node which is a leaf node, the TPF is an interior node with at most one child.
- **TPF SubQuery node (TPF only)** – child of the TPF node. Represents the subquery for the input table argument.
- **UDF Library** – UDF library file name. Shows the full path on disk from which the dynamic library implementing the UDF was loaded.
- **Uniqueness of an output column** – reflects the value set by `extfnapi4_describe_col_is_unique`.
- **TABLE\_UDF\_ROW\_BLOCK\_SIZE\_KB** – option value displays in query plan statistics if you specify a value other than 128KB.

## Enabling Memory Tracking

---

Enable memory tracking to help you locate memory leaks in your UDFs, and to free the leaked memory. Memory tracking imposes a performance penalty.

Enabling memory tracking tracks all invocations of `a_v4_extfn_proc_context alloc` and `a_v4_extfn_proc_context free`. An allocations without a matching free is logged to the `iqmsg` file.

1. Ensure the **external\_UDF\_execution\_mode** is set to 1 or 2 (validation mode or tracing mode).
2. Use the `alloc` and `free` methods of `a_v4_extfn_proc_context`.

### See also

- *alloc* on page 297
- *free* on page 298

## Table Parameterized Functions

---

A Table parameterized function (TPF) is an extension of a table UDF that accepts table input parameters in addition to scalar input parameters.

You can configure user-specified partitioning for your TPF. The UDF developer can declare a partitioning scheme that breaks down the dataset into smaller pieces of query processing that you distribute across multiplex nodes. This enables you to execute the TPF in parallel in a distributed server environment over partitions of row sets. The query engine supports massive parallelization of TPF processing.

---

**Note:** Multiplex requires a separate license. See *Using Sybase IQ Multiplex*.

---

### Learning Roadmap for TPF Developers

Develop a C or C++ TPF.

This roadmap assumes:

- You have a C or C++ development environment on your machine.
- For the optional data partitioning capability, you have a multiplex environment. See the *Using Sybase IQ Multiplex* guide.

Task	See
Familiarize yourself with table UDF development.	<i>Learning Roadmap for Table UDF Developers</i> on page 97
Follow the recommended procedure for creating a TPF.	<i>Developing a TPF</i> on page 136 <i>TPF Implementation Examples</i> on page 156
Establish a table context for the input table and consume table rows from it.	<i>Consume Table Parameters</i> on page 137
(Optional) Order incoming data.	<i>Order Input Table Data</i> on page 140
(Optional) Partition the incoming data to enable parallel TPF processing in your multiplex.	<i>Partitioning Input Data</i> on page 140

### Developing a TPF

Review the major steps required to develop a TPF.

1. Perform the same steps required to develop a table UDF.
2. Consume input parameters.
3. (Optional) Order input table data.
4. (Optional) Partition input table data.



## 5. (Optional) Enable parallel TPF processing.

### See also

- *Consume Table Parameters* on page 137
- *Order Input Table Data* on page 140
- *Partitioning Input Data* on page 140
- *\_open\_extfn* on page 316
- *\_fetch\_into\_extfn* on page 317
- *\_fetch\_block\_extfn* on page 318
- *\_rewind\_extfn* on page 318
- *\_evaluate\_extfn* on page 286
- *Developing a Table UDF* on page 103

## Consume Table Parameters

A table parameter is a non-constant parameter. This means that the TPF must be in the execution state to retrieve table parameters.

The TPF can retrieve the table parameter from these methods:

- *\_open\_extfn*
- *\_fetch\_into\_extfn*
- *\_fetch\_block\_extfn*
- *\_rewind\_extfn*
- *\_evaluate\_extfn*

To consume a table parameter, the TPF must:

### Obtain a Table Object

The TPF obtains a table object for the table parameter using the `get_value` method of `a_v4_extfn_proc_context`.

A table object (`a_v4_extfn_table`) can initiate retrieving rows from an input table. The following code snippet illustrates how `get_value` obtains a table object for parameter `1`. For simplicity, this code assumes that parameter `1` is a table.

```
a_v4_extfn_value      value;
a_v4_extfn_table *    table;

ctx->get_value( args_handle,
                1,
                &value );

table = (a_v4_extfn_table *)value.data;
```

### See also

- *get\_value* on page 289

### Open the Result Set

Once a table object has been obtained using `get_value`, the TPF must open a result set on the table object using the `open_result_set` method of `a_v4_extfn_proc_context` before it can fetch any rows.

Calling `open_result_set` returns an instance of `a_v4_extfn_table_context` that the TPF can use to process table data. It also saves the table object in the `table` member of the `a_v4_extfn_table_context` object.

The following code snippet illustrates how `open_result_set` gets an instance of `a_v4_extfn_table_context` for fetching rows:

```
a_v4_extfn_table_context * rs = NULL;
ctx->open_result_set( ctx,
                    (a_v4_extfn_table *)value.data,
                    &rs );
```

### See also

- *open\_result\_set* on page 298
- *Table Context (a\_v4\_extfn\_table\_context)* on page 306

### Fetch from the Result Set

The TPF fetches table data from an input table using an open result set.

Fetching is accomplished by calling either `fetch_block` or `fetch_into` on the `a_v4_extfn_table_context` object returned from `open_result_set`. The TPF can choose which fetch method to use. If `fetch_block` is used, the server is responsible for rowblock allocation. If `fetch_into` is used, the TPF is responsible for row block allocation.

Each call to the fetch method returns either nothing, which is indicated by a return value of false, or returns a populated row block structure. The row block structure can then be used to consume the table data.

### See also

- *fetch\_into* on page 308
- *fetch\_block* on page 311
- *Row Block Data Exchange* on page 128

### **Consume Table Data Using a Row Block**

The TPF consumes table data using the `fetch_into` or `fetch_block` row block structures.

Each successful call to either `fetch_into` or `fetch_block` populates a `a_v4_extfn_row_block` structure.

The `a_v4_extfn_row_block` members are:

- **max\_rows** – the number of table rows the row block can store in a piece of memory.
- **num\_rows** – the number of rows actually produced or available for consumption. Cannot be larger than `max_rows`.
- **row\_data** – the array of rows produced or available for consumption. Each row is an `a_v4_extfn_row` structure.

Each row of table data in `row_data` has these members:

- **row\_stats** – indicates whether values for this row are present. A value of 1 means the values are present; 0 means the values are not.
- **column\_data** – the column data associated with this row.

The `column_data` members are:

Member	Description
<code>null_value</code>	The value representing NULL
<code>null_mask</code>	One or more bits used to represent the NULL value
<code>data</code>	Pointer to the data for the column. Depending on the type of fetch mechanism, either points to an address in the consumer, or an address where the data is stored in the UDF.
<code>piece_len</code>	The actual length of data for variable-length data types
<code>blob</code>	A non-NULL value means that the data for this column must be read using the <code>blob</code> API

### **See also**

- *Column Data (`a_v4_extfn_column_data`)* on page 200
- *Row Block (`a_v4_extfn_row_block`)* on page 305
- *Row (`a_v4_extfn_row`)* on page 304
- *get\_blob* on page 314

### **Close the Result Set**

Once the TPF is finished processing table data, it closes the open result set using the `close_result_set` method of `a_v4_extfn_proc_context`.

This code snippet illustrates `close_result_set` closing a result set.

```
ctx->close_result_set( ctx,  
                      rs ) );
```

### **Order Input Table Data**

Either a SQL Analyst or the UDF developer can order incoming data.

A SQL Analyst controls ordering by including the **ORDER BY** clause in a **SELECT** statement.

The UDF developer controls ordering by using the **DESCRIBE\_PARM\_TABLE\_ORDERBY** attribute.

Both methods result in the server ordering the incoming data, the results of which can be seen in the query plan in the Order node.

#### **See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Get)* on page 251
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Set)* on page 266

### **Partitioning Input Data**

Use the **PARTITION BY** clause to express and declare invocation partitioning in your parallel TPF.

As a SQL analyst, you can efficiently utilize system resources by leveraging the server query parallelism and distribution features available with the **PARTITION BY** clause in your SQL queries. Depending on the clause specified, the server may partition data into distinct value-based sets of rows or by row-range sets of rows.

- **Value-based partitions** – determined by key values on an expression. These partitions provide value when a computation depends on seeing all rows of the same value for aggregation.
- **Row-based partitions** – simple and efficient means to divide a computation into multiple streams of work. Used when a query must be executed in parallel.

You can express a design for partition via the **PARTITION BY <expr>** clause on the **TABLE** parameter to a TPF. UDF developers can utilize the table parameter metadata attribute **EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY** to programmatically declare that the UDF requires partitioning before invocation can proceed. The UDF can inquire to the partition to enforce it, or to dynamically adapt the partitioning.

**See also**

- *Parallel TPF PARTITION BY Examples Using EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 143
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Get)* on page 252
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Set)* on page 268
- *V4 API describe\_parameter and EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 141

### **V4 API describe\_parameter and EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY**

You can use `describe_parameter_set` and `describe_parameter_get` for partitioning an input table parameter for required columns.

*Declaration*

The `describe_parameter` API has two declarations.

*describe\_parameter\_set Declaration*

```
a_sql_int32 (SQL_CALLBACK *describe_parameter_set)(
    a_v4_extfn_proc_context    *cntxt,
    a_sql_uint32               arg_num,
    a_v4_extfn_describe_parm_type describe_type,
    void                       *describe_buffer,
    size_t                     describe_buffer_
)
```

*describe\_parameter\_get Declaration*

```
a_sql_int32 (SQL_CALLBACK *describe_parameter_get)(
    a_v4_extfn_proc_context    *cntxt,
    a_sql_uint32               arg_num,
    a_v4_extfn_describe_parm_type describe_type,
    const void                 *describe_buffer,
    size_t                     describe_buffer_
)
```

*Usage*

In order to use these APIs, the **arg\_num** must refer to a TABLE parameter, and the **describe\_buffer** must refer to the type of memory block `a_v4_extfn_column_list` structure.

```
typedef struct a_v4_extfn_column_list {
    a_sql_int32    number_of_columns;
    a_sql_uint32  column_indexes[1];
} a_v4_extfn_column_list;
```

The structure field **number\_of\_columns** must have one of these values:

- Positive integer N, where N indicates the number of columns present in the partition by list.

## Table UDFs and TPFs

- 0, which indicates **PARTITION BY ANY**.
- -1, which indicates **NO PARTITION BY**.

This enumerated type is defined in the `extfnapi4.h` header file:

```
typedef struct a_v4_extfn_column_list {
    a_sql_int32      number_of_columns;
    a_sql_uint32    column_indexes[1];
} a_v4_extfn_column_list;
```

You can use the `v4_extfn_partitionby_col_num` enumerated type to build the column list structure and execute the `describe_parameter_set` and `describe_parameter_get` API to inform the server of its requirements and to determine which input columns have been partitioned. The execution of `describe_parameter_set` and `describe_parameter_get` APIs can have following scenarios:

### *describe\_parameter\_set* Scenarios

column Index Scenarios	Description
{ 1, 1 }	Input table column #1 is partitioned as per UDF request.
{ 2, 3, 1 }	Input table columns #3 and #1 are partitioned as per UDF request.
{ 0 }	UDF can support any form of input table partitioning as per UDF request.

### *describe\_parameter\_get* Scenarios

column Index Scenarios	Description
{ 1, 2 }	Input table column #2 is being partitioned on.
{ 2, 1, 2 }	Input table columns #1 and #2 are being partitioned on.
{ 0 }	Input table being partitioned by a noncolumn based scheme.
NULL	No runtime partitioning is provided.

**Note:** A **PARTITION BY** expression other than **PARTITION BY ANY** or **PARTITION BY NONE** must appear in the select list for the input query.

### See also

- *Describe API* on page 204
- *Partition By Column Number (a\_v4\_extfn\_partitionby\_col\_num)* on page 303
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Get)* on page 240

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS Attribute (Get)* on page 249
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Set)* on page 259
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS Attribute (Set)* on page 264
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224

### **Parallel TPF PARTITION BY Examples Using EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY**

Develop partitioning using the **PARTITION BY <expr>** clause on the **TABLE** parameter to a TPF function. As a UDF developer, use the table parameter metadata attribute `EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY` to programmatically declare that the UDF requires partitioning before invoking it.

The examples illustrate:

- Various SQL writer scenarios where the UDF describes partitioning requirements to the server
- Valid queries and invalid queries (SQL exceptions) for each scenario
- How the server detects mismatches
- The various possible combinations that arise from usage of the **PARTITION BY** SQL clause and the `EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY` UDF attribute

### **See also**

- *Partitioning Input Data* on page 140

### **Example Procedure Definition**

An example procedure definition that supports TPF **PARTITION BY** clause examples.

All TPF **PARTITION BY** clause examples in this section assume that you first execute this procedure definition:

```
CREATE PROCEDURE my_tpf( arg1 TABLE( c1 INT, c2 INT ) )
  RESULTS ( r1 INT, r2 INT, r3 INT )
  EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );
```

### **See also**

- *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1* on page 144
- *describe\_parameter\_set Example # 2: Two-Column Partitioning* on page 146
- *describe\_parameter\_set Example # 3: Any-Column Partitioning* on page 148
- *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause* on page 150

## Table UDFs and TPFs

- *describe\_parameter\_set Example # 5: No Partitioning Support* on page 151
- *describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2* on page 153

### *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1*

An example UDF that informs the server to perform partitioning on column 1 (c1).

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
    if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
        a_sql_int32 rc = 0;
        a_v4_extfn_column_list pbcoll =
            { 1, // 1 column in the partition by list
              1 }; // column index 1 requires partitioning

        // Describe partitioning for argument 1 (the table)
        rc = ctx->describe_parameter_set(
ctx,
        1,
        EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
        &pbcoll,
        sizeof(pbcoll) );

        if( rc == 0 ) {
            ctx->set_error( ctx, 17000,
                "Runtime error, unable set partitioning requirements for
column." );
        }
    }
}
```

### See also

- *Example Procedure Definition* on page 143
- *describe\_parameter\_set Example # 2: Two-Column Partitioning* on page 146
- *describe\_parameter\_set Example # 3: Any-Column Partitioning* on page 148
- *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause* on page 150
- *describe\_parameter\_set Example # 5: No Partitioning Support* on page 151
- *describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2* on page 153



*SQL Writer Semantics for One-Column Partitioning on Column 1*

Example queries valid for one-column partitioning on column 1 (c1).

*Example 1*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.x ) )
```

In this example, the UDF describes to the server that the data is partitioned by the first column (T.x) and the SQL writer also explicitly requests partitioning on the same column. When the two columns match, the above query proceeds without any errors using this negotiated query:

```
my_tpf( TABLE( SELECT T.x, T.y FROM T )
  OVER ( PARTITION BY T.x ) )
V4 describe_parameter_get API returns: { 1, 1 }
```

*Example 2*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY ANY ) )
```

In this example, the UDF describes to the server that the data is partitioned by the first column (T.x) and the SQL writer only wants the query engine to execute the UDF on partitions. The server uses the UDF's preference for partitioning and as a result the same effective query in Example 1 is executed.

*Example 3*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T ) )

SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER ( PARTITION BY DEFAULT ) )
```

This example shows that the SQL writer does not include the **PARTITION BY** clause or the **PARTITION BY DEFAULT** clause as part of the input table query specification. In this case, the specification requested by the UDF applies, which is to perform partitioning on column T.x.

*SQL Exceptions for One Column Partitioning on Column 1*

Example queries not valid for one column partitioning on column 1 (c1). Each example raises a SQL exception.

*Example 1*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y ) )
```

## Table UDFs and TPFs

In this example the UDF describes to the server that the data is partitioned by the first column (T.x) and that the SQL writer is also explicitly requesting partitioning on a different column (T.y) which conflicts with what the UDF is requesting and as a result the server returns a SQL error.

### Example 2

```
SELECT * FROM my_tpf(  
  TABLE( SELECT T.x, T.y FROM T )  
  OVER( NO PARTITION BY ) )
```

This example conflicts with the request made by the UDF because the SQL writer does not want the input table partitioned and as a result the server returns a SQL error.

### Example 3

```
SELECT * FROM my_tpf(  
  TABLE( SELECT T.x, T.y FROM T )  
  OVER( PARTITION BY T.x, T.y ) )
```

In this example the UDF describes to the server that the data is partitioned by the first column (T.x) and the SQL writer requests partitioning on columns (T.x and T.y) which conflicts with what the UDF is requesting and as a result the server returns a SQL error.

### describe\_parameter\_set Example # 2: Two-Column Partitioning

An example UDF that informs the server to perform partitioning on column 1 (c1) and column 2 (c2).

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context  
*ctx )  
{  
  if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {  
    a_sql_int32 rc = 0;  
    a_v4_extfn_column_list pbcoll =  
{ EXTFNAPIV4_PARTITION_BY_COLUMN_ANY };  
  
    // Describe partitioning for argument 1 (the table)  
    rc = ctx->describe_parameter_set(  
ctx,  
    1,  
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,  
    &pbcoll,  
    sizeof(pbcoll) );  
  
    if( rc == 0 ) {  
      ctx->set_error( ctx, 17000,  
"Runtime error, unable set partitioning requirements for  
column." );  
    }  
  }  
}
```

**See also**

- *Example Procedure Definition* on page 143
- *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1* on page 144
- *describe\_parameter\_set Example # 3: Any-Column Partitioning* on page 148
- *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause* on page 150
- *describe\_parameter\_set Example # 5: No Partitioning Support* on page 151
- *describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2* on page 153

**SQL Writer Semantics for Two-Column Partitioning**

Example queries valid for two-column partitioning on column 1 (c1) and column 2 (c2).

**Example 1**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y, T.x ) )
```

In this example, the UDF describes to the server that the data is partitioned by columns T.y and T.x. The SQL writer also requests partitioning on the same column. When the two columns match, the above query proceeds without any errors using this negotiated query:

```
my_tpf( TABLE( SELECT T.x, T.y FROM T )
  OVER ( PARTITION BY T.y, T.x ) )
V4 describe_parameter_get API returns: { 2, 2, 1 }
```

**Example 2**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY ANY ) )
```

In this example, the SQL writer does not specify a specific column for partitioning. Instead the SQL writer partitions the input table. The UDF requests partitioning on columns T.y and T.x, and as a result, the server partitions the input data on the columns T.y and T.x.

**Example 3**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )

SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER ( PARTITION BY DEFAULT ) )
```

This example shows that the SQL writer does not include **PARTITION BY** clause or the **PARTITION BY DEFAULT** clause. The server uses the partition requested by the UDF, and since

## Table UDFs and TPFs

the UDF describes that it requires partitioning on columns T.y and T.x, the server executes the query by performing partitioning on columns T.y and T.x.

### *Example 4*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.x,T.y ) )
```

This example is semantically identical to Example 1. The ordering of the two columns are different, but within a given partition, the values for columns T.x and T.y stay the same. Both columns (T.x, T.y) and columns (T.y, T.x) result in the same logical partitioning of data.

### *SQL Exceptions for Two-Column Partitioning*

Invalid example queries for two-column partitioning on column 1 (c1) and column 2 (c2). Each example raises a SQL exception.

### *Example 1*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( NO PARTITION BY ) )
```

This example conflicts with the request made by the UDF because the SQL writer does not want the input table partitioned. As a result, the server returns a SQL error.

### *Example 2*

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.x ) )

SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y ) )
```

In this example the UDF describes to the server that the data is partitioned by columns T.y and T.x, while the SQL writer requests the partitioning on either column T.y or T.x. which conflicts with what the UDF is requesting. As a result, the server returns a SQL error.

### *describe\_parameter\_set Example # 3: Any-Column Partitioning*

An example UDF that informs the server it can perform partitioning on any column.

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
  if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
    a_sql_int32 rc = 0;
    a_v4_extfn_column_list pbcol =
    { EXTFNAPIV4_PARTITION_BY_COLUMN_ANY };
  }
}
```

```

        // Describe partitioning for argument 1 (the table)
        rc = ctx->describe_parameter_set(
ctx,
        1,
        EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
        &pbcol,
        sizeof(pbcol) );

        if( rc == 0 ) {
            ctx->set_error( ctx, 17000,
                "Runtime error, unable set partitioning requirements for
column." );
        }
    }
}

```

### See also

- *Example Procedure Definition* on page 143
- *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1* on page 144
- *describe\_parameter\_set Example # 2: Two-Column Partitioning* on page 146
- *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause* on page 150
- *describe\_parameter\_set Example # 5: No Partitioning Support* on page 151
- *describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2* on page 153

### SQL Writer Semantics for Any-Column Partitioning

Example queries valid for any-column partitioning.

#### Example 1

```

SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T )
    OVER( PARTITION BY T.x ) )

```

In this example, the UDF describes to the server that the data is partitioned by the first column (T.x) and the SQL writer also explicitly requests partitioning on the same column. When the two columns match, the above query proceeds without any errors using this negotiated query:

```

my_tpf( TABLE( SELECT T.x, T.y FROM T )
    OVER ( PARTITION BY T.y, T.x ) )
V4 describe_parameter_get API returns: { 2, 2, 1 }

```

#### Example 2

```

SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T )
    OVER( PARTITION BY ANY ) )

```

In this example, neither the SQL writer nor the UDF specify a specific column for partitioning. Instead the SQL writer partitions the input table and, as a result, the server arranges partitioning in a nonvalue-based scheme and the data is partitioned over ranges of rows.

#### *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause*

An example UDF that informs the server that it cannot perform partitioning on any columns, because the UDF does not support the **PARTITION BY ANY** clause.

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
    // No describe calls
}
```

#### See also

- *Example Procedure Definition* on page 143
- *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1* on page 144
- *describe\_parameter\_set Example # 2: Two-Column Partitioning* on page 146
- *describe\_parameter\_set Example # 3: Any-Column Partitioning* on page 148
- *describe\_parameter\_set Example # 5: No Partitioning Support* on page 151
- *describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2* on page 153

#### *SQL Writer Semantics for No Support for PARTITION BY ANY Clause*

Example queries valid when the UDF does not support the **PARTITION BY ANY** clause.

##### *Example 1*

```
SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T ))
```

This example shows that the SQL writer does not include **PARTITION BY** clause. The server uses the partition requested by the UDF and since the UDF does not supports any partitioning requirements, the server executes the query without performing any partitioning.

##### *Example 2*

```
SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T )
    OVER( NO PARTITION BY ))
```

In this example, the SQL writer requests the **NO PARTITION BY** clause as part of the input table query specification. As a result, the server executes the query with no runtime partitioning.

##### *Example 3*

```
SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T ))
```

```
OVER( PARTITION BY T.x))
```

In this example the UDF does not describe any partitioning requirements. However, the SQL writer requests partitioning by column T.x and as a result the server executes the query by performing partitioning on column T.x.

#### Example 4

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y))
```

In this example, the UDF does not describe any partitioning requirements. However, the SQL writer requests partitioning by column T.y. As a result, the server executes the query by performing partitioning on column T.y.

#### Example 5

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y, T.x))
```

In this example, the UDF does not describe any partitioning requirements. However, the SQL writer requests partitioning by columns T.y and T.x. As a result, the server executes the query by performing partitioning on columns T.y and T.x.

#### Example 6

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY ANY ))
```

In this example, the SQL writer requests **PARTITION BY ANY** partitioning. However, the UDF does not support any partitioning requirements. As a result, the server executes the query by performing row range partitioning.

#### describe parameter set Example # 5: No Partitioning Support

An example UDF that informs the server that it does not support any partitioning.

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
  if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
    a_sql_int32 rc = 0;
    a_v4_extfn_column_list pbcoll =
    { EXTFNAPIV4_PARTITION_BY_COLUMN_NONE };

    // Describe partitioning for argument 1 (the table)
    rc = ctx->describe_parameter_set(
```

```

ctx,
        1,
        EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
        &pbcol,
        sizeof(pbcol) );

    if( rc == 0 ) {
        ctx->set_error( ctx, 17000,
            "Runtime error, unable set partitioning requirements for
column." );
    }
}
}

```

### See also

- *Example Procedure Definition* on page 143
- *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1* on page 144
- *describe\_parameter\_set Example # 2: Two-Column Partitioning* on page 146
- *describe\_parameter\_set Example # 3: Any-Column Partitioning* on page 148
- *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause* on page 150
- *describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2* on page 153

### SQL Writer Semantics for No Partitioning Support

Valid example queries when the UDF does not support any partitioning.

#### Example 1

```

SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T )
    OVER( PARTITION BY ANY )

```

In this example, the SQL writer requests **PARTITION BY ANY** partitioning. However, the UDF does not support any partitioning, and as a result, the server executes the query without runtime partitioning.

#### Example 2

```

SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T )

SELECT * FROM my_tpf(
    TABLE( SELECT T.x, T.y FROM T )
    OVER ( PARTITION BY DEFAULT )

```

This example shows that the SQL writer does not include the **PARTITION BY** clause or the **PARTITION BY DEFAULT** clause. The server uses the partition requested by the UDF and since the UDF does not supports any partitioning, the server executes the query without performing any partitioning.



**Example 3**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( NO PARTITION BY )
```

In this example, the SQL writer requests no partitioning, and as a result, the server executes the query without runtime partitioning.

**SQL Exceptions for No Partitioning Support**

Invalid example queries, because the UDF does not support any partitioning. Each example raises a SQL exception.

**Example 1**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.x ))
```

This example results in a SQL error because the SQL writer requested partitioning on column T.x, and the UDF does not support any partitioning on any columns.

**Example 2**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y ))
```

This example results in a SQL error because the SQL writer requested partitioning on column T.y, and the UDF does not support any partitioning on any columns.

**Example 3**

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.y, T.x ))
```

This example results in a SQL error because the SQL writer requested partitioning on columns T.y and T.x, and the UDF does not support any partitioning on any columns.

**describe\_parameter\_set Example # 6: One-Column Partitioning on Column 2**

An example UDF that informs the server to perform partitioning on column 2 (c2).

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
  if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
    a_sql_int32 rc = 0;
    a_v4_extfn_column_list pbcoll =
      { 1, // 1 column in the partition by list
        2 }; // column index 2 requires partitioning
```

```

        // Describe partitioning for argument 1 (the table)
        rc = ctx->describe_parameter_set(
ctx,
        1,
        EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
        &pbcol,
        sizeof(pbcol) );

        if( rc == 0 ) {
            ctx->set_error( ctx, 17000,
                "Runtime error, unable set partitioning requirements for
column." );
        }
    }
}

```

### See also

- *Example Procedure Definition* on page 143
- *describe\_parameter\_set Example # 1: One-Column Partitioning on Column 1* on page 144
- *describe\_parameter\_set Example # 2: Two-Column Partitioning* on page 146
- *describe\_parameter\_set Example # 3: Any-Column Partitioning* on page 148
- *describe\_parameter\_set Example # 4: No Support for PARTITION BY ANY Clause* on page 150
- *describe\_parameter\_set Example # 5: No Partitioning Support* on page 151

### SQL Writer Semantics for One-Column Partitioning on Column 2

Valid example queries for one-column partitioning on column 2 (c2).

#### Example 1

```

SELECT * FROM my_tpf(
TABLE( SELECT T.x, T.y FROM T )
OVER( PARTITION BY T.y )

```

In this example, the UDF describes to the server that the data is partitioned by the first column (T.y), and the SQL writer also explicitly requests partitioning on the same column. When the two columns match, the above query proceeds without any errors using this negotiated query:

```

my_tpf( TABLE( SELECT T.x, T.y FROM T )
        OVER ( PARTITION BY T.y ) )
V4 describe_parameter_get API returns: { 1, 2 }

```

#### Example 2

```

SELECT * FROM my_tpf(
TABLE( SELECT T.x, T.y FROM T )
OVER( PARTITION BY ANY )

```

In this example the SQL writer does not specify a specific column for partitioning. Instead the SQL writer partitions the input table. The UDF requests partitioning on column T.y, and as a result, the server partitions the input data on the column T.y.

### Example 3

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )

SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER ( PARTITION BY DEFAULT )
```

This example shows that the SQL writer does not include the **PARTITION BY** clause or the **PARTITION BY DEFAULT** clause as part of the input table query specification. In this case, the specification requested by the UDF applies, which is to perform partitioning on column T.y.

### SQL Exceptions for One-Column Partitioning on Column 2

Invalid example queries for one-column partitioning on column 2 (c2). Each example raises a SQL exception.

#### Example 1

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.x )
```

In this example, the UDF describes to the server that the data is partitioned by the first column (T.y), and that the SQL writer is also explicitly requesting partitioning on a different column (T.x), which conflicts with what the UDF is requesting. As a result the server returns a SQL error.

#### Example 2

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( NO PARTITION BY )
```

This example conflicts with the request made by the UDF because the SQL writer does not want the input table partitioned. As a result the server returns a SQL error.

#### Example 3

```
SELECT * FROM my_tpf(
  TABLE( SELECT T.x, T.y FROM T )
  OVER( PARTITION BY T.x, T.y )
```

In this example, the UDF describes to the server that the data is partitioned by the first column (T.y), and the SQL writer requests partitioning on columns (T.x and T.y), which conflicts with what the UDF is requesting. As a result the server returns a SQL error.

## TPF Implementation Examples

Implementation examples start with a simple TPF and increase in complexity and functionality as the examples progress.

The TPF implementation examples are in the `samples` directory.

The examples are available in a precompiled dynamic library called `libv4apiex`. The extension of this library name is platform-dependent. This library includes the functions defined in `udf_main.cxx`, which contains the library-level functions, such as `extfn_use_new_api`. Put `libv4apiex` in a directory the server can read.

### tpf\_rg\_1

TPF sample `tpf_rg_1.cxx` is similar to the table UDF sample `udf_rg_2.cxx`. It produces rows of data based on an input parameter.

The number of rows generated is the sum of the values of the rows in a single input table. The output is the same as `udf_rg_2.cxx`.

The majority of the code for this sample is the same as `udf_rg_2.cxx`. The main differences are:

- The names of the implementing functions have the prefix `tpf_rg` instead of `udf_rg`. See the file `tpf_rg_1.cxx` for details.
- The implementation of `_describe_extfn` validates the schema of this example but does not estimate the number of rows generated.
- The implementation of `_open_extfn` reads rows from an input table to determine the number of rows to generate.

The `_describe_extfn` method accommodates the schema differences between `udf_rg_2.cxx` and this example. In particular, parameter 1 is a table with one integer column. This code snippet illustrates `_describe_extfn`:

```
static void UDF_CALLBACK tpf_rg_describe(
    a_v4_extfn_proc_context *ctx )
/*****/
{
    a_sql_int32      desc_rc;

    // The following describes will ensure that the schema defined
    // by the user matches the schema supported by this TPF
    // This is achieved by telling the server what our schema is
    // using describe_xxxx_set methods.
    if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
        ...

        // Inform the server that the type of parameter 1 is a TABLE
        type = DT_EXTFN_TABLE;
        desc_rc = ctx->describe_parameter_set
```

```

        ( ctx,
          1,
          EXTFNAPIV4_DESCRIBE_PARM_TYPE,
          &type,
          sizeof( type ) );

    UDF_CHECK_DESCRIBE( ctx, desc_rc );

    // Inform the server that the input table should have a single
    // column.
    num_cols = 1;
    desc_rc = ctx->describe_parameter_set
        ( ctx,
          1,
          EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS,
          &num_cols,
          sizeof( num_cols ) );

    UDF_CHECK_DESCRIBE( ctx, desc_rc );

    // Inform the server that the input table column is an integer
    type = DT_INT;
    desc_rc = ctx->describe_column_set
        ( ctx,
          1,
          1,
          EXTFNAPIV4_DESCRIBE_COL_TYPE,
          &type,
          sizeof( type ) );

    UDF_CHECK_DESCRIBE( ctx, desc_rc );

...
...
}
}

```

In `udf_rg_2.cxx`, the number of rows generated by the UDF may be available during the describe phase if the value is a constant. A table argument can never be constant, so its value is unavailable until the Execution state. For this reason, no optimizer estimate for the number of rows being generated is provided during the describe phase.

Calls to the describe only during the Annotation state have an effect in this example. Such calls will do nothing in the other states.

The `_open_extfn` method reads in rows from the input table and sums their values. As in the `udf_rg_2.cxx` example, the value of the first input parameter is retrieved using `get_value`. The difference here is that the type of the parameter is `a_v4_extfn_table` pointer. This code snippet illustrates `_open_extfn`:

```

static short UDF_CALLBACK tpf_rg_open(
    a_v4_extfn_table_context * tctx )
/*****
{

```

## Table UDFs and TPFs

```

an_extfn_value          value;
tpf_rg_state *          state          = NULL;
a_v4_extfn_table_context * rs          = NULL;
a_sql_uint32           num_to_generate = 0;

// Read in the value of the input parameter and store it away in a
// state object.  Save the state object in the context.
if( !tctx->proc_context->get_value( tctx->args_handle,
                                   1,
                                   &value ) ) {

    // Send an error to the client if we could not get the value.
    tctx->proc_context->set_error(
        tctx->proc_context,
        17001,
        "Error: Could not get the value of parameter 1" );

    return 0;
}

// Open a result set for the input table.
if( !tctx->proc_context->open_result_set( tctx->proc_context,
                                         ( a_v4_extfn_table * )value.data,
                                         &rs ) ) {
    // Send an error to the client if we could not open the result
    // set.
    tctx->proc_context->set_error(
        tctx->proc_context,
        17001,
        "Error: Could not open result set on input table." );

    return 0;
}

a_v4_extfn_row_block *  rbf           = NULL;
a_v4_extfn_row *        rfb           = NULL;
a_v4_extfn_column_data * cdfb         = NULL;
// When using fetch block to read rows from an input table, the
// server will manage the row block allocation.
while( rs->fetch_block( rs, &rbfb ) ) {

    // Each successful call to fetch will fill rows in the server
    // allocated row block.  The number of rows retrieved is
    // indicated by the num_rows member.
    for( unsigned int i = 0; i < rbf->num_rows; i++ ) {
        rfb = &(rbfb->row_data[i]);
        cdfb = &(rfb->column_data[0]);

        // Only consider non-null values.  To determine null we
        // have to use the following logic.
        if( (*cdfb->is_null & cdfb->>null_mask) != cdfb->>null_value )
        {
            num_to_generate += *(a_sql_int32 *)cdfb->data;
        }
    }
}

```

```

if( !tctx->proc_context->close_result_set( tctx->proc_context, rs ) )
{
    // Send an error to the client if we could not close the
    // result set.
    tctx->proc_context->set_error(
        tctx->proc_context,
        17001,
        "Error: Could not close result set on input table." );

    return 0;
}

// Allocate memory for the state using the a_v4_extfn_proc_context
// function alloc.
state = (tpf_rg_state *)
tctx->proc_context->alloc( tctx->proc_context,
    sizeof( tpf_rg_state ) );
// Start generating at row zero.
state->next_row = 0;

// Save the value of parameter 1
state->max_row = num_to_generate;

// Save the state on the context
tctx->user_data = state;

return 1;
}

```

Once you retrieve the table object using `get_value`, call `open_result_set` to read in rows of data from the table.

To read rows from the input table, the UDF can use `fetch_into` or `fetch_block`. When a UDF is fetching rows from an input table, it becomes a consumer of data. If the consumer (the UDF in this case) wants to be responsible for managing the row block structure, then the consumer must allocate their own row block structure and use `fetch_into` to retrieve the data. Alternatively, if the consumer wants the producer (the server in this case) to manage the row block structure, then use `fetch_block`. `tpf_rg_1` demonstrates the latter.

Using an open result set, `tpf_rg_1` retrieves rows of data from the server by calling `fetch_block` repeatedly. Each successful call to `fetch_block` populates the server allocated row block structure with up to `num_rows` rows. In `tpf_rg_1`, the value of column 1 for each row is added to a total. As in the `udf_rg_2.cxx` example, this total is saved in the `a_v4_extfn_proc_context` state to be used later.

### See also

- *Describe API* on page 204
- *\_open\_extfn* on page 316
- *Table (a\_v4\_extfn\_table)* on page 306
- *\_fetch\_block\_extfn* on page 318

### Running the Sample TPF in tpf\_rg\_1

The sample `tpf_rg_1` is included in a precompiled dynamic library called `libv4apiex` (extension is platform-dependent). Its implementation is in the `samples` directory in `tpf_rg_1.cxx`.

1. Declare the TPF to the server.

```
CREATE OR REPLACE PROCEDURE tpf_rg_1( IN tab TABLE( num INT ) )
RESULT( c1 INT )
EXTERNAL NAME 'tpf_rg_1@libv4apiex';
```

2. Declare a table to use as input to the TPF.

```
CREATE TABLE test_table( val int );
```

3. Insert rows into the table:

```
INSERT INTO test_table values(1);
INSERT INTO test_table values(2);
INSERT INTO test_table values(3);
COMMIT;
```

4. Select rows from the TPF.

The table `test_table` has three rows with values 1,2,3. The sum of these values is 6. The example generates 6 rows.

```
SELECT * from tpf_rg_1( TABLE( select val from test_table ) );
```

- a) To see how the `describe` affects the behavior, issue a **CREATE PROCEDURE** statement that has a different schema than the schema the TPF publishes in the `describe`:

```
CREATE OR REPLACE PROCEDURE tpf_rg_1( IN tab TABLE( num INT,
num2 INT ) )
RESULT( c1 INT )
EXTERNAL NAME 'tpf_rg_1@libv4apiex';
```

- b) Select rows from the TPF:

```
// This will return an error that the number of columns in
select list
does not match input table param schema
SELECT * from tpf_rg_1( TABLE( select val from test_table ) );
```

### tpf\_rg\_2

TPF sample `tpf_rg_2.cxx` builds on the sample in `tpf_rg_1.cxx` and has similar behavior. It produces rows of data based on an input parameter.

This sample provides an alternate implementation of the `_open_extfn` method in the `a_v4_extfn_func` descriptor. The behavior is the same as `tpf_rg_1` but the TPF uses `fetch_into` instead of `fetch_block` to read rows from the input table.

This code snippet from the `_open_extfn` method shows `fetch_into` retrieving rows from the input table:



```

static short UDF_CALLBACK tpf_rg_open(
    a_v4_extfn_table_context * tctx )
/*****/
{
...
...
...
    // This block of code will create a statically allocated row block
    // that can contain at most 1 row of data.

    a_sql_uint32          c1_data;
    a_sql_byte            c1_null    = 0x0;
    a_sql_uint32          c1_len     = 0;
    a_sql_byte            null_mask  = 0x1;
    a_sql_byte            null_value = 0x1;
    a_v4_extfn_column_data cd[1] =
    {
        { &c1_null,          // is_null
          null_mask,        // null_mask
          null_value,      // null_value
          &c1_data,         // data
          &c1_len,         // piece_len
          sizeof(c1_data), // max_piece_len
          NULL              // blob
        }
    };

    a_sql_uint32          r_status;

    a_v4_extfn_row        row =
    {
        { &r_status, &cd[0]
        };
    };

    a_v4_extfn_row_block  rb =
    {
        { 1, 0, &row
        };
    };

    // We are providing a row block structure that was statically
    // allocated to have a single row. This means that each call to
    // fetch_into will return at most 1 row.
    while( rs->fetch_into( rs, &rb ) ) {

// Only consider non-null rows. They way the column data has
// been defined allows us to treat c1_null as a boolean.
if( !c1_null ) {
    num_to_generate += c1_data;
}

    }

    ...
    ...
}

```

When using `fetch_into` to retrieve rows from an input table, the TPF manages the row block structure. In this example, a static row block structure is created that retrieves one row at a time. Alternatively, you can allocate a dynamic structure that simultaneously supports an arbitrary number of rows.

In the code snippet, the row block structure defined to store the value of the column from the input table in the variable `c1_data`. If a NULL row is encountered, the variable `c1_null` is set to 1 to indicate this.

### See also

- `_open_extfn` on page 316
- `_fetch_into_extfn` on page 317

### Running the Sample TPF in `tpf_rg_2`

The sample `tpf_rg_2` is included in a pre-compiled dynamic library called `libv4apiex` (extension is platform-dependent). Its implementation is in the `samples` directory in `tpf_rg_2.cxx`.

1. Issue a **CREATE PROCEDURE** statement to declare the TPF to the server.

```
CREATE OR REPLACE PROCEDURE tpf_rg_2( IN tab TABLE( num INT ) )
RESULT( c1 INT )
EXTERNAL NAME 'tpf_rg_2@libv4apiex';
```

2. Issue a **CREATE TABLE** statement to declare a table to use as input to the TPF.

```
CREATE TABLE test_table( val INT );
```

3. Insert rows into the table.

```
INSERT INTO test_table VALUES(1);
INSERT INTO test_table VALUES(2);
INSERT INTO test_table VALUES(3);
COMMIT;
```

4. Select rows from the TPF.

```
SELECT * FROM tpf_rg_2( TABLE( SELECT val FROM test_table ) );
```

The table `test_table` has three rows with values 1,2,3. The sum of these values is 6. The example generates 6 rows.

### Pass-Through TPF in `tpf_blob`

The TPF sample `tpf_blob.cxx` demonstrates advanced UDF LOB and CLOB handling. The example is available in the `samples` directory. This examination of `tpf_blob` illustrates concepts not already covered by the simpler examples in `tpf_rg_1` and `tpf_rg_2`; only the relevant portions are discussed.

A table UDF or TPF can not produce LOB or CLOB data. However, using a concept known as *pass-through*, LOB or CLOB data can be passed from an input table to an output table. In fact, any data type can be passed through from an input table to the result set. This allows a TPF to *filter* rows, meaning that the output is a subset of the input table rows.

The **CREATE PROCEDURE** statement supported by `tpf_blob` is:

```
CREATE PROCEDURE tpf_blob( IN tab TABLE( num INT, s [LONG] <VARCHAR |
BINARY >,
                        IN pattern char(1) )
RESULT SET ( num INT, s [LONG] <VARCHAR | BINARY > )
EXTERNAL NAME 'tpf_blob@libv4apiex'
```

The procedure supports multiple schemas. The data types for column **s** in the result set and input table can be one of `VARCHAR`, `BINARY`, `LONG VARCHAR`, or `LONG BINARY`.

### Dynamic Schema Support

The schema for the `tpf_blob` procedure is dynamic.

The data types for column **s** in the result set and input table can be one of `VARCHAR`, `BINARY`, `LONG VARCHAR`, or `LONG BINARY`. You accomplish this using the `describe_column_get` method of `a_v4_extfn_proc_context` to get the data type of the input table column. The implementation of the TPF is adjusted according to what the actual defined schema is. The interpretation of the *pattern* argument of the procedure differs depending on the data type of column **s**. For character data types, the argument is interpreted as a letter; for binary data types it is interpreted as a digit.

### See also

- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287
- *\*describe\_column\_get* on page 204

### Processing LOB and CLOB Columns in Input Tables

A `tpf_blob` example counting the number of occurrences of a pattern in each data row in the input table.

When the procedure is defined as having `LONG VARCHAR` or `LONG BINARY` for column **s**, the data must be processed using the `blob` API. This code snippet from the `fetch_into_extfn` method illustrates how a TPF can use the `blob` API to process LOB and CLOB data from an input table:

```
if( EXTFN_COL_IS_BLOB(cd, 1) ) {
    ret = state->rs->get_blob( state->rs, &cd[1], &blob );

    UDF_SQLERROR_RT( tctx->proc_context,
                    "Failed to get blob",
                    (ret == 1 && blob != NULL),
                    0 );

    if( state->data_type == DT_VARCHAR ||
        state->data_type == DT_LONGVARCHAR ) {

        num = ProcessBlob( tctx->proc_context, blob, state->pattern );
    } else {
        char i = (char)atoi( &(state->pattern) );
        num = ProcessBlob( tctx->proc_context, blob, i );
    }
}
```

## Table UDFs and TPFs

```
    }
    ret = blob->release( blob );
    UDF_SQLERROR_RT( tctx->proc_context,
                    "Failed to release blob",
                    (ret == 1),
                    0 );
    } else {
    if( state->data_type == DT_VARCHAR ||
        state->data_type == DT_LONGVARCHAR ) {
        num = CountNum( (char *)cd[1].data,
*(cd[1].piece_len),
state>pattern );
    } else {
        char i = (char)atoi( &(amp;state->pattern) );
        num = CountNum( (char *)cd[1].data, *(cd[1].piece_len), i );
    }
    }
}
```

For each of the rows in the input table, the TPF checks if it is a blob using the macro `EXTFN_COL_IS_BLOB`. If it is a blob, then the TPF uses the `get_blob` method of `a_v4_extfn_table_context` to create a blob object for the specified column. On success, the `get_blob` method provides the TPF with an instance of `a_v4_extfn_blob`, which allows the TPF to read in the blob data. Once the TPF is finished with the blob, it should call `release` on it.

The `ProcessBlob` method illustrates how a blob object processes the data:

```
static a_sql_uint64 ProcessBlob(
    a_v4_extfn_proc_context *ctx,
    a_v4_extfn_blob *blob,
    char pattern)
/*****/
{
    char buffer[BLOB_ISTREAM_BUFFER_LEN];
    size_t len = 0;
    short ret = 0;

    a_sql_uint64 num = 0;

    a_v4_extfn_blob_istream *is = NULL;

    ret = blob->open_istream( blob, &is );
    UDF_SQLERROR_RT( ctx,
                    "Failed to open blob istream",
                    (ret == 1 && is != NULL),
                    0 );

    for(;;) {
        len = is->get( is, buffer, BLOB_ISTREAM_BUFFER_LEN );
        if( len == 0 ) {
            break;
        }
        num += CountNum( buffer, len, pattern );
    }
}
```

```

ret = blob->close_istream( blob, is );
UDF_SQLERROR_RT( ctx,
    "Failed to close blob istream",
    (ret == 1),
    0 );
return num;
}

```

The `open_istream` method on the blob object creates an instance of `a_v4_extfn_blob_istream`, which can then be used to read a specified amount of the blob into a buffer using the `get` method.

### See also

- *Blob Input Stream (a\_v4\_extfn\_blob\_istream)* on page 199
- *Blob (a\_v4\_extfn\_blob)* on page 195
- *get\_blob* on page 314
- *fetch\_into* on page 308

### Passing Input Table Columns to the Result Set

A `tpf_blob` illustrating how a TPF can pass the rows from an input table to the result table, and how to use the `row_status` flag to indicate if a row is present.

This allows the TPF to filter out unwanted rows.

1. During the describe phase, ensure the TPF uses the `describe_column_set` method of `EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT` to inform the server that specific result-set rows are a subset of rows from an input table.

This code snippet from the `describe_extfn` method illustrates filtering:

```

else if( ctx->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
    // The output columns of this TPF are the same as the first
    // argument's input table columns. The following describe
    // informs the consumer of this fact.
    a_v4_extfn_col_subset_of_input colMap;

    for( short i = 1; i <= 2; i++ ) {
        colMap.source_table_parameter_arg_num = 1;
        colMap.source_column_number = i;

        desc_rc = ctx->describe_column_set( ctx,
            0, i,
            EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT,
            &colMap, sizeof(a_v4_extfn_col_subset_of_input) );

        UDF_CHECK_DESCRIBE( ctx, desc_rc );
    }
}

```

2. Pass the call to `fetch_into` for the input table the same rowblock structure that was passed into the `fetch_into_extfn` method. This ensures that the rowblock structure for the result set is the same as for the input tables.

### See also

- *EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Set)* on page 236
- *fetch\_into* on page 308
- *\_fetch\_into\_extfn* on page 317

### Running the Sample TPF in `tpf_blob.cxx`

The sample `tpf_blob` is included in a precompiled dynamic library called `libv4apiex` (extension is platform-dependent). Its implementation is in the `samples` directory in `tpf_blob.cxx`.

1. Declare the TPF to the server:

```
CREATE OR REPLACE PROCEDURE tpf_blob( IN tab TABLE( num INT,
                                                    s long
      varchar ),
                                     IN pattern char(1) )
RESULT( num INT, s long varchar )
EXTERNAL NAME 'tpf_blob@libv4apiex';
```

2. Declare a table to use as input to the TPF:

```
CREATE TABLE test_table( val INT, str LONG VARCHAR );
```

3. Insert rows into the table:

```
INSERT INTO test_table VALUES(1, 'aaaaaaaaaabbmbmbmbmb');
INSERT INTO test_table VALUES(2, 'aaaaaaaaaabbmbmbmbmb');
INSERT INTO test_table VALUES(3, 'aaaaaaaaaabbmbmbmbmb');
INSERT INTO test_table VALUES(4, 'aaaaaaaaaabbmbmbmbmb');
INSERT INTO test_table VALUES(5, 'aaaaaaaaaabbmbmbmbmb');
COMMIT;
```

4. Select rows from the TPF:

```
SELECT * FROM tpf_blob( TABLE( SELECT val,str FROM test_table ),
'a' );
```

The table `test_table` has three rows with an even number of `as`. Row 1 has 10, row 3 has 12, and row 5 has 14.

---

## SQL Reference for Table UDF and TPF Queries

SQL statement reference for queries referencing table UDFs and TPFs.

## **ALTER PROCEDURE Statement**

Replaces an existing procedure with a modified version. Include the entire modified procedure in the **ALTER PROCEDURE** statement, and reassign user permissions on the procedure.

### **Syntax**

Syntax 1

```
ALTER PROCEDURE [ owner.]procedure-name procedure-definition
```

Syntax 2

```
ALTER PROCEDURE [ owner.]procedure-name  
REPLICATE { ON | OFF }
```

Syntax 3

```
ALTER PROCEDURE [ owner.]procedure-name  
SET HIDDEN
```

Syntax 4

```
ALTER PROCEDURE [ owner.]procedure-name  
RECOMPILE
```

Syntax 5

```
ALTER PROCEDURE  
[ owner.]procedure-name ( [ parameter, ... ] )  
[ RESULT (result-column, ...)]  
EXTERNAL NAME 'external-call' [ LANGUAGE environment-name ] }
```

### **Parameters**

- **procedure-definition** – **CREATE PROCEDURE** syntax following the name
- **environment-name** – **JAVA** [ **DISALLOW** | **ALLOW SERVER SIDE REQUESTS** ]
- **external-call** – [*column-name:*]*function-name*@*library*, ...

### **Usage**

The **ALTER PROCEDURE** statement must include the entire new procedure. You can use **PROC** as a synonym for **PROCEDURE**.

The **ALTER PROCEDURE** statement is identical in syntax to the **CREATE PROCEDURE** statement.

- **Syntax 1** – The **ALTER PROCEDURE** statement is identical in syntax to the **CREATE PROCEDURE** statement except for the first word. Both Watcom and Transact-SQL dialect procedures can be altered through the use of **ALTER PROCEDURE**.

With **ALTER PROCEDURE**, existing permissions on the procedure are not changed. If you execute **DROP PROCEDURE** followed by **CREATE PROCEDURE**, execute permissions are reassigned.

- **Syntax 2** – If a procedure needs to be relocated to other sites using Sybase Replication Server, set **REPLICATE ON** for the procedure.
- **Syntax 3** – Use **SET HIDDEN** to obfuscate the definition of the associated procedure and cause it to become unreadable. The procedure can be unloaded and reloaded into other databases.

---

**Note:** This setting is irreversible. It is recommended that you retain the original procedure definition outside of the database.

---

If **SET HIDDEN** is used, debugging using the debugger does not show the procedure definition, and the definition is not available through procedure profiling.

You cannot combine Syntax 2 with Syntax 1.

- **Syntax 4** – Use the **RECOMPILE** syntax to recompile a stored procedure. When you recompile a procedure, the definition stored in the catalog is re-parsed and the syntax is verified. For procedures that generate a result set but do not include a **RESULT** clause, the database server attempts to determine the result set characteristics for the procedure and stores the information in the catalog. This can be useful if a table referenced by the procedure has been altered to add, remove, or rename columns since the procedure was created.

The procedure definition is not changed by recompiling. You can recompile procedures with definitions hidden with the **SET HIDDEN** clause, but their definitions remain hidden.

- **Syntax 5** – Use this syntax for Java UDFs.

**iq-environment-name: JAVA [ DISALLOW | ALLOW SERVER SIDE REQUESTS ]:**

**DISALLOW** is the default.

**ALLOW** indicates that server-side connections are allowed.

---

**Note:** Do not specify **ALLOW** unless necessary. A setting of **ALLOW** slows down certain types of Sybase IQ table joins.

Do not use UDFs with both **ALLOW SERVER SIDE REQUESTS** and **DISALLOW SERVER SIDE REQUESTS** in the same query.

---

When using the **ALTER PROCEDURE** statement for table UDFs, the same set of restrictions apply as for the **CREATE PROCEDURE Statement (External Procedures)**.

### **Standards**

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.



## Permissions

Must be the owner of the procedure or a DBA. Automatic commit.

## See also

- *Table UDF Restrictions* on page 99
- *CREATE PROCEDURE Statement (Table UDF)* on page 169

## CREATE PROCEDURE Statement (Table UDF)

Creates an interface to an external table user-defined function (table UDF). Users must be specifically licensed to use table UDFs.

You define table UDFs using the `a_v4_extfn` API. **CREATE PROCEDURE** statement reference information for external procedures that do not use the `a_v3_extfn` or `a_v4_extfn` APIs is located in a separate topic. **CREATE PROCEDURE** statement reference information for Java UDFs is located in a separate topic.

## Syntax

```

CREATE [ OR REPLACE ] PROCEDURE
[ owner. ] procedure-name ( [ parameter [ , ... ] ] )
| RESULT result-column [ , ... ] )
[ SQL SECURITY { INVOKER | DEFINER } ]
EXTERNAL NAME 'external-call'

```

## Parameters

- **parameter:** – [ **IN** ] *parameter-name data-type* [ **DEFAULT** *expression* ]  
| [ **IN** ] *parameter-name table-type*
- **table-type:** – **TABLE**( *column-name data-type* [ , ... ] )
- **external-call:** – [ *column-name:* ] *function-name@library*, ...

## Usage

The **CREATE PROCEDURE** statement creates a procedure in the database. Users with DBA authority can create procedures for other users by specifying an owner.

If a stored procedure returns a result set, it cannot also set output parameters or return a return value.

When referencing a temporary table from multiple procedures, a potential issue can arise if the temporary table definitions are inconsistent and statements referencing the table are cached. Use caution when referencing temporary tables within procedures.

## Table UDFs and TPFs

You can use the **CREATE PROCEDURE** statement to create external table UDFs implemented in a different programming language than SQL. However, be aware of the table UDF restrictions before creating external UDFs.

The data type for a scalar parameter, a result column, and a column of a TABLE parameter must be a valid SQL data type.

Parameter names must conform to the rules for other database identifiers such as column names. They must be a valid SQL data type.

TPFs support a mix scalar parameters and single table parameter. A TABLE parameter must define a schema for an input set of rows to be processed by the UDF. The definition of a TABLE parameter includes column names and column data types.

```
TABLE(c1 INT, c2 CHAR(20))
```

The above example defines a schema with the two columns c1 and c2 of types INT and CHAR(20). Each row processed by the UDF must be a tuple with two (2) values. Table parameters, unlike scalar parameters cannot be assigned a default value.

- **IN keyword** – Parameters can be prefixed with the keyword IN:
  - IN—The parameter is an object that provides a value for a scalar parameter or a set of values for a TABLE parameter to the UDF.

---

**Note:** TABLE parameters cannot be declared as INOUT or OUT. You can only have one table parameter (the position of which is not important).

---

- **OR REPLACE** – Specifying **OR REPLACE (CREATE OR REPLACE PROCEDURE)** creates a new procedure, or replaces an existing procedure with the same name. This clause changes the definition of the procedure, but preserves existing permissions. An error is returned if you attempt to replace a procedure that is already in use.
- **RESULT** – Declares the column names and their data types for the result set of the external UDF. The data types of the columns must be a valid SQL data type (e.g., a column in the result set cannot have TABLE as data type). The set of datums in the result implies the TABLE. External UDFs can only have one result set of type TABLE.

---

**Note:** TABLE is not an output value. A table UDF cannot have LONG VARBINARY or LONG VARCHAR data types in its result set, but a table parameterized function (TPF) can have large object (LOB) data in its result set.

A TPF cannot produce LOB data, but can have columns in the result set as LOB data types. However, the only way to get LOB data in the output is to pass a column from an input table to the output table. The describe attribute `EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT` allows this, as illustrated in the sample file `tpf_blob.cxx`.

---

For more information on returning result sets from procedures, see *System Administration Guide: Volume 2 > Using Procedures and Batches*.

- **SQL SECURITY** – Defines whether the procedure is executed as the INVOKER (the user who is calling the UDF), or as the DEFINER (the user who owns the UDF). The default is DEFINER.

When SQL SECURITY INVOKER is specified, more memory is used because annotation must be done for each user that calls the procedure. Also, when SQL SECURITY INVOKER is specified, name resolution is done as the invoker as well. Therefore, care should be taken to qualify all object names (tables, procedures, and so on) with their appropriate owner. For example, suppose user1 creates this procedure:

```
CREATE PROCEDURE user1.myProcedure(
    RESULT( columnA INT )
    SQL SECURITY INVOKER
    BEGIN
        SELECT columnA FROM table1;
    END;
```

If user2 attempts to run this procedure and a table user2.table1 does not exist, a table lookup error results. Additionally, if a user2.table1 does exist, that table is used instead of the intended user1.table1. To prevent this situation, qualify the table reference in the statement (user1.table1, instead of just table1).

- **EXTERNAL NAME** – EXTERNAL NAME *'external-call': external-call: [operating-system: ]function-name@library; ...*

An external UDF must have **EXTERNAL NAME** clause which defines an interface to a function written in a programming language such as C. The function is loaded by the database server into its address space.

The library name can include the file extension, which is typically .dll on Windows and .so on UNIX. In the absence of the extension, the software appends the platform-specific default file extension for libraries. This is a formal example.

```
CREATE PROCEDURE mystring( IN instr CHAR(255),
    IN input_table TABLE(A INT) )
    RESULT (CHAR(255))
    EXTERNAL NAME
    'mystring@mylib.dll;Unix:mystring@mylib.so'
```

A simpler way to write the preceding EXTERNAL NAME clause, using platform-specific defaults, is as follows:

```
CREATE PROCEDURE mystring( IN instr CHAR(255),
    IN input_table TABLE(A INT) )
    RESULT (CHAR(255))
    EXTERNAL NAME 'mystring@mylib'
```

## **Standards**

- SQL—ISO/ANSI SQL compliant.
- Sybase—The Transact-SQL **CREATE PROCEDURE** statement is different.

## Table UDFs and TPFs

- SQLJ—The syntax extensions for Java result sets are as specified in the proposed SQLJ1 standard.

### Permissions

Must have RESOURCE authority, unless creating a temporary procedure. Users with DBA authority can create UDFs for other users by specifying an owner. Users must have DBA authority to create an external UDFs or to create an external UDF for another user

### **See also**

- *Sample Files* on page 99

## CREATE FUNCTION Statement

Creates a new function in the database.

### Syntax

Syntax 1

```
CREATE [ OR REPLACE ] [ TEMPORARY ] FUNCTION [ owner.]function-name
( [ parameter, ... ] )
RETURNS data-type routine-characteristics
[ SQL SECURITY { INVOKER | DEFINER } ]
{ compound-statement
| AS tsql-compound-statement
| external-name }
```

Syntax 2

```
CREATE FUNCTION [ owner.]function-name ( [ parameter, ... ] )
RETURNS data-type
URL url-string
[ HEADER header-string ]
[ SOAPHEADER soap-header-string ]
[ TYPE { 'HTTP[:{ GET | POST } ]' | 'SOAP[:{ RPC | DOC } ]' } ]
[ NAMESPACE namespace-string ]
[ CERTIFICATE certificate-string ]
[ CLIENTPORT clientport-string ]
[ PROXY proxy-string ]
```

### Parameters

- **url-string:** – ' { HTTP | HTTPS | HTTPS\_FIPS } ://[user:password@]hostname[:port][/path] '
- **parameter:** – IN *parameter-name* *data-type* [ DEFAULT *expression* ]
- **routine-characteristics:** – ON EXCEPTION RESUME | [ NOT ] DETERMINISTIC
- **tsql-compound-statement:** – *sql-statement sql-statement ...*
- **external-name:** – EXTERNAL NAME *library-call* | EXTERNAL NAME *java-call*  
LANGUAGE JAVA

- **library-call:** – '[ *operating-system*:]*function-name*@*library*, ...'
- **operating-system:** – **UNIX**
- **java-call:** – '[ *package-name*.]*class-name*.*method-name* *method-signature*'
- **method-signature:** – ( [ *field-descriptor*, ... ] ) *return-descriptor*
- **field-descriptor and return-descriptor:** – **Z | B | S | I | J | F | D | C | V** | [ *descriptor* | **L** *class-name*;

## Examples

- **Example 1** – Concatenate a firstname string and a lastname string:

```
CREATE FUNCTION fullname (
    firstname CHAR(30),
    lastname CHAR(30) )
RETURNS CHAR(61)
BEGIN
    DECLARE name CHAR(61);
    SET name = firstname || ' ' || lastname;
    RETURN (name);
END
```

This examples illustrate the use of the **fullname** function.

- Return a full name from two supplied strings:

```
SELECT fullname ('joe', 'smith')
```

<b>fullname('joe', 'smith')</b>
joe smith

- List the names of all employees:

```
SELECT fullname (givenname, surname)
FROM Employees
```

<b>fullname (givenname, surname)</b>
Fran Whitney
Matthew Cobb
Philip Chin
Julie Jordan
Robert Breault
...

- **Example 2** – Use Transact-SQL syntax:

```
CREATE FUNCTION DoubleIt ( @Input INT )
RETURNS INT
AS
```

## Table UDFs and TPFs

```
DECLARE @Result INT
SELECT @Result = @Input * 2
RETURN @Result
```

The statement `SELECT DoubleIt( 5 )` returns a value of 10.

- **Example 3** – Create an external function written in Java:

```
CREATE FUNCTION dba.encrypt( IN name char(254) )
RETURNS VARCHAR
EXTERNAL NAME
'Scramble.encrypt (Ljava/lang/String;)Ljava/lang/String;'
LANGUAGE JAVA
```

### Usage

The **CREATE FUNCTION** statement creates a user-defined function in the database. A function can be created for another user by specifying an owner name. Subject to permissions, a user-defined function can be used in exactly the same way as other non-aggregate functions.

**CREATE FUNCTION**—Parameter names must conform to the rules for database identifiers. They must have a valid SQL data type and be prefixed by the keyword **IN**, signifying that the argument is an expression that provides a value to the function.

Specifying **CREATE OR REPLACE FUNCTION** creates a new function, or replaces an existing function with the same name. When a function is replaced, the definition of the function is changed but the existing permissions are preserved.

You cannot use the **OR REPLACE** clause with temporary functions.

When functions are executed, not all parameters need to be specified. If a default value is provided in the **CREATE FUNCTION** statement, missing parameters are assigned the default values. If an argument is not provided by the caller and no default is set, an error is given.

Specifying **TEMPORARY (CREATE TEMPORARY FUNCTION)** means that the function is visible only by the connection that created it, and that it is automatically dropped when the connection is dropped. Temporary functions can also be explicitly dropped. You cannot perform **ALTER**, **GRANT**, or **REVOKE** operations on them, and unlike other functions, temporary functions are not recorded in the catalog or transaction log.

Temporary functions execute with the permissions of their creator (current user), and can only be owned by their creator. Therefore, do not specify owner when creating a temporary function.

Temporary functions can be created and dropped when connected to a read-only database.

**SQL SECURITY**—Defines whether the function is executed as the **INVOKER**, the user who is calling the function, or as the **DEFINER**, the user who owns the function. The default is **DEFINER**.

When **SQL SECURITY INVOKER** is specified, more memory is used because annotation must be done for each user that calls the procedure. Also, when **SQL SECURITY INVOKER** is

specified, name resolution is done as the invoker as well. Therefore, take care to qualify all object names (tables, procedures, and so on) with their appropriate owner.

**compound-statement**—A set of SQL statements bracketed by **BEGIN** and **END**, and separated by semicolons. See *BEGIN ... END Statement*.

**tsql-compound-statement**—A batch of Transact-SQL statements. See *Reference: Building Blocks, Tables, and Procedures > Compatibility with Other Sybase Databases > Transact-SQL Procedure Language Overview > Transact-SQL Batch Overview* and *CREATE PROCEDURE Statement [T-SQL]*.

**EXTERNAL NAME**—A function using the **EXTERNAL NAME** clause is a wrapper around a call to a function in an external library. A function using **EXTERNAL NAME** can have no other clauses following the **RETURNS** clause. The library name may include the file extension, which is typically `.dll` on Windows and `.so` on UNIX. In the absence of the extension, the software appends the platform-specific default file extension for libraries.

The **EXTERNAL NAME** clause is not supported for temporary functions. See *SQL Anywhere Server – Programming > SQL Anywhere external call interface*.

---

**Note:** This reference points to SQL Anywhere documentation.

---

**EXTERNAL NAME LANGUAGE JAVA**—A function that uses **EXTERNAL NAME** with a **LANGUAGE JAVA** clause is a wrapper around a Java method. For information on calling Java procedures, see *CREATE PROCEDURE Statement*.

**ON EXCEPTION RESUME**—Uses Transact-SQL-like error handling. See *CREATE PROCEDURE Statement*.

**NOT DETERMINISTIC**—A function specified as **NOT DETERMINISTIC** is re-evaluated each time it is called in a query. The results of functions not specified in this manner may be cached for better performance, and re-used each time the function is called with the same parameters during query evaluation.

Functions that have side effects, such as modifying the underlying data, should be declared as **NOT DETERMINISTIC**. For example, a function that generates primary key values and is used in an **INSERT ... SELECT** statement should be declared **NOT DETERMINISTIC**:

```
CREATE FUNCTION keygen( increment INTEGER )
RETURNS INTEGER
NOT DETERMINISTIC
BEGIN
    DECLARE keyval INTEGER;
    UPDATE counter SET x = x + increment;
    SELECT counter.x INTO keyval FROM counter;
    RETURN keyval
END
INSERT INTO new_table
SELECT keygen(1), ...
FROM old_table
```

Functions may be declared as **DETERMINISTIC** if they always return the same value for given input parameters.

All user-defined functions are treated as deterministic unless they are declared **NOT DETERMINISTIC**. Deterministic functions return a consistent result for the same parameters and are free of side effects. That is, the database server assumes that two successive calls to the same function with the same parameters will return the same result without unwanted side-effects on the semantics of the query.

---

**Note:** User-defined functions in the CLR, ESQL and ODBC, Perl, and PHP external environments are processed by the Catalog Store (SQL Anywhere). They do not take advantage of the performance features of Sybase IQ. User-defined functions in Java are processed by Sybase IQ and offer increased performance.

In certain cases, differences in semantics between SQL Anywhere and Sybase IQ can produce different results for a query if the query is issued in a user-defined function. For example, Sybase IQ treats the CHAR and VARCHAR data types as distinct and different, while SQL Anywhere treats CHAR data as if it were VARCHAR.

---

To modify a user-defined function, or to hide the contents of a function by scrambling its definition, use the **ALTER FUNCTION** statement.

**URL**—For use only when defining an **HTTP** or **SOAP** web services client function. Specifies the URL of the web service. The optional user name and password parameters provide a means of supplying the credentials needed for **HTTP** basic authentication. **HTTP** basic authentication base-64 encodes the user and password information and passes it in the “Authentication” header of the **HTTP** request.

For web service client functions, the return type of **SOAP** and **HTTP** functions must one of the character data types, such as VARCHAR. The value returned is the body of the **HTTP** response. No **HTTP** header information is included. If more information is required, such as status information, use a procedure instead of a function.

Parameter values are passed as part of the request. The syntax used depends on the type of request. For **HTTP:GET**, the parameters are passed as part of the URL; for **HTTP:POST** requests, the values are placed in the body of the request. Parameters to **SOAP** requests are always bundled in the request body.

**HEADER**—When creating **HTTP** web service client functions, use this clause to add or modify **HTTP** request header entries. Only printable ASCII characters can be specified for **HTTP** headers, and they are case-insensitive. For more information about how to use this clause, see the **HEADER** clause of the *CREATE PROCEDURE Statement*.

For more information about using HTTP headers, see *SQL Anywhere Server – Programming > HTTP web services » Accessing web services using web clients » Developing web client applications » Web client function and procedure requirements and recommendations > HTTP request header management*.



---

**Note:** This reference points to SQL Anywhere documentation.

---

**SOAPHEADER**—When declaring a **SOAP** Web service as a function, use this clause to specify one or more **SOAP** request header entries. A **SOAP** header can be declared as a static constant, or can be dynamically set using the parameter substitution mechanism (declaring **IN**, **OUT**, or **INOUT** parameters for hd1, hd2, and so on). A web service function can define one or more **IN** mode substitution parameters, but cannot define an **INOUT** or **OUT** substitution parameter. For more information about how to use this clause, see the **SOAPHEADER** clause in *SQL Anywhere Server – SQL Reference > SQL statements > SQL statements > CREATE PROCEDURE statement (web clients)*.

---

**Note:** This reference points to SQL Anywhere documentation.

---

**TYPE**—Specifies the format used when making the web service request. If **SOAP** is specified or no type clause is included, the default type **SOAP:RPC** is used. **HTTP** implies **HTTP:POST**. Since **SOAP** requests are always sent as XML documents, **HTTP:POST** is always used to send **SOAP** requests.

**NAMESPACE**—Applies to **SOAP** client functions only and identifies the method namespace usually required for both **SOAP:RPC** and **SOAP:DOC** requests. The **SOAP** server handling the request uses this namespace to interpret the names of the entities in the **SOAP** request message body. The namespace can be obtained from the WSDL description of the **SOAP** service available from the web service server. The default value is the procedure's URL, up to but not including the optional path component.

**CERTIFICATE**—To make a secure (HTTPS) request, a client must have access to the certificate used by the HTTPS server. The necessary information is specified in a string of semicolon-separated key/value pairs. The certificate can be placed in a file and the name of the file provided using the file key, or the whole certificate can be placed in a string, but not both. These keys are available:

Key	Abbreviation	Description
file		File name of certificate
certificate	cert	The certificate
company	co	Company specified in the certificate
unit		Company unit specified in the certificate
name		Common name specified in the certificate

Certificates are required only for requests that are either directed to an HTTPS server or can be redirected from an insecure to a secure server.

**CLIENTPORT**—Identifies the port number on which the HTTP client procedure communicates using TCP/IP. It is provided for and recommended only for connections across firewalls, as firewalls filter according to the TCP/UDP port. You can specify a single port

## Table UDFs and TPFs

number, ranges of port numbers, or a combination of both; for example, CLIENTPORT '85,90-97'.

See *System Administration Guide: Volume 1 > Connection and Communication Parameters > Network Communications Parameters > ClientPort Communication Parameter [CPort]*.

PROXY—Specifies the URI of a proxy server. For use when the client must access the network through a proxy. Indicates that the procedure is to connect to the proxy server and send the request to the web service through it.

### Side Effects

- Automatic commit

### **Standards**

- SQL—ISO/ANSI SQL compliant.
- Sybase—Not supported by Adaptive Server Enterprise.

### **Permissions**

Must have RESOURCE authority.

External functions, including Java functions, must have DBA authority.

## **DEFAULT\_TABLE\_UDF\_ROW\_COUNT Option**

Enables you to override the default estimate of the number of rows to return from a table UDF (either a C, C++, or Java table UDF).

### *Allowed Values*

0 to 4294967295

### *Default*

200000

### *Scope*

DBA permissions are required to set this option. Temporary only for the PUBLIC group.

### *Description*

A table UDF can use the **DEFAULT\_TABLE\_UDF\_ROW\_COUNT** option to give the query processor an estimate for the number of rows that a table UDF will return. This is the only way a Java table UDF can convey this information. However, for a C or C++ table UDF, the UDF developer should consider publishing this information in the describe phase using the `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS` describe parameter to publish the number of rows it expects to return. The value of `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS` always overrides the value of the **DEFAULT\_PROXY\_TABLE\_UDF\_ROW\_COUNT** option.

**See also**

- *Query Processing States* on page 121

**TABLE\_UDF\_ROW\_BLOCK\_SIZE\_KB Option**

Controls the size, in kilobytes, for server-allocated row blocks. Row blocks are used by Table UDFs and TPFs.

*Allowed Values*

0 to 4294967295

*Default*

128

*Scope*

DBA permissions are required to set this option. Temporary only for the PUBLIC group.

*Description*

Specifies the row block size, in kilobytes, to fetch from the server.

The server allocates row blocks when you use `fetch_into` to fetch rows from a table UDF, and when you use `fetch_block` to fetch rows from a TPF input table.

The row block contains as many rows as will fit into the specified size. If you specify a row block size smaller than the size required for a single row, the server allocates the size of one row.

**FROM Clause**

Specifies the database tables or views involved in a **SELECT** statement.

**Syntax**

```
... FROM table-expression [, ...]
```

```
table-expression :
  table-name
  | view-name
  | procedure-name
  | common-table-expression
  | (subquery) [[ AS] derived-table-name [column_name, ...]]
  | derived-table
  | join-expression
  | ( table-expression, ... )
  | openstring-expression
  | apply-expression
  | contains-expression
  | dml-derived-table
```

```
table-name :
  [ userid.]table-name ]
```

## Table UDFs and TPFs

```
[ [ AS ] correlation-name ]  
[ FORCE INDEX ( index-name ) ]
```

```
view-name :  
[ userid.view-name [ [ AS ] correlation-name ]
```

```
procedure-name :  
[ owner, ] procedure-name ([ parameter, ...])  
[ WITH( column-name datatype, )]  
[ [ AS] correlation-name ]
```

```
parameter :  
scalar-expression | table-parameter
```

```
table-parameter :  
TABLE(select-statement) [ OVER (table-parameter-over)]
```

```
table-parameter-over :  
[ PARTITION BY {ANY| NONE| table-expression } ]  
[ ORDER BY { expression | integer } [ ASC | DESC ] [, ... ]
```

```
derived-table :  
( select-statement )  
[ AS ] correlation-name [ ( column-name, ... ) ]
```

```
join-expression :  
table-expression join-operator table-expression  
[ ON join-condition ]
```

```
join-operator :  
[ KEY | NATURAL ] [ join-type ] JOIN  
| CROSS JOIN
```

```
join-type :  
INNER  
| LEFT [ OUTER ]  
| RIGHT [ OUTER ]  
| FULL [ OUTER ]
```

```
openstring-expression :  
OPENSTRING ( { FILE | VALUE } string-expression )  
WITH ( rowset-schema )  
[ OPTION ( scan-option ... ) ]  
[ AS ] correlation-name
```

```
apply-expression :  
table-expression { CROSS | OUTER } APPLY table-expression
```

```
contains-expression :  
{ table-name | view-name } CONTAINS ( column-name [,...], contains-  
query ) [ [ AS ] score-correlation-name ]
```

```
rowset-schema :  
column-schema-list  
| TABLE [owner.]table-name [ ( column-list ) ]
```

```
column-schema-list :
  { column-name user-or-base-type | filler( ) } [ , ... ]
```

```
column-list :
  { column-name | filler( ) } [ , ... ]
```

```
scan-option :
  BYTE ORDER MARK { ON | OFF }
  | COMMENTS INTRODUCED BY comment-prefix
  | DELIMITED BY string
  | ENCODING encoding
  | ESCAPE CHARACTER character
  | ESCAPES { ON | OFF }
  | FORMAT { TEXT | BCP }
  | HEXADECIMAL { ON | OFF }
  | QUOTE string
  | QUOTES { ON | OFF }
  | ROW DELIMITED BY string
  | SKIP integer
  | STRIP { ON | OFF | LTRIM | RTRIM | BOTH }
```

```
contains-query :
  string
```

```
dml-derived-table :
  ( dml-statement ) REFERENCING ( [ table-version-names | NONE ] )
```

```
dml-statement :
  insert-statement
  delete-statement
  update-statement
  merge-statement
```

```
table-version-names :
  OLD [ AS ] correlation-name [ FINAL [ AS ] correlation-name ]
  | FINAL [ AS ] correlation-name
```

## Examples

- **Example 1** – These are valid **FROM** clauses:

```
...
FROM Employees
...
...
FROM Employees NATURAL JOIN Departments
...
...
FROM Customers
KEY JOIN SalesOrders
KEY JOIN SalesOrderItems
KEY JOIN Products
...
```

- **Example 2** – This query illustrates how to use derived tables in a query:

```
SELECT Surname, GivenName, number_of_orders
FROM Customers JOIN
```

```
( SELECT CustomerID, count(*)
  FROM SalesOrders
  GROUP BY CustomerID )
AS sales_order_counts ( CustomerID,
                        number_of_orders )
ON ( Customers.ID = sales_order_counts.cust_id )
WHERE number_of_orders > 3
```

### Usage

The **SELECT** statement requires a table list to specify which tables are used by the statement.

**Note:** Although this description refers to tables, it also applies to views unless otherwise noted.

The **FROM** table list creates a result set consisting of all the columns from all the tables specified. Initially, all combinations of rows in the component tables are in the result set, and the number of combinations is usually reduced by join conditions and/or **WHERE** conditions.

- **SCALAR** – *scalar-parameter* are any objects of a valid SQL datatype.
- **TABLE** – A **TABLE** parameter can be specified using a table, view or common *table-expression* name which are treated as new instance of this object if the object is also used outside the **TABLE** parameter.

This query illustrates a valid **FROM** clause where the two references to the same table T are treated as two different instances of the same table T.

```
SELECT * FROM T, my_proc(TABLE(SELECT T.Z, T.X FROM T)
OVER(PARTITION BY T.Z));
```

Table Parameterized Function (TPF) Example—This query illustrates a valid **FROM** clause.

```
SELECT * FROM R, SELECT * FROM my_udf(1);
SELECT * FROM my_tpf(1, TABLE(SELECT c1, c2 FROM t))
(my_proc(R.X, TABLE T OVER PARTITION BY T.X)) AS XX;
```

If a subquery is used to define the **TABLE** parameter, then the following restrictions must hold:

- The **TABLE** parameter must be of type **IN**.
- **PARTITION BY** or **ORDER BY** clauses must refer to the columns of the derived table and outer references. An expression in the *expression-list* can be an integer *K* which refers to the *K*th column of the **TABLE** input parameter.

**Note:** A Table UDF can only be referenced in a **FROM** clause of a SQL statement.

- **PARTITION BY** – The **PARTITION BY** clause logically specifies how the invocation of the function will be performed by the execution engine. The execution engine must invoke the function for each partition and the function must process a whole partition in each invocation.

**PARTITION BY** also specifies how the input data must be partitioned such that each invocation of the function will process exactly one partition of data. The function must be

invoked the number of times equal to the number of partitions. For TPF, the parallelism characteristics are established through dynamic negotiation between the server and the UDF at the runtime. If the TPF can be executed in parallel, for  $N$  input partitions, the function can be instantiated  $M$  times, with  $M \leq N$ . Each instantiation of the function can be invoked more than once, each invocation consuming exactly one partition.

---

**Note:** The execution engine can invoke the function in any order of the partitions and the function is assumed to return the same result sets regardless of the partitions order. Partitions cannot be split among two invocations of the function.

---

You can specify only one **TABLE** input parameter for **PARTITION BY** *expression-list* or **PARTITION BY ANY** clause. For all other **TABLE** input parameters you must specify, explicit or implicit **PARTITION BY NONE** clause.

- **ORDER BY** – The **ORDER BY** clause specifies that the input data in each partition is expected to be sorted by *expression-list* by the execution engine. The UDF expects each partition to have this physical property. If only one partition exists, the whole input data is ordered based on the **ORDER BY** specification. **ORDER BY** clause can be specified for any of the **TABLE** input parameters with **PARTITION BY NONE** or without **PARTITION BY** clause.

For information on the *contains-expression* used in the **FROM** clause for full text searches, see *Unstructured Data Analytics in Sybase IQ*.

- **Joins** – The *join-type* keywords are:

**Table 5. FROM clause join-type keywords**

join-type keyword	Description
CROSS JOIN	Returns the Cartesian product (cross product) of the two source tables
NATURAL JOIN	Compares for equality all corresponding columns with the same names in two tables (a special case equijoin; columns are of same length and data type)
KEY JOIN	Restricts foreign-key values in the first table to be equal to the primary-key values in the second table
INNER JOIN	Discards all rows from the result table that do not have corresponding rows in both tables
LEFT OUTER JOIN	Preserves unmatched rows from the left table, but discards unmatched rows from the right table
RIGHT OUTER JOIN	Preserves unmatched rows from the right table, but discards unmatched rows from the left table
FULL OUTER JOIN	Retains unmatched rows from both the left and the right tables

Do not mix comma-style joins and keyword-style joins in the **FROM** clause. The same query can be written two ways, each using *one* of the join styles. The ANSI syntax keyword style join is preferable.

This query uses a comma-style join:

```
SELECT *
  FROM Products pr, SalesOrders so, SalesOrderItems si
 WHERE pr.ProductID = so.ProductID
      AND pr.ProductID = si.ProductID;
```

The same query can use the preferable keyword-style join:

```
SELECT *
  FROM Products pr INNER JOIN SalesOrders so
      ON (pr.ProductID = so.ProductID)
 INNER JOIN SalesOrderItems si
      ON (pr.ProductID = si.ProductID);
```

The **ON** clause filters the data of inner, left, right, and full joins. Cross joins do not have an **ON** clause. In an inner join, the **ON** clause is equivalent to a **WHERE** clause. In outer joins, however, the **ON** and **WHERE** clauses are different. The **ON** clause in an outer join filters the rows of a cross product and then includes in the result the unmatched rows extended with nulls. The **WHERE** clause then eliminates rows from both the matched and unmatched rows produced by the outer join. You must take care to ensure that unmatched rows you want are not eliminated by the predicates in the **WHERE** clause.

You cannot use subqueries inside an outer join **ON** clause.

For information on writing Transact-SQL compatible joins, see *Reference: Building Blocks, Tables, and Procedures > Compatibility with Other Sybase Databases*.

Tables owned by a different user can be qualified by specifying the *userid*. Tables owned by groups to which the current user belongs are found by default without specifying the user ID.

The correlation name is used to give a temporary name to the table for this SQL statement only. This is useful when referencing columns that must be qualified by a table name but the table name is long and cumbersome to type. The correlation name is also necessary to distinguish between table instances when referencing the same table more than once in the same query. If no correlation name is specified, then the table name is used as the correlation name for the current statement.

If the same correlation name is used twice for the same table in a table expression, that table is treated as if it were only listed once. For example, in:

```
SELECT *
  FROM SalesOrders
 KEY JOIN SalesOrderItems,
 SalesOrders
 KEY JOIN Employees
```



The two instances of the `SalesOrders` table are treated as one instance that is equivalent to:

```
SELECT *
FROM SalesOrderItems
KEY JOIN SalesOrders
KEY JOIN Employees
```

By contrast, the following is treated as two instances of the `Person` table, with different correlation names `HUSBAND` and `WIFE`.

```
SELECT *
FROM Person HUSBAND, Person WIFE
```

You can supply a **SELECT** statement instead of one or more tables or views in the **FROM** clause, letting you use groups on groups, or joins with groups, without creating a view. This use of **SELECT** statements is called derived tables.

Join columns require like data types for optimal performance.

- **Performance Considerations** – Depending on the query, Sybase IQ allows between 16 and 64 tables in the **FROM** clause with the optimizer turned on; however, performance might suffer if you have more than 16 to 18 tables in the **FROM** clause in very complex queries.

---

**Note:** If you omit the **FROM** clause, or if all tables in the query are in the `SYSTEM` dbspace, the query is processed by SQL Anywhere instead of Sybase IQ and might behave differently, especially with respect to syntactic and semantic restrictions and the effects of option settings.

If you have a query that does not require a **FROM** clause, you can force the query to be processed by Sybase IQ by adding the clause **FROM iq\_dummy**, where `iq_dummy` is a one-row, one-column table that you create in your database.

---

### Standards

- SQL—ISO/ANSI SQL compliant.
- Sybase—The **JOIN** clause is not supported in some versions of Adaptive Server Enterprise. Instead, you must use the **WHERE** clause to build joins.

### Permissions

Must be connected to the database.

## SELECT Statement

Retrieves information from the database.

### Syntax

```
SELECT [ ALL | DISTINCT ] [ FIRST | TOP number-of-rows ] select-list
... [ INTO { host-variable-list | variable-list | table-name } ]
```

## Table UDFs and TPFs

```
... [ INTO LOCAL TEMPORARY TABLE { table-name } ]
... [ FROM table-list ]
... [ WHERE search-condition ]
... [ GROUP BY [ expression [, ...]
      | ROLLUP ( expression [, ...] )
      | CUBE ( expression [, ...] ) ] ]
... [ HAVING search-condition ]
... [ ORDER BY { expression | integer } [ ASC | DESC ] [, ...] ]
... [ row-limitation-option ]
```

### Parameters

- **select-list:** –

```
{ column-name
| expression [ [ AS ] alias-name ]
| * }
```

- **row-limitation-option:** –

```
LIMIT { [ offset-expression, ] limit-expression
| limit-expression OFFSET offset-expression }
```

- **limit-expression:** – *simple-expression*
- **offset-expression:** – *simple-expression*
- **simple-expression:** –

```
integer
| variable
| ( simple-expression )
| ( simple-expression { + | - | * } simple-expression )
```

### Examples

- **Example 1** – List all tables and views in the system catalog:

```
SELECT tname
FROM SYS.SYSCATALOG
WHERE tname LIKE 'SYS%' ;
```

- **Example 2** – List all customers and the total value of their orders:

```
SELECT CompanyName,
       CAST( sum(SalesOrderItems.Quantity *
                Products.UnitPrice) AS INTEGER) VALUE
FROM Customers
   LEFT OUTER JOIN SalesOrders
   LEFT OUTER JOIN SalesOrderItems
   LEFT OUTER JOIN Products
GROUP BY CompanyName
ORDER BY VALUE DESC
```

- **Example 3** – List the number of employees:

```
SELECT count(*)
FROM Employees;
```

- **Example 4** – An Embedded SQL SELECT statement:

```
SELECT count(*) INTO :size FROM Employees;
```

- **Example 5** – List the total sales by year, model, and color:

```
SELECT year, model, color, sum(sales)
FROM sales_tab
GROUP BY ROLLUP (year, model, color);
```

- **Example 6** – Select all items with a certain discount into a temporary table:

```
SELECT * INTO #TableTemp FROM lineitem
WHERE l_discount < 0.5
```

- **Example 7** – Return information about the employee that appears first when employees are sorted by last name:

```
SELECT FIRST *
FROM Employees
ORDER BY Surname;
```

- **Example 8** – Return the first five employees when their names are sorted by last name:

```
SELECT TOP 5 *
FROM Employees
ORDER BY Surname;
```

```
SELECT *
FROM Employees
ORDER BY Surname
LIMIT 5;
```

- **Example 9** – List the fifth and sixth employees sorted in descending order by last name:

```
SELECT *
FROM Employees
ORDER BY Surname DESC
LIMIT 4,2;
```

- **Example 10** – Return the first five employees when their names are sorted by last name:

```
CREATE OR REPLACE VARIABLE atop INT = 10;
```

```
SELECT TOP (atop -5) *
FROM Employees
ORDER BY Surname;
```

```
SELECT *
FROM Employees
ORDER BY Surname
LIMIT (atop-5);
```

- **Example 11** – List the fifth and sixth employees sorted in descending order by last name:

```
CREATE OR REPLACE VARIABLE atop INT = 10;
```

```
SELECT *
FROM Employees
ORDER BY Surname DESC
LIMIT (atop - 8) OFFSET (atop -2 -3 -1);
```

**Usage**

You can use a **SELECT** statement in Interactive SQL to browse data in the database or to export data from the database to an external file.

You can also use a **SELECT** statement in procedures or in Embedded SQL. The **SELECT** statement with an **INTO** clause is used for retrieving results from the database when the **SELECT** statement returns only one row. (Tables created with **SELECT INTO** do not inherit **IDENTITY/AUTOINCREMENT** tables.) For multiple-row queries, you must use cursors. When you select more than one column and do not use *#table*, **SELECT INTO** creates a permanent base table. **SELECT INTO #table** always creates a temporary table regardless of the number of columns. **SELECT INTO** table with a single column selects into a host variable.

---

**Note:** When writing scripts and stored procedures that **SELECT INTO** a temporary table, wrap any select list item that is not a base column in a **CAST** expression. This guarantees that the column data type of the temporary table is the required data type.

---

Tables with the same name but different owners require aliases. A query without aliases returns incorrect results:

```
SELECT * FROM user1.t1
WHERE NOT EXISTS
(SELECT *
FROM user2.t1
WHERE user2.t1.col1 = user1.t.col1);
```

For correct results, use an alias for each table:

```
SELECT * FROM user1.t1 U1
WHERE NOT EXISTS
(SELECT *
FROM user2.t1 U2
WHERE U2.col1 = U1.col1);
```

The **INTO** clause with a *variable-list* is used only in procedures.

In **SELECT** statements, a stored procedure call can appear anywhere a base table or view is allowed. Note that CIS functional compensation performance considerations apply. For example, a **SELECT** statement can also return a result set from a procedure. For syntax and an example, see *SQL Anywhere Server – SQL Reference > SQL statements > SQL statements > FROM clause*.

---

**Note:** This reference points to SQL Anywhere documentation.

---

See *System Administration Guide: Volume 2 > Using Procedures and Batches > Introduction to Procedures > Returning Procedure Results in Result Sets* for a restriction that affects selecting from temporary tables within stored procedures.

**ALL** or **DISTINCT**—If neither is specified, all rows that satisfy the clauses of the **SELECT** statement are retrieved. If **DISTINCT** is specified, duplicate output rows are eliminated. This is

called the projection of the result of the statement. In many cases, statements take significantly longer to execute when **DISTINCT** is specified, so reserve the use of **DISTINCT** for cases where it is necessary.

If **DISTINCT** is used, the statement cannot contain an aggregate function with a **DISTINCT** parameter.

**FIRST** or **TOP** number-of-rows—Specifies the number of rows returned from a query. **FIRST** returns the first row selected from the query. **TOP** returns the specified number of rows from the query where *number-of-rows* is in the range 1 – 2147483647 and can be an integer constant or integer variable.

**FIRST** and **TOP** are used primarily with the **ORDER BY** clause. If you use these keywords without an **ORDER BY** clause, the result might vary from run to run of the same query, as the optimizer might choose a different query plan.

**FIRST** and **TOP** are permitted only in the top-level **SELECT** of a query, so they cannot be used in derived tables or view definitions. Using **FIRST** or **TOP** in a view definition might result in the keyword being ignored when a query is run on the view.

Using **FIRST** is the same as setting the `ROW_COUNT` database option to 1. Using **TOP** is the same as setting the `ROW_COUNT` option to the same number of rows. If both **TOP** and `ROW_COUNT` are set, then the value of **TOP** takes precedence.

The `ROW_COUNT` option could produce inconsistent results when used in a query involving global variables, system functions or proxy tables. See *ROW\_COUNT Option* for details.

*select-list*—The *select-list* is a list of expressions, separated by commas, specifying what is retrieved from the database. If an asterisk (\*) is specified, all columns of all tables in the **FROM** clause (`table-name` all columns of the named table) are selected. Aggregate functions and analytical functions are allowed in the *select-list*. See *Reference: Building Blocks, Tables, and Procedures > SQL Functions*.

---

**Note:** In Sybase IQ, scalar subqueries (nested selects) are allowed in the select list of the top level **SELECT**, as in SQL Anywhere and Adaptive Server Enterprise. Subqueries cannot be used inside a conditional value expression (for example, in a **CASE** statement).

Subqueries can also be used in a **WHERE** or **HAVING** clause predicate (one of the supported predicate types). However, inside the **WHERE** or **HAVING** clause, subqueries cannot be used inside a value expression or inside a **CONTAINS** or **LIKE** predicate. Subqueries are not allowed in the **ON** clause of outer joins or in the **GROUP BY** clause.

For more details on the use of subqueries, see *Reference: Building Blocks, Tables, and Procedures > SQL Language Elements > Expressions > Subqueries in Expressions* and *Reference: Building Blocks, Tables, and Procedures > SQL Language Elements > Search Conditions > Subqueries in Search Conditions*.

---

*alias-names* can be used throughout the query to represent the aliased expression. Alias names are also displayed by Interactive SQL at the top of each column of output from the **SELECT**

statement. If the optional *alias-name* is not specified after an expression, Interactive SQL displays the expression. If you use the same name or expression for a column alias as the column name, the name is processed as an aliased column, not a table column name.

**INTO host-variable-list**—Used in Embedded SQL only. It specifies where the results of the **SELECT** statement goes. There must be one *host-variable* item for each item in the *select-list*. Select list items are put into the host variables in order. An indicator host variable is also allowed with each *host-variable* so the program can tell if the select list item was NULL.

**INTO variable-list**—Used in procedures only. It specifies where the results of the **SELECT** statement go. There must be one variable for each item in the select list. Select list items are put into the variables in order.

**INTO table-name**—Used to create a table and fill the table with data.

If the table name starts with #, the table is created as a temporary table. Otherwise, the table is created as a permanent base table. For permanent tables to be created, the query must satisfy these conditions:

- The *select-list* contains more than one item, and the **INTO** target is a single *table-name* identifier, or
- The select-list contains a \* and the **INTO** target is specified as *owner.table*.

To create a permanent table with one column, the table name must be specified as *owner.table*. Omit the owner specification for a temporary table.

This statement causes a **COMMIT** before execution as a side effect of creating the table. RESOURCE authority is required to execute this statement. No permissions are granted on the new table: the statement is a short form for **CREATE TABLE** followed by **INSERT... SELECT**.

A **SELECT INTO** from a stored procedure or function is not permitted, as **SELECT INTO** is an atomic statement and you cannot do **COMMIT**, **ROLLBACK**, or some **ROLLBACK TO SAVEPOINT** statements in an atomic statement. For more information, see *System Administration Guide: Volume 2 > Using Procedures and Batches > Control Statements > Atomic Compound Statements* and *System Administration Guide: Volume 2 > Using Procedures and Batches > Transactions and Savepoints in Procedures*.

Tables created using this statement do not have a primary key defined. You can add a primary key using **ALTER TABLE**. A primary key should be added before applying any updates or deletes to the table; otherwise, these operations result in all column values being logged in the transaction log for the affected rows.

Use of this clause is restricted to valid SQL Anywhere queries. Sybase IQ extensions are not supported.

**INTO LOCAL TEMPORARY TABLE**—Creates a local, temporary table and populates it with the results of the query. When you use this clause, you do not need to start the temporary table name with #.

FROM table-list—Rows are retrieved from the tables and views specified in the *table-list*. Joins can be specified using join operators. For more information, see *FROM Clause*. A **SELECT** statement with no **FROM** clause can be used to display the values of expressions not derived from tables. For example:

```
SELECT @@version
```

displays the value of the global variable @@version. This is equivalent to:

```
SELECT @@version
FROM DUMMY
```

---

**Note:** If you omit the **FROM** clause, or if all tables in the query are in the `SYSTEM` dbspace, the query is processed by SQL Anywhere instead of Sybase IQ and might behave differently, especially with respect to syntactic and semantic restrictions and the effects of option settings. See the SQL Anywhere documentation for rules that might apply to processing.

If you have a query that does not require a **FROM** clause, you can force the query to be processed by Sybase IQ by adding the clause “FROM `iq_dummy`,” where `iq_dummy` is a one-row, one-column table that you create in your database.

---

WHERE search-condition—Specifies which rows are selected from the tables named in the **FROM** clause. It is also used to do joins between multiple tables. This is accomplished by putting a condition in the **WHERE** clause that relates a column or group of columns from one table with a column or group of columns from another table. Both tables must be listed in the **FROM** clause.

The use of the same **CASE** statement is not allowed in both the **SELECT** and the **WHERE** clause of a grouped query. See *Reference: Building Blocks, Tables, and Procedures > SQL Language Elements > Search Conditions*.

Sybase IQ also supports the disjunction of subquery predicates. Each subquery can appear within the **WHERE** or **HAVING** clause with other predicates and can be combined using the **AND** or **OR** operators. See *Reference: Building Blocks, Tables, and Procedures > SQL Language Elements > Search Conditions > Subqueries in Search Conditions > Disjunction of Subquery Predicates*.

GROUP BY—You can group by columns or alias names or functions. **GROUP BY** expressions must also appear in the select list. The result of the query contains one row for each distinct set of values in the named columns, aliases, or functions. The resulting rows are often referred to as groups since there is one row in the result for each group of rows from the table list. In the case of **GROUP BY**, all **NULL** values are treated as identical. Aggregate functions can then be applied to these groups to get meaningful results.

**GROUP BY** must contain more than a single constant. You do not need to add constants to the **GROUP BY** clause to select the constants in grouped queries. If the **GROUP BY** expression contains only a single constant, an error is returned and the query is rejected.

When **GROUP BY** is used, the select list, **HAVING** clause, and **ORDER BY** clause cannot reference any identifiers except those named in the **GROUP BY** clause. This exception applies: The *select-list* and **HAVING** clause may contain aggregate functions.

**ROLLUP** operator—The **ROLLUP** operator in the **GROUP BY** clause lets you analyze subtotals using different levels of detail. It creates subtotals that roll up from a detailed level to a grand total.

The **ROLLUP** operator requires an ordered list of grouping expressions to be supplied as arguments. **ROLLUP** first calculates the standard aggregate values specified in the **GROUP BY**. Then **ROLLUP** moves from right to left through the list of grouping columns and creates progressively higher-level subtotals. A grand total is created at the end. If *n* is the number of grouping columns, **ROLLUP** creates *n+1* levels of subtotals.

Restrictions on the **ROLLUP** operator:

- **ROLLUP** supports all of the aggregate functions available to the **GROUP BY** clause, but **ROLLUP** does not currently support **COUNT DISTINCT** and **SUM DISTINCT**.
- **ROLLUP** can be used only in the **SELECT** statement; you cannot use **ROLLUP** in a **SELECT** subquery.
- A multiple grouping specification that combines **ROLLUP**, **CUBE**, and **GROUP BY** columns in the same **GROUP BY** clause is not currently supported.
- Constant expressions as **GROUP BY** keys are not supported.

**GROUPING** is used with the **ROLLUP** operator to distinguish between stored NULL values and NULL values in query results created by **ROLLUP**.

**ROLLUP** syntax:

```
SELECT ... [ GROUPING ( column-name ) ... ] ...
GROUP BY [ expression [, ...]
| ROLLUP ( expression [, ...] ) ]
```

See *Reference: Building Blocks, Tables, and Procedures > SQL Language Elements > Expressions* for the format of an operator expression.

**GROUPING** takes a column name as a parameter and returns a Boolean value:

**Table 6. Values Returned by GROUPING with the ROLLUP Operator**

If the value of the result is	GROUPING returns
NULL created by a <b>ROLLUP</b> operation	1 (TRUE)
NULL indicating the row is a subtotal	1 (TRUE)
not created by a <b>ROLLUP</b> operation	0 (FALSE)
a stored NULL	0 (FALSE)

For **ROLLUP** examples, see *System Administration Guide: Volume 2 > Using OLAP*.

**CUBE** operator—The **CUBE** operator in the **GROUP BY** clause analyzes data by forming the data into groups in more than one dimension. **CUBE** requires an ordered list of grouping



expressions (dimensions) as arguments and enables the **SELECT** statement to calculate subtotals for all possible combinations of the group of dimensions.

Restrictions on the **CUBE** operator:

- **CUBE** supports all of the aggregate functions available to the **GROUP BY** clause, but **CUBE** does not currently support **COUNT DISTINCT** or **SUM DISTINCT**.
- **CUBE** does not currently support the inverse distribution analytical functions **PERCENTILE\_CONT** and **PERCENTILE\_DISC**.
- **CUBE** can be used only in the **SELECT** statement; you cannot use **CUBE** in a **SELECT** subquery.
- A multiple **GROUPING** specification that combines **ROLLUP**, **CUBE**, and **GROUP BY** columns in the same **GROUP BY** clause is not currently supported.
- Constant expressions as **GROUP BY** keys are not supported.

**GROUPING** is used with the **CUBE** operator to distinguish between stored NULL values and NULL values in query results created by **CUBE**.

**CUBE** syntax:

```
SELECT ... [ GROUPING ( column-name ) ... ] ...
GROUP BY [ expression [, ...]
| CUBE ( expression [, ...] ) ]
```

**GROUPING** takes a column name as a parameter and returns a Boolean value:

**Table 7. Values Returned by GROUPING with the CUBE Operator**

If the value of the result is	GROUPING returns
NULL created by a <b>CUBE</b> operation	1 (TRUE)
NULL indicating the row is a subtotal	1 (TRUE)
not created by a <b>CUBE</b> operation	0 (FALSE)
a stored NULL	0 (FALSE)

When generating a query plan, the Sybase IQ optimizer estimates the total number of groups generated by the **GROUP BY CUBE** hash operation. The **MAX\_CUBE\_RESULTS** database option sets an upper boundary for the number of estimated rows the optimizer considers for a hash algorithm that can be run. If the actual number of rows exceeds the **MAX\_CUBE\_RESULT** option value, the optimizer stops processing the query and returns the error message “Estimate number: *nnn* exceed the **DEFAULT\_MAX\_CUBE\_RESULT** of **GROUP BY CUBE** or **ROLLUP**”, where *nnn* is the number estimated by the optimizer. See *MAX\_CUBE\_RESULT Option* for information on setting the **MAX\_CUBE\_RESULT** option.

For **CUBE** examples, see *System Administration Guide: Volume 2 > Using OLAP*.

**HAVING** search-condition—Based on the group values and not on the individual row values. The **HAVING** clause can be used only if either the statement has a **GROUP BY** clause or if the

## Table UDFs and TPFs

select list consists solely of aggregate functions. Any column names referenced in the **HAVING** clause must either be in the **GROUP BY** clause or be used as a parameter to an aggregate function in the **HAVING** clause.

**ORDER BY**—Orders the results of a query. Each item in the **ORDER BY** list can be labeled as **ASC** for ascending order or **DESC** for descending order. Ascending is assumed if neither is specified. If the expression is an integer **n**, then the query results are sorted by the **n**th item in the select list.

In Embedded SQL, the **SELECT** statement is used for retrieving results from the database and placing the values into host variables with the **INTO** clause. The **SELECT** statement must return only one row. For multiple row queries, you must use cursors.

You cannot include a Java class in the **SELECT** list, but you can, for example, create a function or variable that acts as a wrapper for the Java class and then select it.

*row-limitation* clause—The row limitation clause allows you to return only a subset of the rows that satisfy the **WHERE** clause. Only one *row-limitation* clause can be specified at a time. When specifying this clause, an **ORDER BY** clause is required to order the rows in a meaningful manner.

The **LIMIT** and **OFFSET** arguments can be simple arithmetic expressions over host variables, integer constants, or integer variables. The **LIMIT** argument must evaluate to a value greater than or equal to 0. The **OFFSET** argument must evaluate to a value greater than or equal to 0. If *offset-expression* is not specified, the default is 0.

The row limitation clause **LIMIT** *offset-expression*, *limit-expression* is equivalent to **LIMIT** *limit-expression* **OFFSET** *offset-expression*.

The **LIMIT** keyword is disabled by default. Use the **RESERVED\_KEYWORDS** option to enable the **LIMIT** keyword.

### **Standards**

- SQL—ISO/ANSI SQL compliant.
- Sybase—Supported by Adaptive Server Enterprise, with some differences in syntax.

### **Permissions**

Must have **SELECT** permission on the named tables and views.

# API Reference for a\_v4\_extfn

Reference information for a\_v4\_extfn functions, methods, and attributes.

## Blob (a\_v4\_extfn\_blob)

---

Use the a\_v4\_extfn\_blob structure to represent a free-standing blob object.

### Implementation

```
typedef struct a_v4_extfn_blob {
    a_sql_uint64 (SQL_CALLBACK *blob_length)(a_v4_extfn_blob *blob);
    void (SQL_CALLBACK *open_istream)(a_v4_extfn_blob *blob,
a_v4_extfn_blob_istream **is);
    void (SQL_CALLBACK *close_istream)(a_v4_extfn_blob *blob,
a_v4_extfn_blob_istream *is);
    void (SQL_CALLBACK *release)(a_v4_extfn_blob *blob);
} a_v4_extfn_blob;
```

### Method Summary

Method Name	Data Type	Description
<b>blob_length</b>	a_sql_uint64	Returns the length, in bytes, of the specified blob.
<b>open_istream</b>	void	Opens an input stream that can be used to begin reading from the specified blob.
<b>close_istream</b>	void	Closes the input stream for the specified blob.

Method Name	Data Type	Description
<b>release</b>	void	Indicates that the caller is done with this blob and that the blob owner is free to release resources. After <b>release()</b> , referencing the blob results in an error. The owner usually deletes the memory when <b>release()</b> is called.

*Description*

The object `a_v4_extfn_blob` is used when:

- a table UDF needs to read LOB or CLOB data from a scalar input value
- a TPF needs to read LOB or CLOB data from a column in an input table

*Restrictions and Limitations*

None.

**blob\_length**

Use the `blob_length` v4 API method to return the length, in bytes, of the specified blob.

*Declaration*

```
a_sql_uint64 blob_length(
    a_v4_extfn_blob *
```

*Usage*

Returns the length, in bytes, of the specified blob.

*Parameters*

Parameter	Description
<b>blob</b>	The blob to return the length of.

*Returns*

The length of the specified blob.

**See also**

- *open\_istream* on page 197
- *close\_istream* on page 197

- *release* on page 198

## open\_istream

Use the `open_istream` v4 API method to open an input stream to read from a blob.

### *Declaration*

```
void open_istream(
    a_v4_extfn_blob *blob,
    a_v4_extfn_blob_istream **is
)
```

### *Usage*

Opens an input stream that can be used to begin reading from the specified blob.

### *Parameters*

Parameter	Description
<b>blob</b>	The blob to open the input stream on.
<b>is</b>	An output parameter identifying the returned open input stream.

### *Returns*

Nothing.

### **See also**

- *blob\_length* on page 196
- *close\_istream* on page 197
- *release* on page 198

## close\_istream

Use the `close_istream` v4 API method to close the input stream for the specified blob.

### *Declaration*

```
void close_istream(
    a_v4_extfn_blob *blob,
    a_v4_extfn_blob_istream *is
)
```

### *Usage*

Closes the input stream previously opened with the `open_istream` API.

*Parameters*

Parameter	Description
<b>blob</b>	The blob to close the input stream on.
<b>is</b>	A parameter identifying the input stream to close.

*Returns*

Nothing.

**See also**

- *blob\_length* on page 196
- *open\_istream* on page 197
- *release* on page 198

**release**

Use the `release v4` API method to indicate that the caller is done with the currently selected blob. Releasing enables the owner to free memory.

*Declaration*

```
void release(
a_v4_extfn_blob *blob
)
```

*Usage*

Indicates that the caller is done with this blob and that the blob owner is free to release resources. After **release()**, referencing the blob results in an error. The owner usually deletes the memory when **release()** is called.

*Parameters*

Parameter	Description
blob	The blob to release.

*Returns*

Nothing.

**See also**

- *blob\_length* on page 196
- *open\_istream* on page 197
- *close\_istream* on page 197

## Blob Input Stream (a\_v4\_extfn\_blob\_istream)

Use the `a_v4_extfn_blob_istream` structure to read blob data for a LOB or CLOB scalar input column, or LOB or CLOB column in an input table.

### Implementation

```
typedef struct a_v4_extfn_blob_istream {
    size_t (SQL_CALLBACK *get)( a_v4_extfn_blob_istream *is, void
*buf, size_t len );
    a_v4_extfn_blob          *blob;
    a_sql_byte               *beg;
    a_sql_byte               *ptr;
    a_sql_byte               *lim;
} a_v4_extfn_blob_istream;
```

### Method Summary

Method Name	Data Type	Description
<code>get</code>	<code>size_t</code>	Gets a specified amount of data from a blob input stream.

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>Blob</i>	<code>a_v4_extfn_blob</code>	The underlying blob structure for which this input stream was created.
<i>Beg</i>	<code>a_sql_byte</code>	A pointer to the beginning of the current chunk of data.
<i>Ptr</i>	<code>a_sql_byte</code>	A pointer to the current byte in the chunk of data.
<i>Lim</i>	<code>a_sql_byte</code>	A pointer to the end of the current chunk of data.

### get

Use the `get` v4 API method to get a specified amount of data from a blob input stream.

### Declaration

```
size_t get(
    a_v4_extfn_blob_istream *is,
```

## API Reference for a\_v4\_extfn

```
void *buf,  
size_t len  
)
```

### Usage

Gets a specified amount of data from a blob input stream.

### Parameters

Parameter	Description
<b>is</b>	The input stream to retrieve data from.
<b>buf</b>	The buffer to store the data in.
<b>len</b>	The amount of data to retrieve.

### Returns

The amount of data received.

## Column Data (a\_v4\_extfn\_column\_data)

The structure `a_v4_extfn_column_data` represents a single column's worth of data. This is used by the producer when generating result set data, or by the consumer when reading input table column data.

### Implementation

```
typedef struct a_v4_extfn_column_data {  
    a_sql_byte      *is_null;  
    a_sql_byte      null_mask;  
    a_sql_byte      null_value;  
  
    void            *data;  
    a_sql_uint32    *piece_len;  
    size_t          max_piece_len;  
  
    void            *blob_handle;  
} a_v4_extfn_column_data;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>is_null</i>	a_sql_byte *	Points to a byte where the NULL information for the value is stored.
<i>null_mask</i>	a_sql_byte	One or more bits used to represent the NULL value



Data Member	Data Type	Description
<i>null_value</i>	a_sql_byte	The value representing NULL
<i>data</i>	void *	Pointer to the data for the column. Depending on the type of fetch mechanism, either points to an address in the consumer, or an address where the data is stored in the UDF.
<i>piece_len</i>	a_sql_uint32 *	The actual length of data for variable-length data types
<i>max_piece_len</i>	size_t	The maximum data length allowed for this column.
<i>blob_handle</i>	void *	A non-NULL value means that the data for this column must be read using the <code>blob</code> API

### Description

The `a_v4_extfn_column_data` structure represents the data values and related attributes for a specific data column. This structure is used by the producer when generating result set data. Data producers are also expected to create storage for *data*, *piece\_len*, and the *is\_null* flag.

The *is\_null*, *null\_mask*, and *null\_value* data members indicate null in a column, and handle situations in which the null-bits are encoded into one byte for eight columns, or other cases in which a full byte is used for each column.

This example shows how to interpret the three fields used to represent NULL: *is\_null*, *null\_mask*, and *null\_value*.

```
is_value_null()
    return( (*is_null & null_mask) == null_value )

set_value_null()
    *is_null = ( *is_null & ~null_mask) | null_value

set_value_not_null()
    *is_null = *is_null & ~null_mask | (~null_value & null_mask)
```

### See also

- *get\_blob* on page 314

## Column List (a\_v4\_extfn\_column\_list)

---

Use the `a_v4_extfn_column_list` structure to provide a list of columns when describing **PARTITION BY** or to provide a list of columns when describing **TABLE\_UNUSED\_COLUMNS**.

### Implementation

```
typedef struct a_v4_extfn_column_list {
    a_sql_int32      number_of_columns;
    a_sql_uint32    column_indexes[1];    // there are
number_of_columns entries
} a_v4_extfn_column_list;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>number_of_columns</i>	a_sql_uint32	The number of columns in the list.
<i>column_indexes</i>	a_sql_uint32 *	A contiguous array of size <i>number_of_columns</i> with the column indexes (1-based).

### Description

The meaning of the contents of the column list changes, depending on whether the list is used with **TABLE\_PARTITIONBY** or **TABLE\_UNUSED\_COLUMNS**.

### See also

- *V4 API describe\_parameter and EXT FNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 141
- *Parallel TPF PARTITION BY Examples Using EXT FNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 143
- *EXT FNAPIV4\_DESCRIBE\_PARM\_TABLE\_UNUSED\_COLUMNS Attribute (Get)* on page 256
- *EXT FNAPIV4\_DESCRIBE\_PARM\_TABLE\_UNUSED\_COLUMNS Attribute (Set)* on page 272

## Column Order (a\_v4\_extfn\_order\_el)

---

Use the `a_v4_extfn_order_el` structure to describe the element order in a column.

### Implementation

```
typedef struct a_v4_extfn_order_el {
    a_sql_uint32    column_index;    // Index of the column in the
```

```

table (1-based)
  a_sql_byte      ascending;           // Nonzero if the column
is ordered "ascending".
} a_v4_extfn_order_el;

```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>column_index</i>	a_sql_uint32	Index of the column in the table (1-based).
<i>ascending</i>	a_sql_byte	Nonzero, if the column order is "ascending."

### Description

The `a_v4_extfn_order_el` structure describes a column and tells whether it should be in ascending or descending order. The `a_v4_extfn_orderby_list` structure holds an array of these structures. There is one `a_v4_extfn_order_el` structure for each column in the **ORDERBY** clause.

### See also

- *Order By List (a\_v4\_extfn\_orderby\_list)* on page 302

## Column Subset (a\_v4\_extfn\_col\_subset\_of\_input)

Use the `a_v4_extfn_col_subset_of_input` structure to declare that an output column has a value that is always taken from a particular input column to the UDF.

### Implementation

```

typedef struct a_v4_extfn_col_subset_of_input {
  a_sql_uint32  source_table_parameter_arg_num;  // arg_num of
the source table parameter
  a_sql_uint32  source_column_number;           // source column of
the source table
} a_v4_extfn_col_subset_of_input;

```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>source_table_parameter_arg_num</i>	a_sql_uint32 *	<i>arg_num</i> of the source table parameter
<i>source_column_number</i>	a_sql_uint32 *	Source column of the source table

### Description

The query optimizer uses the subset of input to infer logical properties of the values in the output column. For example, the number of distinct values in the input column is an upper

bound on the distinct values in the output column, and any local predicates on the input column also hold on the output column.

### See also

- *Describe Column Type (a\_v4\_extfn\_describe\_col\_type)* on page 277

## Describe API

---

The `_describe_extfn` function is a member of `a_v4_extfn_proc`. A UDF gets and sets logical properties using the `describe_column`, `describe_parameter`, and `describe_udf` properties in the `a_v4_extfn_proc_context` object.

### *\_describe\_extfn Declaration*

```
void (UDF_CALLBACK *_describe_extfn)(a_v4_extfn_proc_context
*cntxt );
)
```

### *Usage*

The `_describe_extfn` function describes the procedure evaluation to the server.

Each of the `describe_column`, `describe_parameter`, and `describe_udf` properties has an associated get and set method, a set of attribute types, and an associated data type for each attribute. The get methods retrieve information from the server; the set methods describe the logical properties of the UDF (such as the number of output columns or the number of distinct values for a output column) to the server.

### See also

- *\*describe\_column\_get* on page 204
- *\*describe\_column\_set* on page 221
- *\*describe\_parameter\_get* on page 238
- *\*describe\_parameter\_set* on page 257
- *\*describe\_udf\_get* on page 273
- *\*describe\_udf\_set* on page 275
- *External Function (a\_v4\_extfn\_proc)* on page 284

## \*describe\_column\_get

The `describe_column_get` v4 API method is used by the table UDF to retrieve properties about an individual column of a table parameter.

### *Declaration*

```
a_sql_int32 (SQL_CALLBACK *describe_column_get)(
    a_v4_extfn_proc_context    *cntxt,
    a_sql_uint32                arg_num,
```

```

a_sql_uint32          column_num,
a_v4_extfn_describe_parm_type describe_type,
void                 *describe_buffer,
size_t              describe_buffer_len );

```

### Parameters

Parameter	Description
<b>cntxt</b>	The procedure context object for this UDF.
<b>arg_num</b>	The ordinal of the table parameter (0 is the result table, 1 for first input argument).
<b>column_num</b>	The ordinal of the column starting at 1.
<b>describe_type</b>	A selector indicating what property to retrieve.
<b>describe_buffer</b>	A structure that holds the describe information for the specified property to get from the server. The specific structure or data type is indicated by the <b>describe_type</b> parameter.
<b>describe_buffer_length</b>	The length, in bytes, of the <b>describe_buffer</b> .

### Returns

On success, returns the number of bytes written to the **describe\_buffer**. If an error occurs, or no property is retrieved, this function returns one of the generic `describe_column` errors.

### See also

- *\*describe\_column\_set* on page 221

### Attributes for \*describe\_column\_get

Code showing the attributes for `describe_column_get` v4 API method.

```

typedef enum a_v4_extfn_describe_col_type {
    EXTFNAPIV4_DESCRIBE_COL_NAME,
    EXTFNAPIV4_DESCRIBE_COL_TYPE,
    EXTFNAPIV4_DESCRIBE_COL_WIDTH,
    EXTFNAPIV4_DESCRIBE_COL_SCALE,
    EXTFNAPIV4_DESCRIBE_COL_CAN_BE_NULL,
    EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES,
    EXTFNAPIV4_DESCRIBE_COL_IS_UNIQUE,
    EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT,
    EXTFNAPIV4_DESCRIBE_COL_CONSTANT_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER,
    EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT,
} a_v4_extfn_describe_col_type;

```

### *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Get)*

The **EXTFNAPIV4\_DESCRIBE\_COL\_NAME** attribute indicates the column name. Used in a `describe_column_get` scenario.

#### *Data Type*

char [ ]

#### *Description*

The column name. This property is valid only for table arguments.

#### *Usage*

If a UDF gets this property, then the name of the specified column is returned.

#### *Returns*

On success, returns the length of the column name.

On failure, returns one of the generic `describe_column` errors, or:

- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – get error returned if the state is not greater than Initial.
- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – get error returned if the buffer length has insufficient characters or is 0 length.
- **EXTFNAPIV4\_DESCRIBE\_NON\_TABLE\_PARAMETER** – get error returned if the parameter is not a **TABLE** parameter.

#### *Query Processing States*

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

#### **See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Set)* on page 222
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)**

The `EXTFNAPIV4_DESCRIBE_COL_TYPE` attribute indicates the data type of the column. Used in a `describe_column_get` scenario.

***Data Type***

`a_sql_data_type`

***Description***

The data type of the column. This property is valid only for table arguments.

***Usage***

If a UDF gets this property, then returns the data type of the specified column.

***Returns***

On success, the `sizeof(a_sql_data_type)` is returned.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_sql_data_type`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.

***Query Processing States***

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

**See also**

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)**

The `EXTFNAPIV4_DESCRIBE_COL_WIDTH` attribute indicates the width of the column. Used in a `describe_column_get` scenario.

***Data Type***

`a_sql_uint32`

### *Description*

The width of a column. Column width is the amount of storage, in bytes, required to store a value of the associated data type. This property is valid only for table arguments.

### *Usage*

If a UDF gets this property, then returns the width of the column as defined in the **CREATE PROCEDURE** statement.

### *Returns*

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the query processing state is not greater than Initial.

### *Query Processing States*

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### **See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

### *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Get)*

The `EXTFNAPIV4_DESCRIBE_COL_SCALE` attribute indicates the scale of the column. Used in a `describe_column_get` scenario.

### *Data Type*

`a_sql_uint32`

### *Description*

The scale of a column. For arithmetic data types, parameter scale is the number of digits to the right of the decimal point in a number. This property is valid only for table arguments.

### *Usage*

If the UDF gets this property, returns the scale of the column as defined in the **CREATE PROCEDURE** statement. This property is valid only for arithmetic data types.



### Returns

On success, returns the `sizeof(a_sql_uint32)` if the value was returned, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – get error returned if the scale is unavailable for the data type of the specified column.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the query processing state is not greater than `Initial`.

### Query Processing States

- Annotation state
- Query Optimization State
- Plan Building state
- Execution state

### See also

- *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Set)* on page 226
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL (Get)

The `EXTFNAPIV4_DESCRIBE_COL_CAN_BE_NULL` attribute indicates if the column can be NULL. Used in a `describe_column_get` scenario.

### Data Type

`a_sql_byte`

### Description

True, if the column can be NULL. This property is valid only for table arguments. This property is valid only for argument 0.

### Usage

If a UDF gets this property, returns 1 if the column can be NULL, and returns 0 if otherwise.

### Returns

On success, returns the `sizeof(a_sql_byte)` if the attribute is available, or:

## API Reference for a\_v4\_extfn

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – returned if the attribute was not available to get. This can happen if the column was not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the specified argument is an input table and the query processing state is not greater than Plan building state.

### *Query Processing States*

Valid in:

- Execution state

### **See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL (Set)* on page 227
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_COL\_DISTINCT\_VALUES (Get)

The `EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES` attribute describes the distinct values for a column. Used in a `describe_column_get` scenario.

### *Data Type*

`a_v4_extfn_estimate`

### *Description*

The estimated number of distinct values for a column. This property is valid only for table arguments.

### *Usage*

If a UDF gets this property, it returns the estimated number of distinct values for a column.

### *Returns*

On success, returns the `sizeof(a_v4_extfn_estimate)`, if it returns a value, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – returned if the attribute was not available to get. This can happen if the column was not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_v4_extfn_estimate`.

- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the specified argument is an input table and the query processing state is greater than Optimization.

### Query Processing States

Valid in:

- Plan Building state
- Execution state

### Example

Consider this procedure definition and code fragment in the `_describe_extfn` API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2 INT ) )
RESULTS ( r1 INT, r2 INT, r3 INT )
EXTERNAL 'my_tpf_proc@mylibrary';
```

```
CREATE TABLE T( x INT, y INT, z INT );
```

```
select * from my_tpf( 'test', TABLE( select x,y from T ) )
```

This example shows how a TPF gets the number of distinct values for column one of the input table. A TPF may want to get this value, if it is beneficial for choosing an appropriate processing algorithm.

```
my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state == EXTFNAPIV4_STATE_PLAN_BUILDING ) {
        a_v4_extfn_estimate num_distinct;

        a_sql_int32 ret = 0;

        // Get the number of distinct values expected from the first
column
        // of the table input parameter 'col_table'
        ret = cntxt->describe_column_get( cntxt, 2, 1
            EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES,
            &num_distinct,
            sizeof(a_v4_extfn_estimate) );

        // default algorithm is 1
        _algorithm = 1;

        if( ret > 0 ) {
            // choose the best algorithm for sample size.

            if ( num_distinct.value < 100 ) {
                // use faster algorithm for small distinct values.
                _algorithm = 2;
            }
        }
        else {
            if ( ret < 0 ) {
```

```

        // Handle the error
        // or continue with default algorithm
    } else {
        // Attribute was unavailable
        // We will use the default algorithm.
    }
}
}
}

```

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_DISTINCT\_VALUES (Set)* on page 228
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE (Get)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE** attribute indicates if a column is unique in the table. Used in a `describe_column_get` scenario.

**Data Type**

a\_sql\_byte

**Description**

True, if the column is unique within the table. This property is valid only for table arguments.

**Usage**

If the UDF gets this property, then returns 1 if the column is unique, and 0 otherwise.

**Returns**

On success, returns the `sizeof(a_sql_byte)` or:

- **EXTFNAPIV4\_DESCRIBE\_NOT\_AVAILABLE** – if the attribute was unavailable to get. This can happen if the column was not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – get error returned if the describe buffer is not the size of `a_sql_byte`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – get error returned if the query processing state is not greater than Initial.

**Query Processing States**

Valid in:

- Annotation state
- Query Optimization state

- Plan Building state
- Execution state

### See also

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE (Set)* on page 229
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Get)

The **EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT** attribute indicates if a column is constant. Used in a `describe_column_get` scenario.

#### *Data Type*

`a_sql_byte`

#### *Description*

True, if the column is constant for the lifetime of the statement. This property is valid only for input table arguments.

#### *Usage*

If a UDF gets this property, the return value is 1 if the column is constant for the lifetime of the statement and 0 otherwise. Input table columns are constant, if the column in the select list for the input table is a constant expression or NULL.

#### *Returns*

On success, returns the `sizeof(a_sql_byte)`, if the value was returned, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – the attribute is not available to get. Returned, if the column is not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned, if the describe buffer is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned, if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned, if the specified argument is not an input table.

#### *Query Processing States*

Valid in:

- Annotation state
- Query Optimization state

- Plan Building state
- Execution state

### See also

- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Set)* on page 230
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Get)

The **EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE** attribute indicates the constant value of a column. Used in a `describe_column_get` scenario.

### Data Type

`an_extfn_value`

### Description

The value of the column, if it is constant for the statement lifetime. If `EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT` for this column returns true, this value is available. This property is valid only for table arguments.

### Usage

For columns of input tables that have a constant value, the value is returned. If the value is unavailable, then NULL is returned.

### Returns

On success, returns the `sizeof(a_sql_byte)`, if the value was returned, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – the attribute is not available to get. Returned, if the column is not involved in the query, or if the value is not considered constant.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned, if the describe buffer is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned, if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned, if the specified argument is not an input table.

### Query Processing States

Valid in:

- Annotation state

- Query Optimization state
- Plan Building state
- Execution state

### See also

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Set)* on page 230
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER (Get)

The **EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER** attribute indicates if a column in the result table is used by the consumer. Used in a `describe_column_get` scenario.

### Data Type

`a_sql_byte`

### Description

Used either to determine whether a column in the result table is used by the consumer, or to indicate that a column in an input is not needed. Valid for table arguments. Allows the user to set or retrieve information about a single column, whereas the similar attribute `EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS` sets or retrieves information about all columns in a single call.

### Usage

The UDF queries this property to determine if a result table column is required by the consumer. This can help the UDF avoid unnecessary work for unused columns.

### Returns

On success, returns the `sizeof(a_sql_byte)` or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – if the attribute was unavailable to get. This can happen if the column was not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_v4_extfn_estimate`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned if the argument specified is not argument 0.

### Query Processing States

Valid during:

## API Reference for a\_v4\_extfn

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

The PROCEDURE definition and code fragment in the `_describe_extfn` API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2
INT ) )
    RESULTS ( r1 INT, r2 INT, r3 INT )
    EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );

select r2,r3 from my_tpf( 'test', TABLE( select x,y from T ) )
```

When this TPF runs, it is beneficial to know if the user has selected column `r1` of the result set. If the user does not need `r1`, calculations for `r1` may be unnecessary and we do not need to produce it for the server.

```
my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state > EXTFNAPIV4_STATE_INITIAL ) {
        a_sql_byte col_is_used = 0;
        a_sql_int32 ret = 0;

        ret = cntxt->describe_column_get( cntxt, 0, 1,
            EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER,
            &col_is_used,
            sizeof(a_sql_byte) );

        if( ret < 0 ) {
            // Handle the error.
        }
    }
}
```

### See also

- [EXTFNAPIV4\\_DESCRIBE\\_COL\\_IS\\_USED\\_BY\\_CONSUMER \(Set\)](#) on page 231
- [Generic describe\\_column Errors](#) on page 321
- [Query Processing States](#) on page 121

### EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE (Get)

The **EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE** attribute indicates the minimum value for a column. Used in a `describe_column_get` scenario.

#### Data Type

`an_extfn_value`



### Description

The minimum value for a column, if available. Valid only for argument 0 and table arguments.

### Usage

If a UDF gets the **EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE** property, the minimum value of the column data is returned in the **describe\_buffer**. If the input table is a base table, the minimum value is based on all of the column data in the table and is accessible only if there is an index on the table column. If the input table is the result of another UDF, the minimum value is the **EXTFNAPIV4\_DESCRIBE\_COL\_TYPE** set by that UDF.

The data type for this property is different for different columns. The UDF can use **EXTFNAPIV4\_DESCRIBE\_COL\_TYPE** to determine the data type of the column. The UDF can also use **EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH** to determine the storage requirements of the column, to provide an equivalently sized buffer to hold the value.

**describe\_buffer\_length** allows the server to determine if the buffer is valid.

If the **EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE** property is unavailable, **describe\_buffer** is NULL.

### Returns

On success, returns the **describe\_buffer\_length**, or:

- **EXTFNAPIV4\_DESCRIBE\_NOT\_AVAILABLE** – if the attribute was unavailable to get. Returned if the column was not involved in the query or the minimum value was unavailable for the requested column.

On failure, returns one of the generic **describe\_column** errors, or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – Get error returned, if the describe buffer is not large enough to hold the minimum value.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – Get error returned if the state is not greater than Initial.

### Query Processing States

Valid in any state except Initial state:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### Example

The procedure definition and code fragment in the **\_describe\_extfn** API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2 INT ) )
```

```

RESULTS ( r1 INT, r2 INT, r3 INT )
EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );

select * from my_tpf( 'test', TABLE( select x,y from T ) )

```

This example illustrates how a TPF would get the minimum value for column two of the input table, for internal optimization purposes.

```

my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state > EXTFNAPIV4_STATE_INITIAL ) {
        a_sql_int32 min_value = 0;
        a_sql_int32 ret = 0;

        // Get the minimum value of the second column of the
        // table input parameter 'col_table'

        ret = cntxt->describe_column_get( cntxt, 2, 2
            EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE,
            &min_value,
            sizeof(a_sql_int32) );

        if( ret < 0 ) {
            // Handle the error.
        }
    }
}

```

### See also

- *Query Processing States* on page 121
- *EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE (Set)* on page 233
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *Generic describe\_column Errors* on page 321

### EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE (Get)

The **EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE** attribute indicates the maximum value for the column. Used in a `describe_column_get` scenario.

#### Data Type

an\_extfn\_value

#### Description

The maximum value for a column. This property is valid only for argument 0 and table arguments.

### Usage

If a UDF gets the `EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE` property, then the maximum value of the column data is returned in the `describe_buffer`. If the input table is a base table, the maximum value is based on all of the column data in the table and is accessible only if there is an index on the table column. If the input table is the result of another UDF, the maximum value is the `COL_MAXIMUM_VALUE` set by that UDF.

The data type for this property is different for different columns. The UDF can use `EXTFNAPIV4_DESCRIBE_COL_TYPE` to determine the data type of the column. The UDF can also use `EXTFNAPIV4_DESCRIBE_COL_WIDTH` to determine the storage requirements of the column, to provide an equivalently sized buffer to hold the value.

`describe_buffer_length` allows the server to determine if the buffer is valid.

If `EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE` is unavailable, `describe_buffer` is NULL.

### Returns

On success, returns the `describe_buffer_length` or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – If the attribute was unavailable to get. This can happen if the column was uninvolved in the query, or if the maximum value was unavailable for the requested column.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – Get error returned if the describe buffer is not large enough to hold the maximum value.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – Get error returned if the state is not greater than Initial.

### Query Processing States

Valid in any state except Initial state:

- Annotation state
- Query Optimization state
- Plan building state
- Execution state

### Example

The **PROCEDURE** definition and code fragment in the `_describe_extfn` API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2 INT ) )
    RESULTS ( r1 INT, r2 INT, r3 INT )
    EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );
```

```
select * from my_tpf( 'test', TABLE( select x,y from T ) )
```

This example illustrates how a TPF would get the maximum value for column two of the input table, for internal optimization purposes.

```
my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state > EXTFNAPIV4_STATE_INITIAL ) {
        a_sql_int32 max_value = 0;
        a_sql_int32 ret = 0;

        // Get the maximum value of the second column of the
        // table input parameter 'col_table'
        ret = cntxt->describe_column_get( cntxt, 2, 2
            EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE,
            &max_value,
            sizeof(a_sql_int32) );

        if( ret < 0 ) {
            // Handle the error.
        }
    }
}
```

### See also

- *Query Processing States* on page 121
- *EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE (Set)* on page 235
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *Generic describe\_column Errors* on page 321

### EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Get)

The **EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT** attribute sets a subset of the values specified in an input column. Using this attribute in a `describe_column_get` scenario returns an error.

#### *Data Type*

`a_v4_extfn_col_subset_of_input`

#### *Description*

Column values are a subset of the values specified in an input column.

#### *Usage*

This attribute can be set only.

### Returns

Returns the error EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE.

### Query Processing States

Error EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE is returned in any state.

### See also

- *EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Set)* on page 236
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

## **\*describe\_column\_set**

The `describe_column_set` v4 API method sets UDF column-level properties on the server.

### Description

Column-level properties describe various characteristics about columns in the result set or input tables in a TPF. For example, a UDF can tell the server that a column in its result set will have only ten distinct values.

### Declaration

```
a_sql_int32 (SQL_CALLBACK *describe_column_set)(
    a_v4_extfn_proc_context    *cntxt,
    a_sql_uint32               arg_num,
    a_sql_uint32               column_num,
    a_v4_extfn_describe_udf_type describe_type,
    const void                 *describe_buffer,
    size_t                     describe_buffer_len );
```

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object for this UDF.
<code>arg_num</code>	The ordinal of the table parameter (0 is the result table, 1 for first input argument).
<code>column_num</code>	The ordinal of the column starting at 1.
<code>describe_type</code>	A selector indicating what property to set.

Parameter	Description
<b>describe_buffer</b>	A structure that holds the describe information for the specified property to set on the server. The specific structure or data type is indicated by the <b>describe_type</b> parameter.
<b>describe_buffer_length</b>	The length, in bytes, of the <b>describe_buffer</b> .

### Returns

On success, returns the number of bytes written to the **describe\_buffer**. If an error occurs, or no property is retrieved, this function returns one of the generic `describe_column` errors.

### See also

- *\*describe\_column\_get* on page 204

### Attributes for \*describe\_column\_set

Code showing the attributes for `describe_column_set`.

```
typedef enum a_v4_extfn_describe_col_type {
    EXTFNAPIV4_DESCRIBE_COL_NAME,
    EXTFNAPIV4_DESCRIBE_COL_TYPE,
    EXTFNAPIV4_DESCRIBE_COL_WIDTH,
    EXTFNAPIV4_DESCRIBE_COL_SCALE,
    EXTFNAPIV4_DESCRIBE_COL_CAN_BE_NULL,
    EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES,
    EXTFNAPIV4_DESCRIBE_COL_IS_UNIQUE,
    EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT,
    EXTFNAPIV4_DESCRIBE_COL_CONSTANT_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER,
    EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT,
} a_v4_extfn_describe_col_type;
```

### EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Set)

The **EXTFNAPIV4\_DESCRIBE\_COL\_NAME** attribute indicates a column name. Used in a `describe_column_set` scenario.

### Data Type

char[ ]

### Description

The column name. This property is valid only for table arguments.

### Usage

For argument 0, if the UDF sets this property, the server compares the value with the name of the column supplied in the **CREATE PROCEDURE** statement. The comparison ensures that the **CREATE PROCEDURE** statement has the same column name as expected by the UDF.

### Returns

On success, returns the length of the column name.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not Annotation.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – set error returned if the parameter is not a Table parameter.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the length of input column name exceeds 128 characters or if the input column name and column name stored in the catalog do not match.

### Query Processing States

- Annotation state

### Example

```
short desc_rc = 0;
char name[7] = 'column1';
// Verify that the procedure was created with the second column
of the result table as an int
if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
    desc_rc = ctx->describe_column_set( ctx, 0, 2,
EXTFNAPIV4_DESCRIBE_COL_NAME,
                                name,
                                sizeof(name) );

    if( desc_rc < 0 ) {
        // handle the error.
    }
}
```

### See also

- *EXTFNAPIV4\_DESCRIBE\_COL\_NAME (Get)* on page 206
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

***EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)***

The `EXTFNAPIV4_DESCRIBE_COL_TYPE` attribute indicates the data type of the column. Used in a `describe_column_set` scenario.

***Data Type***

`a_sql_data_type`

***Description***

The data type of the column. This property is valid only for table arguments.

***Usage***

For argument zero, if the UDF sets this property, then the server compares the value with the data type of the column supplied in the **CREATE PROCEDURE** statement. This allows the UDF to ensure the **CREATE PROCEDURE** statement has the same data type as expected by the UDF.

***Returns***

On success, returns the `a_sql_data_type`.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – Set error returned if the describe buffer is not the size of `a_sql_data_type`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – Set error returned if the state is not Annotation.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – Set error returned if the input data type and the data type stored in the catalog do not match,.

***Query processing states***

- Annotation state

***Example***

```
short desc_rc = 0;
a_sql_data_type type = DT_INT;

    // Verify that the procedure was created with the second column of
    the result table as an int
    if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
        desc_rc = ctx->describe_column_set( ctx, 0, 2,
EXTFNAPIV4_DESCRIBE_COL_TYPE,
        &type,
        sizeof(a_sql_data_type) );
        if( desc_rc < 0 ) {
            // handle the error.
        }
    }
}
```



**See also**

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH** attribute indicates the width of the column. Used in a `describe_column_set` scenario.

*Data Type*

a\_sql\_uint32

*Description*

The width of a column. Column width is the amount of storage, in bytes, required to store a value of the associated data type. This property is valid only for table arguments.

*Usage*

If the UDF sets this property, the server compares the value with the width of the column supplied in the **CREATE PROCEDURE** statement. This allows the UDF to ensure the **CREATE PROCEDURE** statement has the same column width as expected by the UDF.

*Returns*

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_column` errors, or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – set error returned if the describe buffer is not the size of `a_sql_uint32`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – set error returned if the query processing state is not Annotation.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE\_VALUE** – set error returned if the input width and width stored in the catalog do not match.

*Query Processing States*

Valid in:

- Annotation state

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_SCALE** attribute indicates the scale of the column. Used in a `describe_column_set` scenario.

**Data Type**

`a_sql_uint32`

**Description**

The scale of a column. For arithmetic data types, parameter `scale` is the number of digits to the right of the decimal point in a number. This property is valid only for table arguments.

**Usage**

If the UDF sets this property, the server compares the value with the scale of the column supplied in the **CREATE PROCEDURE** statement. This allows the UDF to ensure the **CREATE PROCEDURE** statement has the same column width as expected by the UDF. This property is valid only for arithmetic data types.

**Returns**

On success, returns the `sizeof(a_sql_uint32)`, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – set error returned if the scale is not available for the data type of the specified column.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the `describe` buffer is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the query processing state is not Annotation.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the input `scale` and `scale` stored in the catalog do not match.

**Query Processing States**

Valid in:

- Annotation state

**Example**

```
short desc_rc = 0;
a_sql_uint32 scale = 0;

// Verify that the procedure has a scale of zero for the
second result table column.
if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
    desc_rc = ctx->describe_column_set( ctx, 0, 2,
EXTFNAPIV4_DESCRIBE_COL_SCALE,
```

```

        &scale,
        sizeof(a_sql_data_type) );
    if( desc_rc < 0 ) {
        // handle the error.
    }
}

```

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_SCALE (Get)* on page 208
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL** attribute indicates if the column can be null. Used in a `describe_column_set` scenario.

***Data Type***

`a_sql_byte`

***Description***

True, if the column can be NULL. This property is valid only for table arguments. This property is valid only for argument 0.

***Usage***

The UDF can set this property for a result table column if that column can be NULL. If the UDF does not explicitly set this property, it is assumed that the column can be NULL. The server can use this information during the Optimization state.

***Returns***

On success, returns the `sizeof(a_sql_byte)` if the attribute was set or:

- **EXTFNAPIV4\_DESCRIBE\_NOT\_AVAILABLE** – returned if the attribute was unavailable to set, which may happen if the column was uninvolved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – set error returned if the describe buffer is not the size of `a_sql_byte`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – set error returned if the state is not equal to **OPTIMIZATION**.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE\_VALUE** – set error returned if the UDF attempts to set this attribute to a value other than 0 or 1.

***Query Processing States***

Valid in:

- Optimization state

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL (Get)* on page 209
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_DISTINCT\_VALUES (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_DISTINCT\_VALUES** attribute describes the distinct values for a column. Used in a `describe_column_set` scenario.

*Data Type*

`a_v4_extfn_estimate`

*Description*

The estimated number of distinct values for a column. This property is valid only for table arguments.

*Usage*

The UDF can set this property if it knows how many distinct values a column can have in its result table. The server uses this information during the Optimization state.

*Returns*

On success, returns the `sizeof(a_v4_extfn_estimate)`, if it sets the value, or:

- **EXTFNAPIV4\_DESCRIBE\_NOT\_AVAILABLE** – returned if the attribute was unavailable to set. This can happen if the column was not involved in the query.

On failure, returns:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – set error returned if the describe buffer is not the size of `a_v4_extfn_estimate`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – set error returned if the state is not equal to Optimization.

*Query Processing States*

Valid in:

- Optimization state

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_DISTINCT\_VALUES (Get)* on page 210
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE** attribute indicates if the column is unique in the table. Used in a `describe_column_set` scenario.

***Data Type***

`a_sql_byte`

***Description***

True, if the column is unique within the table. This property is valid only for table arguments.

***Usage***

The UDF can set this property if it knows the result table column value is unique. The server uses this information during the Optimization state. The UDF can set this property only for argument 0.

***Returns***

On success, returns the `sizeof(a_sql_byte)` or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – if the attribute was not available to set. This can happen if the column was not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the describe buffer is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the query processing state is not Optimization.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the **arg\_num** is not zero.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF attempts to set this attribute to a value other than 0 or 1.

***Query Processing States***

Valid in:

- Optimization state

**See also**

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE (Get)* on page 212
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT** attribute indicates if the column is constant. Used in a `describe_column_set` scenario.

*Data Type*

a\_sql\_byte

*Description*

True, if the column is constant for the lifetime of the statement. This property is valid only for input table arguments.

*Usage*

This is a read only property. All attempts to set it return **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE**.

*Returns*

- **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE** – this is a read-only property; all attempts to set return this error.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – set error returned, if the state is not Optimization.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER** – set error returned, if the **arg\_num** is not zero.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE\_VALUE** – set error returned, if the UDF attempts to set this attribute to a value other than 0 or 1.

*Query Processing States*

Not applicable.

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT (Get)* on page 213
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE** attribute indicates the constant value of the column. Used in a `describe_column_set` scenario.

*Data Type*

an\_extfn\_value

*Description*

The value of the column, if it is constant for the statement lifetime. If `EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT` for this column returns true, this value is available. This property is valid only for table arguments.

*Usage*

This property is read-only.

*Returns*

- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE` – this is a read-only property; all attempts to set return this error.

*Query Processing States*

Not applicable.

**See also**

- *Generic describe\_column Errors* on page 321
- *EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE (Get)* on page 214
- *Query Processing States* on page 121

*EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER (Set)*

The `EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER` attribute indicates if the column in the result table is used by the consumer. Used in a `describe_column_set` scenario.

*Data Type*

`a_sql_byte`

*Description*

Used either to determine whether a column in the result table is used by the consumer, or to indicate that a column in an input is not needed. Valid for table arguments. Allows the user to set or retrieve information about a single column, whereas the similar attribute `EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS` sets or retrieves information about all columns in a single call.

*Usage*

The UDF sets `EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER` on columns in an input table to inform the producer that it does not need values for the column.

*Returns*

On success, returns the `sizeof(a_sql_byte)` or:

## API Reference for a\_v4\_extfn

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – if the attribute was not available to set. This can happen if the column was not involved in the query.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the describe buffer is not the size of `a_v4_extfn_estimate`.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the argument specified is argument 0.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not equal to Optimization.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the value the UDF is setting is not 0 or 1.

### Query Processing States

Valid during:

- Optimization state

The PROCEDURE definition and code fragment in the `_describe_extfn` API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2
INT ) )
    RESULTS ( r1 INT, r2 INT, r3 INT )
    EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );

select r2,r3 from my_tpf( 'test', TABLE( select x,y from T ) )
```

When this TPF runs, it is beneficial for the server to know if column `y` is used by this TPF. If the TPF does not need `y`, the server can use this knowledge for optimization and does not send this column information to the TPF.

```
my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_byte col_is_used = 0;
        a_sql_int32 ret = 0;

        ret = cntxt->describe_column_get( cntxt, 2, 2,
            EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER,
            &col_is_used,
            sizeof(a_sql_byte) );

        if( ret < 0 ) {
            // Handle the error.
        }
    }
}
```



**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER (Get)* on page 215
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE** attribute indicates the minimum value for the column. Used in a `describe_column_set` scenario.

***Data Type***

`an_extfn_value`

***Description***

The minimum value a column can have, if available. Only valid for argument 0.

***Usage***

The UDF can set `EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE`, if it knows what the minimum data value of the column is. The server can use this information during optimization.

The UDF can use `EXTFNAPIV4_DESCRIBE_COL_TYPE` to determine the data type of the column, and `EXTFNAPIV4_DESCRIBE_COL_WIDTH` to determine the storage requirements of the column, to provide an equivalently sized buffer to hold the value to set.

***Returns***

On success, returns the `describe_buffer_length`, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – if the attribute cannot be set. Returned if the column was not involved in the query or the minimum value was not available for the requested column.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned, if the describe buffer is not large enough to hold the minimum value.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned, if the state is not equal to `Optimization`.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned, if the **arg\_num** is not 0.

***Query Processing States***

Valid in:

- `Optimization state`

### Example

The **PROCEDURE** definition and UDF code fragment that implements the `_describe_extfn` callback API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2 INT ) )
  RESULTS ( r1 INT, r2 INT, r3 INT )
  EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );

select * from my_tpf( 'test', TABLE( select x,y from T ) )
```

This example shows a TPF where it is useful to the server (or to another TPF that takes the result of this TPF as input) to know the minimum value of result set column one. In this instance, the minimum output value of column one is 27.

```
my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_int32 min_value = 27;
        a_sql_int32 ret = 0;

// Tell the server what the minimum value of the first column
// of our result set will be.

        ret = cntxt->describe_column_set( cntxt, 0, 1
            EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE,
            &min_value,
            sizeof(a_sql_int32) );

        if( ret < 0 ) {
            // Handle the error.
        }
    }
}
```

### See also

- *Query Processing States* on page 121
- *EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE (Get)* on page 216
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *Generic describe\_column Errors* on page 321

**EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE (Set)**

The **EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE** attribute indicates the maximum value for the column. Used in a `describe_column_set` scenario.

***Data Type***

`an_extfn_value`

***Description***

The maximum value for a column. This property is valid only for argument 0 and table arguments.

***Usage***

The UDF can set `EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE`, if it knows what the maximum data value of the column is. The server can use this information during optimization.

The UDF can use `EXTFNAPIV4_DESCRIBE_COL_TYPE` to determine the data type of the column, and `EXTFNAPIV4_DESCRIBE_COL_WIDTH` to determine the storage requirements of the column, to provide an equivalently sized buffer to hold the value to set.

`describe_buffer_length` is the `sizeof()` this buffer.

***Returns***

On success, returns the `describe_buffer_length`, if the value was set, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE` – if the attribute could not be set. Returned if the column was not involved in the query or the maximum value was not available for the requested column.

On failure, returns one of the generic `describe_column` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned, if the describe buffer is not large enough to hold the maximum value.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – Set error returned, if the query processing state is not equal to Optimization.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned, if the `arg_num` is not 0.

***Query Processing States***

Valid in:

- Optimization state

***Example***

The PROCEDURE definition and and UDF code fragment that implements the `_describe_extfn` callback API function:

```
CREATE PROCEDURE my_tpf( col_char char(10), col_table TABLE( c1 INT,
c2 INT ) )
    RESULTS ( r1 INT, r2 INT, r3 INT )
    EXTERNAL 'my_tpf_proc@mylibrary';

CREATE TABLE T( x INT, y INT, z INT );

select * from my_tpf( 'test', TABLE( select x,y from T ) )
```

This example shows a TPF where it is useful to the server (or to another TPF that takes the result of this TPF as input) to know the maximum value of result set column one. In this instance, the maximum output value of column one is 500000.

```
my_tpf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_int32 max_value = 500000;
        a_sql_int32 ret = 0;

        // Tell the server what the maximum value of the first column
        // of our result set will be.

        ret = cntxt->describe_column_set( cntxt, 0, 1
            EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE,
            &max_value,
            sizeof(a_sql_int32) );

        if( ret < 0 ) {
            // Handle the error.
        }
    }
}
```

### See also

- *Query Processing States* on page 121
- *EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE (Get)* on page 218
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_TYPE (Set)* on page 224
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Get)* on page 207
- *EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH (Set)* on page 225
- *Generic describe\_column Errors* on page 321

### EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Set)

The **EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT** attribute sets a subset of the values specified in an input column. Used in a `describe_column_set` scenario.

### Data Type

`a_v4_extfn_col_subset_of_input`

*Description*

Column values are a subset of the values specified in an input column.

*Usage*

Setting this describe attribute informs the query optimizer that the indicated column values are a subset of those values specified in an input column. For example, consider a filter TPF that consumes a table and filters out rows based on a function. In such a case, the return table is a subset of the input table. Setting

**EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT** for the filter TPF optimizes the query.

*Returns*

On success, returns the `sizeof(a_v4_extfn_col_subset_of_input)`.

On failure, returns one of the generic `describe_column` errors, or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – set error returned if the buffer length is less than `sizeof(a_v4_extfn_col_subset_of_input)`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE\_VALUE** – set error returned if the column index of the source table is out of range.
- **EXTFNAPIV4\_DESCRIBE\_NOT\_AVAILABLE** – set error returned if the column `subset_of_input` is set on is not applicable (for example, if the column is not in the select list).
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – set error returned if the query processing state is not `Optimization`.
- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – set error returned if the buffer length is zero.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER** – set error returned if called on a parameter other than the return table.

*Query Processing States*

Valid in:

- Optimization state

*Example*

```
a_v4_extfn_col_subset_of_input colMap;

colMap.source_table_parameter_arg_num = 4;
colMap.source_column_number = i;

desc_rc = ctx->describe_column_set( ctx,
    0, i,
    EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT,
    &colMap, sizeof(a_v4_extfn_col_subset_of_input) );
```

**See also**

- *EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT (Get)* on page 220
- *Generic describe\_column Errors* on page 321
- *Query Processing States* on page 121

**\*describe\_parameter\_get**

The `describe_parameter_get` v4 API method gets UDF parameter properties from the server.

*Declaration*

```
a_sql_int32 (SQL_CALLBACK *describe_parameter_get)(
    a_v4_extfn_proc_context      *cntxt,
    a_sql_uint32                 arg_num,
    a_v4_extfn_describe_udf_type describe_type,
    const void                   *describe_buffer,
    size_t                       describe_buffer_len );
```

*Parameters*

Parameter	Description
<code>cntxt</code>	The procedure context object.
<code>arg_num</code>	The ordinal of the table parameter (0 is for the result table and 1 is for first input argument)
<code>describe_type</code>	A selector indicating what property to set.
<code>describe_buffer</code>	A structure that holds the describe information for the specified property to set on the server. The specific structure or data type is indicated by the <b>describe_type</b> parameter.
<code>describe_buffer_length</code>	The length, in bytes, of the <b>describe_buffer</b> .

*Returns*

On success, returns 0 or the number of bytes written to the **describe\_buffer**. A value of 0 indicates that the server was unable to get the attribute, but no error condition occurred. If an error occurred, or no property was retrieved, this function returns one of the generic **describe\_parameter** errors.

**Attributes for \*describe\_parameter\_get**

Code showing the attributes for `describe_parameter_get`.

```
typedef enum a_v4_extfn_describe_parm_type {
    EXTFNAPIV4_DESCRIBE_PARM_NAME,
    EXTFNAPIV4_DESCRIBE_PARM_TYPE,
```

```

EXTFNAPIV4_DESCRIBE_PARM_WIDTH,
EXTFNAPIV4_DESCRIBE_PARM_SCALE,
EXTFNAPIV4_DESCRIBE_PARM_CAN_BE_NULL,
EXTFNAPIV4_DESCRIBE_PARM_DISTINCT_VALUES,
EXTFNAPIV4_DESCRIBE_PARM_IS_CONSTANT,
EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE,

EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_ORDERBY,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_REQUEST_REWIND,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS,
} a_v4_extfn_describe_parm_type;

```

### EXTFNAPIV4\_DESCRIBE\_PARM\_NAME Attribute (Get)

The EXTFNAPIV4\_DESCRIBE\_PARM\_NAME attribute indicates the parameter name. Used in a describe\_parameter\_get scenario.

#### *Data Type*

char[ ]

#### *Description*

The name of a parameter to a UDF.

#### *Usage*

Gets the parameter name as defined in the **CREATE PROCEDURE** statement. Invalid for parameter 0.

#### *Returns*

On success, returns the length of the parameter name.

On failure, returns one of the generic describe\_parameter errors or:

- EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH – get error returned if the **describe\_buffer** is not large enough to hold the name.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE – get error returned if the state is not greater than Initial.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER – get error returned if the parameter is the result table.

#### *Query Processing States*

Valid in:

- Annotation state
- Query optimization state

## API Reference for a\_v4\_extfn

- Plan building state
- Execution state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_NAME Attribute (Set)* on page 258
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Get)

The **EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE** attribute returns the data type in a `describe_parameter_get` scenario.

#### *Data Type*

`a_sql_data_type`

#### *Description*

The data type of a parameter to a UDF.

#### *Usage*

Gets the data type of the parameter as defined in the **CREATE PROCEDURE** statement.

#### *Returns*

On success, returns `sizeof(a_sql_data_type)`.

On failure, returns one of the generic `describe_parameter` errors or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – get error returned if the **describe\_buffer** is not the `sizeof(a_sql_data_type)`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – get error returned if the state is not greater than Initial.

#### *Query Processing States*

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Set)* on page 259
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121



***EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH Attribute (Get)***

The **EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH** attribute indicates the width of a parameter. Used in a `describe_parameter_get` scenario.

***Data Type***

`a_sql_uint32`

***Description***

The width of a parameter to a UDF. **EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH** applies only to scalar parameters. Parameter width is the amount of storage, in bytes, required to store a parameter of the associated data type.

- **Fixed length data types** – the bytes required to store the data.
- **Variable length data types** – the maximum length.
- **LOB data types** – the amount of storage required to store a handle to the data.
- **TIME data types** – the amount of storage required to store the encoded time.

***Usage***

Gets the width of the parameter as defined in the **CREATE PROCEDURE** statement.

***Returns***

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_parameter` errors or:

- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – get error returned if the query processing state is not greater than Initial.
- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – get error returned if the **describe\_buffer** is not the size of `a_sql_uint32`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER** – get error returned if the specified parameter is a table parameter. This includes parameter 0, or parameter *n* where *n* is an input table.

***Query Processing States***

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

***Example***

Sample procedure definition:

```
CREATE PROCEDURE my_udf(IN p1 INT, IN p2 char(100))
RESULT (x INT)
EXTERNAL NAME 'my_udf@myudflib';
```

Sample `_describe_extfn` API function code fragment:

```
my_udf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_uint32 width = 0;
        a_sql_int32 ret = 0;

        // Get the width of parameter 1
        ret = cntxt->describe_parameter_get( cntxt, 1,
            EXTFNAPIV4_DESCRIBE_PARM_WIDTH,
            &width,
            sizeof(a_sql_uint32) );

        if( ret < 0 ) {
            // Handle the error.
        }

        //Allocate some storage based on parameter width
        a_sql_byte *p = (a_sql_byte *)cntxt->alloc( cntxt, width )

        // Get the width of parameter 2
        ret = cntxt->describe_parameter_get( cntxt, 2,
            EXTFNAPIV4_DESCRIBE_PARM_WIDTH,
            &width,
            sizeof(a_sql_uint32) );
        if( ret <= 0 ) {
            // Handle the error.
        }

        // Allocate some storage based on parameter width
        char *c = (char *)cntxt->alloc( cntxt, width )

        ...
    }
}
```

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH Attribute (Set)* on page 260
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE Attribute (Get)

The **EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE** attribute indicates the scale of a parameter. Used in a `describe_parameter_get` scenario.

### Data Type

`a_sql_uint32`

*Description*

The scale of a parameter to a UDF. For arithmetic data types, parameter scale is the number of digits to the right of the decimal point in a number.

This attribute is not valid for:

- non-arithmetic data types
- table parameters

*Usage*

Gets the scale of the parameter as defined in the **CREATE PROCEDURE** statement.

*Returns*

On success, returns the size of (a\_sql\_uint32).

On failure, returns one of the generic describe\_parameter errors or:

- EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH – get error returned if the **describe\_buffer** is not the size of a\_sql\_uint32.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE – get error returned if the state is not greater than Initial.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER – get error returned if the specified parameter is a table parameter. This includes parameter 0, or parameter *n* where *n* is an input table.

*Query Processing States*

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

*Example*

Sample \_describe\_extfn API function code fragment that gets the scale of parameter 1:

```
if( cntxt->current_state > EXTFNAPIV4_STATE_ANNOTATION ) {
    a_sql_uint32 scale = 0;
    a_sql_int32 ret = 0;

    ret = ctx->describe_parameter_get( ctx, 1,
EXTFNAPIV4_DESCRIBE_PARM_SCALE,
    &scale, sizeof(a_sql_uint32) );

    if( ret <= 0 ) {
        // Handle the error.
```

```
}
}
```

**See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE Attribute (Set)* on page 261
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_PARM\_CAN\_BE\_NULL Attribute (Get)**

The **EXTFNAPIV4\_DESCRIBE\_PARM\_CAN\_BE\_NULL** attribute indicates whether or not the parameter is null. Used in a `describe_parameter_get` scenario.

**Data Type**

a\_sql\_byte

**Description**

True, if the value of a parameter can be NULL at the time of execution. For a table parameter or parameter 0, the value is false.

**Usage**

Gets whether or not the specified parameter can be null during query execution.

**Returns**

On success, returns the `sizeof(a_sql_byte)`.

On failure, returns one of the generic `describe_parameter` errors or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – Get error returned if the **describe\_buffer** is not the size of `a_sql_byte`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – Get error returned if the state is not greater than Plan Building.

**Query Processing States**

Valid in:

- Execution state

**Examples: EXTFNAPIV4\_DESCRIBE\_PARM\_CAN\_BE\_NULL (Get)**

Example procedure definitions, `_describe_extfn` API function code fragment, and SQL queries for getting **EXTFNAPIV4\_DESCRIBE\_PARM\_CAN\_BE\_NULL** values.

**Procedure Definition**

Sample procedure definition used by the example queries in this topic:

```
CREATE PROCEDURE my_udf(IN p INT)
RESULT (x INT)
EXTERNAL NAME 'my_udf@myudflib';
```

### API Function Code Fragment

Sample `_describe_extfn` API function code fragment used by the example queries in this topic:

```

my_udf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state > EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_byte can_be_null = 0;
        a_sql_int32 ret = 0;

        ret = cntxt->describe_parameter_get( cntxt, 1,
            EXTFNAPIV4_DESCRIBE_PARM_CAN_BE_NULL,
            &can_be_null,
            sizeof(a_sql_byte) );

        if( ret <= 0 ) {
            // Handle the error.
        }
    }
}

```

#### Example 1: Without NOT NULL

This example creates a table with a single integer column without the **NOT NULL** modifier specified. The correlated subquery passes in column `c1` from the table `has_nulls`. When the procedure `my_udf_describe` is called during the Execution state, the call to `describe_parameter_get` populates `can_be_null` with a value of 1.

```

CREATE TABLE has_nulls ( c1 INT );
INSERT INTO has_nulls VALUES(1);
INSERT INTO has_nulls VALUES(NULL);
SELECT * from has_nulls WHERE (SELECT sum(my_udf.x) FROM
my_udf(has_nulls.c1)) > 0;

```

#### Example 2: With NOT NULL

This example creates a table with a single integer column with the **NOT NULL** modifier specified. The correlated subquery passes in column `c1` from the table `no_nulls`. When the procedure `my_udf_describe` is called during the Execution state, the call to `describe_parameter_get` populates `can_be_null` with a value of 0.

```

CREATE TABLE no_nulls ( c1 INT NOT NULL);
INSERT INTO no_nulls VALUES(1);
INSERT INTO no_nulls VALUES(2);
SELECT * from no_nulls WHERE (SELECT sum(my_udf.x) FROM
my_udf(no_nulls.c1)) > 0;

```

#### Example 3: With a Constant

This example calls the procedure `my_udf` with a constant. When the procedure `my_udf_describe` is called, during the Execution state, the call to `describe_parameter_get` populates `can_be_null` with a value of 0.

## API Reference for a\_v4\_extfn

```
SELECT * from my_udf(5);
```

### *Example 4: With a NULL*

This example calls the procedure **my\_udf** with a NULL. When the procedure **my\_udf\_describe** is called, during the Execution state, the call to **describe\_parameter\_get** populates **can\_be\_null** with a value of 1.

```
SELECT * from my_udf(NULL);
```

### EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES Attribute (Get)

The **EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES** attribute returns the number of distinct values. Used in a **describe\_parameter\_get** scenario.

### *Data Type*

a\_v4\_extfn\_estimate

### *Description*

Returns the estimated number of distinct values across all invocations. valid only for scalar parameters.

### *Usage*

If this information is available, the UDF returns the estimated number of distinct values with 100% confidence. If the information is not available, the UDF returns an estimate of 0 with 0% confidence.

### *Returns*

On success, returns the `sizeof(a_v4_extfn_estimate)`.

On failure, returns one of the generic **describe\_parameter** errors or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – get error returned if the **describe\_buffer** is not the size of **a\_v4\_extfn\_estimate**.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – get error returned if the state is not greater than Initial.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER** – get error returned if the parameter is a table parameter.

### *Query Processing States*

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### Example

Sample `_describe_extfn` API function code fragment:

```

if( ctx->current_state >= EXTFNAPIV4_STATE_ANNOTATION ) {
    desc_est.value = 0.0;
    desc_est.confidence = 0.0;

    desc_rc = ctx->describe_parameter_get( ctx,
        1,
        EXTFNAPIV4_DESCRIBE_PARM_DISTINCT_VALUES,
        &desc_est, sizeof(a_v4_extfn_estimate) );
}

```

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES Attribute (Set)* on page 263
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Get)* on page 240
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT Attribute (Get)

The `EXTFNAPIV4_DESCRIBE_PARM_DISTINCT_VALUES` attribute returns whether or not the parameter is constant. Used in a `describe_parameter_get` scenario.

### Data Type

`a_sql_byte`

### Description

True, if the parameter is a constant for the statement. Valid only for scalar parameters.

### Usage

Returns 0 if the value of the specified parameter is not a constant; returns 1 if the value of the specified parameter is a constant.

### Returns

On success, returns the `sizeof(a_sql_byte)`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the **describe\_buffer** is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned if the parameter is a table parameter.

### *Query Processing States*

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### *Example*

Sample `_describe_extfn` API function code fragment:

```
if( ctx->current_state >= EXTFNAPIV4_STATE_ANNOTATION ) {  
    desc_rc = ctx->describe_parameter_get( ctx,  
        1,  
        EXTFNAPIV4_DESCRIBE_PARM_IS_CONSTANT,  
        &desc_byte, sizeof( a_sql_byte ) );  
}
```

### **See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT Attribute (Set)* on page 263
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Set)* on page 259
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### *EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Get)*

The `EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE` attribute indicates the value of the parameter. Used in a `describe_parameter_get` scenario.

### *Data Type*

`an_extfn_value`

### *Description*

The value of the parameter if it is known at describe time. Valid only for scalar parameters.

### *Usage*

Returns the value of the parameters.

### *Returns*

On success, returns the `sizeof( an_extfn_value )` if the value is available, or:

- `EXTFNAPIV4_DESCRIBE_NOT_AVILABLE` – Value returned if the value is not constant.

On failure, returns one of the generic `describe_parameter` errors or:



- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the **describe\_buffer** is not the size of `an_extfn_value`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned if the parameter is a table parameter.

### Query Processing States

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### Example

Sample `_describe_extfn` API function code fragment:

```
if( ctx->current_state >= EXTFNAPIV4_STATE_ANNOTATION ) {
    a_sql_int32 desc_rc;
    desc_rc = ctx->describe_parameter_get( ctx,
        1,
        EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE,
        &arg,
        sizeof( an_extfn_value ) );
}
```

### See also

- [EXTFNAPIV4\\_DESCRIBE\\_PARM\\_IS\\_CONSTANT Attribute \(Set\)](#) on page 263
- [EXTFNAPIV4\\_DESCRIBE\\_PARM\\_TYPE Attribute \(Get\)](#) on page 240
- [Generic describe\\_parameter Errors](#) on page 322
- [Query Processing States](#) on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS Attribute (Get)

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS` attribute indicates the number of columns in the table. Used in a `describe_parameter_get` scenario.

### Data Type

`a_sql_uint32`

### Description

The number of columns in the table. Only valid for argument 0 and table arguments.

### Usage

Returns the number of columns in the specified table argument. Argument 0 returns the number of columns in the result table.

### Returns

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the **describe\_buffer** is not the size of `sizeof(a_sql_uint32)`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – get error returned if the parameter is not a table parameter.

### Query Processing States

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### See also

- *[EXTFNAPIV4\\_DESCRIBE\\_PARM\\_TABLE\\_NUM\\_COLUMNS Attribute \(Set\)](#)* on page 264
- *[Query Processing States](#)* on page 121

### *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS Attribute (Get)*

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS` attribute indicates the number of rows in the table. Used in a `describe_parameter_get` scenario.

### Data Type

`a_v4_extfn_estimate`

### Description

The estimated number of rows in the table. Only valid for argument 0 and table arguments.

### Usage

Returns the estimated number of rows in the specified table argument or result set with a confidence of 100%.

### Returns

On success, returns the size of `a_v4_extfn_estimate`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the **describe\_buffer** is not the size of `a_v4_extfn_estimate`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – get error returned if the parameter is not a table parameter.

### *Query Processing States*

Valid in:

- Annotation state
- Query Optimization state
- Plan Building state
- Execution state

### **See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS Attribute (Set)* on page 265
- *Query Processing States* on page 121

### *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Get)*

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_ORDERBY` attribute indicates the order of rows in the table. Used in a `describe_parameter_get` scenario.

### *Data Type*

`a_v4_extfn_orderby_list`

### *Description*

The order of rows in the table. This property is only valid for argument 0 and table arguments.

### *Usage*

This attribute allows the UDF code to:

- Determine if the input **TABLE** parameter has been ordered
- Declare that the result set is ordered

If the parameter number is 0, then the attribute refers to the outbound result set. If the parameter is  $> 0$  and the parameter type is a table then the attribute refers to the input **TABLE** parameter.

The order is specified by the `a_v4_extfn_orderby_list`, which is a structure supporting a list of column ordinals and their associated ascending or descending property. If the UDF sets the order by property for the outbound result set, the server is then able to perform order by optimizations. For example, if the UDF produced ascending order on the

first result set column, the server will eliminate a redundant order by request on the same column.

If the UDF does not set the orderby property on the outbound result set, the server assumes the data is not ordered.

If the UDF sets the orderby property on the input **TABLE** parameter, the server guarantees data ordering for the input data. In this scenario, the UDF describes to the server that the input data must be ordered. If the server detects a runtime conflict it raises a SQL exception. For example, when the UDF describes that the first column of the input **TABLE** parameter must have ascending order and the SQL statement contains a descending clause, the server raises a SQL exception.

In the event that the SQL did not contain an ordering clause, the server automatically adds the ordering to ensure that input **TABLE** parameter is ordered as required.

### Returns

If successful, returns the number of bytes copied from `a_v4_extfn_orderby_list`.

### Query Processing States

Valid in:

- Annotation state
- Query optimization state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Set)* on page 266
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Get)

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY` attribute indicates that the UDF requires partitioning. Used in a `describe_parameter_get` scenario.

### Data Type

`a_v4_extfn_column_list`

### Description

UDF developers use `EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY` to programmatically declare that the UDF requires partitioning before invocation can proceed.

### Usage

The UDF can inquire to the partition to enforce it, or to dynamically adapt the partitioning. It is the UDF's responsibility to allocate the `a_v4_extfn_column_list`, taking into consideration the total number of columns in the input table, and sending that data to the server.

### Returns

On success, returns the size of `a_v4_extfn_column_list`. This value is equal to:

```
sizeof(a_v4_extfn_column_list) + sizeof(a_sql_uint32) *
number_of_partition_columns
```

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the buffer length is less than the expected size.

### Query Processing States

Valid in:

- Query Optimization state
- Plan Building state
- Execution state

### Example

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
    if( ctx->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_uint32    col_count    = 0;
        a_sql_uint32    buffer_size  = 0;
        a_v4_extfn_column_list *clist = NULL;

        col_count = 3;    // Set to the max number of possible pby
columns

        buffer_size = sizeof( a_v4_extfn_column_list ) + (col_count -
1) * sizeof( a_sql_uint32 );

        clist = (a_v4_extfn_column_list *)ctx->alloc( ctx,
buffer_size );

        clist->number_of_columns = 0;
        clist->column_indexes[0] = 0;
        clist->column_indexes[1] = 0;
        clist->column_indexes[2] = 0;

        args->r_api_rc = ctx->describe_parameter_get( ctx,
args->p3_arg_num,
EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
clist,
buffer_size );
    }
}
```

### See also

- `EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY (Set)` on page 268

- *Generic describe\_parameter Errors* on page 322
- *V4 API describe\_parameter and EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 141
- *Parallel TPF PARTITION BY Examples Using EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 143
- *Query Processing States* on page 121
- *Partitioning Input Data* on page 140

*EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Get)*

The EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND attribute indicates that the consumer requests rewind of an input table. Used in a describe\_parameter\_get scenario.

*Data Type*

a\_sql\_byte

*Description*

Indicates that the consumer wants to rewind an input table. Valid only for table input arguments. By default, this property is false.

*Usage*

The UDF queries this property to retrieve the true/false value.

*Returns*

On success, returns sizeof(a\_sql\_byte).

On failure, returns one of the generic describe\_parameter errors, or:

- EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH – get error returned if the **describe\_buffer** is not the size of a\_sql\_byte.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE – get error returned if the state is not Optimization or Plan Building.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER – get error returned if the UDF attempts to get this attribute on parameter 0.
- EXTFNAPIV4\_DESCRIBE\_NON\_TABLE\_PARAMETER – get error returned if the UDF attempts to get this attribute on a parameter that is not a table.

*Query Processing States*

Valid in:

- Optimization state
- Plan Building state

**See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)* on page 269
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Set)* on page 271
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Get)* on page 255
- *\_rewind\_extfn* on page 318
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Get)**

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND` attribute indicates if the parameter supports rewind. Used in a `describe_parameter_get` scenario.

**Data Type**

`a_sql_byte`

**Description**

Indicates whether a producer can support rewind. Valid only for table arguments.

You must also provide an implementation of the rewind table callback (`_rewind_extfn()`) if you plan on setting `DESCRIBE_PARM_TABLE_HAS_REWIND` to `true`. The server will fail to execute the UDF if the callback method is not provided.

**Usage**

The UDF asks if a table input argument supports rewind. As a prerequisite, the UDF must request rewind using **DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND** before you can use this property.

**Returns**

On success, returns `sizeof(a_sql_byte)`.

On failure, returns one of the generic `describe_parameter` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the **describe\_buffer** is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error error returned if the state is not greater than Annotation.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – get error returned if the UDF attempts to get this attribute on a parameter that is not a table.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned if the UDF attempts to get this attribute on the result table.

**Query Processing States**

Valid in:

- Optimization state
- Plan Building state
- Execution state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Get)* on page 254
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)* on page 269
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Set)* on page 271
- *\_rewind\_extfn* on page 318
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_UNUSED\_COLUMNS Attribute (Get)

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS` attribute lists unconsumed columns. Used in a `describe_parameter_get` scenario.

### Data Type

`a_v4_extfn_column_list`

### Description

The list of output table columns that are not going to be consumed by the server or the UDF.

For the output table parameter, the UDF normally produces the data for all the columns, and the server consumes all the columns. The same holds true for the input table parameter where the server normally produces the data for all the columns, and the UDF consumes all the columns.

However, in some cases the server, or the UDF, may not consume all the columns. The best practice in such a case is for the UDF to perform a **GET** for the output table on the describe attribute `EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS`. This action queries the server for the list of output table columns that are not going to be consumed by the server. The list can then be used by the UDF when populating the column data for the output table; that is, the UDF does not attempt to populate data for unused columns.

In summary, for the output table the UDF polls the list of unused columns. For the input table, the UDF pushes the list of unused columns.

### Usage

The UDF asks the server if all the columns of the output table are going to be used. The UDF must allocate a `a_v4_extfn_column_list` that includes all the columns of the output table, and then must pass it to the server. The server then marks all the unprojected column ordinals as 1. The list returned by the server can be used while producing the data.



### Returns

On success, returns the size of the column list: `sizeof(a_v4_extfn_column_list) + sizeof(a_sql_uint32) * number result columns`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Plan Building.
- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the **describe\_buffer** is not large enough to hold the returned list.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – get error returned if the UDF attempts to get this attribute on an input table.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – get error returned if the UDF attempts to get this attribute on a parameter that is not a table.

### Query Processing States

Valid in:

- Execution state

### See also

- `EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS` Attribute (Set) on page 272

## \*describe\_parameter\_set

The `describe_parameter_set` v4 API method sets properties about a single parameter to the UDF.

### Declaration

```
a_sql_int32 (SQL_CALLBACK *describe_parameter_set)(
    a_v4_extfn_proc_context      *cntxt,
    a_sql_uint32                 arg_num,
    a_v4_extfn_describe_udf_type describe_type,
    const void                   *describe_buffer,
    size_t                       describe_buffer_len );
```

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object.
<code>arg_num</code>	The ordinal of the table parameter (0 is for the result table and 1 is for first input argument)

Parameter	Description
<b>describe_type</b>	A selector indicating what property to set.
<b>describe_buffer</b>	A structure that holds the describe information for the specified property to set on the server. The specific structure or data type is indicated by the <b>describe_type</b> parameter.
<b>describe_buffer_length</b>	The length in bytes of the <b>describe_buffer</b> .

*Returns*

On success, returns 0 or the number of bytes written to the **describe\_buffer**. A value of 0 indicates that the server was unable to set the attribute, but no error condition occurred. If an error occurred, or no property was retrieved, this function returns one of the generic **describe\_parameter** errors.

**Attributes for \*describe\_parameter\_set**

Code showing the attributes for `describe_parameter_set`.

```
typedef enum a_v4_extfn_describe_parm_type {
    EXTFNAPIV4_DESCRIBE_PARM_NAME,
    EXTFNAPIV4_DESCRIBE_PARM_TYPE,
    EXTFNAPIV4_DESCRIBE_PARM_WIDTH,
    EXTFNAPIV4_DESCRIBE_PARM_SCALE,
    EXTFNAPIV4_DESCRIBE_PARM_CAN_BE_NULL,
    EXTFNAPIV4_DESCRIBE_PARM_DISTINCT_VALUES,
    EXTFNAPIV4_DESCRIBE_PARM_IS_CONSTANT,
    EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE,

    EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_ORDERBY,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_REQUEST_REWIND,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS,

} a_v4_extfn_describe_parm_type;
```

**EXTFNAPIV4\_DESCRIBE\_PARM\_NAME Attribute (Set)**

The `EXTFNAPIV4_DESCRIBE_PARM_NAME` attribute indicates the parameter name. Used in a `describe_parameter_set` scenario.

*Data Type*

char [ ]

*Description*

The name of a parameter to a UDF.

### Usage

If the UDF sets this property, the server compares the value with the name of the parameter supplied in the **CREATE PROCEDURE** statement. If the two values do not match, the server returns an error. This allows the UDF to ensure the **CREATE PROCEDURE** statement has the same parameter names as the UDF is expecting.

### Returns

On success, returns the length of the parameter name.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not equal to Annotation.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the parameter is the result table.
- `EXTFNAPI4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF tries to reset the name.

### Query Processing States

Valid in:

- Annotation state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_NAME Attribute (Get)* on page 239
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Set)

The `EXTFNAPIV4_DESCRIBE_PARM_TYPE` attribute indicates the data type of the parameter. Used in a `describe_parameter_set` scenario.

### Data Type

`a_sql_data_type`

### Description

The data type of a parameter to a UDF.

### Usage

When the UDF sets this property, the server compares the value to the parameter type supplied in the **CREATE PROCEDURE** statement. If the two values do not match, the server returns an error. This check ensures that the **CREATE PROCEDURE** statement has the same parameter data types that the UDF expects.

### Returns

On success, returns `sizeof(a_sql_data_type)`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the **describe\_buffer** is not the `sizeof(a_sql_data_type)`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the query processing state is not equal to Annotation.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF tries to set the datatype of a parameter to something other than what it is already defined as.

### Query Processing States

Valid in:

- Annotation state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Get)* on page 240
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH Attribute (Set)

The `EXTFNAPIV4_DESCRIBE_PARM_WIDTH` attribute indicates the width of a parameter. Used in a `describe_parameter_set` scenario.

### Data Type

`a_sql_uint32`

### Description

The width of a parameter to a UDF. `EXTFNAPIV4_DESCRIBE_PARM_WIDTH` applies only to scalar parameters. Parameter width is the amount of storage, in bytes, required to store a parameter of the associated data type.

- **Fixed length data types** – the bytes required to store the data.
- **Variable length data types** – the maximum length.
- **LOB data types** – the amount of storage required to store a handle to the data.
- **TIME data types** – the amount of storage required to store the encoded time.

### Usage

This is a read-only property. The width is derived from the associated column data type. Once the data type is set, you cannot change the width.

### Returns

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the query processing state is not equal to Annotation.
- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the **describe\_buffer** is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the specified parameter is a table parameter. This includes parameter 0, or parameter *n*, where *n* is an input table.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF tries to reset the parameter width.

### Query Processing States

Valid in:

- Annotation state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH Attribute (Get)* on page 241
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE Attribute (Set)

The **EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE** attribute indicates the scale of a parameter. Used in a `describe_parameter_set` scenario.

### Data Type

`a_sql_uint32`

### Description

The scale of a parameter to a UDF. For arithmetic data types, parameter scale is the number of digits to the right of the decimal point in a number.

This attribute is invalid for:

- Nonarithmetic data types
- Table parameters

### Usage

This is a read-only property. The scale is derived from the associated column data type. Once the data type is set, you cannot change the scale.

### Returns

On success, returns `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the **describe\_buffer** is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not Annotation.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the specified parameter is a table parameter. This includes parameter 0, or parameter *n*, where *n* is an input table.

### Query Processing States

Valid in:

- Annotation state

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE Attribute (Get)* on page 242
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_CAN\_BE\_NULL Attribute (Set)

The `EXTFNAPIV4_DESCRIBE_PARM_CAN_BE_NULL` attribute returns whether or not the parameter is null. Using this attribute in a `describe_parameter_set` scenario returns an error.

### Data Type

`a_sql_byte`

### Description

True, if the value of a parameter can be NULL at the time of execution. For a table parameter or parameter 0, the value is false.

### Usage

This is a read-only property.

### Returns

This is a read-only property, so all attempts to set result in an `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE` error.

### Query Processing States

Not applicable.

**EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES Attribute (Set)**

The **EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES** attribute returns the number of distinct values. Using this attribute in a `describe_parameter_set` scenario returns an error.

***Data Type***

`a_v4_extfn_estimate`

***Description***

Returns the estimated number of distinct values across all invocations. valid only for scalar parameters.

***Usage***

This is a read-only property.

***Returns***

This is a read-only property; all attempts to set result in an `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE` error.

***Query Processing States***

Not applicable.

**See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES Attribute (Get)* on page 246
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Get)* on page 240
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT Attribute (Set)**

The **EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT** attribute returns whether or not the parameter is constant. Using this attribute in a `describe_parameter_set` scenario returns an error.

***Data Type***

`a_sql_byte`

***Description***

True, if the parameter is a constant for the statement. Valid only for scalar parameters.

***Usage***

This is a read-only property.

### Returns

This is a read-only property; all attempts to set result in an `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE` error.

### Query Processing States

Not applicable.

### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT Attribute (Get)* on page 247
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Set)* on page 259
- *Generic describe\_parameter Errors* on page 322
- *Query Processing States* on page 121
- *EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Get)* on page 248
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE Attribute (Get)* on page 240

### EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE Attribute (Set)

The `EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE` attribute indicates the value of the parameter. Used in a `describe_parameter_set` scenario.

### Data Type

`an_extfn_value`

### Description

The value of the parameter if it is known at describe time. Valid only for scalar parameters.

### Usage

This is a read-only property.

### Returns

This is a read-only property; all attempts to set result in an `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE` error.

### Query Processing States

Not applicable.

### EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS Attribute (Set)

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS` attribute indicates the number of columns in the table. Used in a `describe_parameter_set` scenario.

### Data Type

`a_sql_uint32`



*Description*

The number of columns in the table. Only valid for argument 0 and table arguments.

*Usage*

If the UDF sets this property, the server compares the value with the name of the parameter supplied in the **CREATE PROCEDURE** statement. If the two values do not match, the server returns an error. This allows the UDF to ensure the **CREATE PROCEDURE** statement has the same parameter names as the UDF is expecting.

*Returns*

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_parameter` errors or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the **describe\_buffer** is not the size of `sizeof a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not `ANNOTATION`.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – set error returned if the parameter is not a table parameter.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF tries to reset the number of columns of the specified table.

*Query Processing States*

Valid in:

- Annotation state

**See also**

- *`EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS` Attribute (Get)* on page 249
- *Query Processing States* on page 121

*`EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS` Attribute (Set)*

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS` attribute indicates the number of rows in the table. Used in a `describe_parameter_set` scenario.

*Data Type*

`a_sql_a_v4_extfn_estimate`

*Description*

The estimated number of rows in the table. Only valid for argument 0 and table arguments.

### *Usage*

The UDF sets this property for argument 0 if it estimates the number of rows in the result set. The server uses the estimate during optimization to make query processing decisions. You cannot set this value for an input table.

If you do not set a value, the server defaults to the number of rows specified by the **DEFAULT\_TABLE\_UDF\_ROW\_COUNT** option.

### *Returns*

On success, returns `a_v4_extfn_estimate`.

On failure, returns one of the generic `describe_parameter` errors or:

- **EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH** – set error returned if the **describe\_buffer** is not the size of `a_v4_extfn_estimate`.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE** – set error returned if the state is not Optimization.
- **EXTFNAPIV4\_DESCRIBE\_NON\_TABLE\_PARAMETER** – get error returned if the parameter is not a table parameter.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_PARAMETER** – get error returned if the table parameter is not the result table.
- **EXTFNAPIV4\_DESCRIBE\_INVALID\_ATTRIBUTE\_VALUE** – get error returned if the UDF tries to reset the number of columns of the specified table.

### *Query Processing States*

Valid in:

- Query Optimization state

### **See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS Attribute (Get)* on page 250
- *Query Processing States* on page 121

### *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Set)*

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_ORDERBY` attribute indicates the order of rows in the table. Used in a `describe_parameter_set` scenario.

### *Data Type*

`a_v4_extfn_orderby_list`

### *Description*

The order of rows in the table. This property is only valid for argument 0 and table arguments.

### *Usage*

This attribute allows the UDF code to:

- Determine if the input **TABLE** parameter has been ordered
- Declare that the result set is ordered.

If the parameter number is 0, then the attribute refers to the outbound result set. If the parameter is > 0 and the parameter type is a table then the attribute refers to the input **TABLE** parameter.

The order is specified by the `a_v4_extfn_orderby_list`, which is a structure supporting a list of column ordinals and their associated ascending or descending property. If the UDF sets the order by property for the outbound result set, the server is then able to perform order by optimizations. For example, if the UDF produced ascending order on the first result set column, the server will eliminate a redundant order by request on the same column.

If the UDF does not set the orderby property on the outbound result set, the server assumes the data is not ordered.

If the UDF sets the orderby property on the input **TABLE** parameter, the server guarantees data ordering for the input data. In this scenario, the UDF describes to the server that the input data must be ordered. If the server detects a runtime conflict it raises a SQL exception. For example, when the UDF describes that the first column of the input **TABLE** parameter must have ascending order and the SQL statement contains a descending clause, the server raises a SQL exception.

In the event that the SQL did not contain an ordering clause, the server automatically adds the ordering to ensure that input **TABLE** parameter is ordered as required.

### *Returns*

If successful, returns the number of bytes copied from `a_v4_extfn_orderby_list`.

### *Query Processing States*

Valid in:

- Annotation state
- Query optimization state

### **See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_ORDERBY Attribute (Get)* on page 251
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Set)**

The EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY attribute indicates that the UDF requires partitioning. Used in a describe\_parameter\_set scenario.

**Data Type**

a\_v4\_extfn\_column\_list

**Description**

UDF developers use **EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY** to programmatically declare that the UDF requires partitioning before invocation can proceed.

**Usage**

The UDF can inquire to the partition to enforce it, or to dynamically adapt the partitioning. The UDF must allocate the **a\_v4\_extfn\_column\_list**, taking into consideration the total number of columns in the input table, and sending that data to the server.

**Returns**

On success, returns the size of a\_v4\_extfn\_column\_list. This value is equal to:

```
sizeof(a_v4_extfn_column_list) + sizeof(a_sql_uint32) *
number_of_partition_columns
```

On failure, returns one of the generic describe\_parameter errors or:

- EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH – Set error returned if the buffer length is less than the expected size.

**Query Processing States**

Valid in:

- Annotation state
- Query Optimization state

**Example**

```
void UDF_CALLBACK my_tpf_proc_describe( a_v4_extfn_proc_context
*ctx )
{
    if( ctx->current_state == EXTFNAPIV4_STATE_ANNOTATION ) {
        a_sql_int32 rc = 0;
        a_v4_extfn_column_list pbcoll =
        { 1,          // 1 column in the partition by list
          2 };      // column index 2 requires partitioning

        // Describe partitioning for argument 1 (the table)
        rc = ctx->describe_parameter_set(
            ctx, 1,
            EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
            &pbcoll,
            sizeof(pbcoll) );
    }
}
```

```

        if( rc == 0 ) {
            ctx->set_error( ctx, 17000,
                "Runtime error, unable set partitioning requirements for
column." );
        }
    }
}

```

**See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY (Get)* on page 252
- *Generic describe\_parameter Errors* on page 322
- *V4 API describe\_parameter and EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 141
- *Parallel TPF PARTITION BY Examples Using EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY* on page 143
- *Query Processing States* on page 121
- *Partitioning Input Data* on page 140

**EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)**

The EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND attribute indicates that the consumer requests rewind of an input table. Used in a describe\_parameter\_set scenario.

**Data Type**

a\_sql\_byte

**Description**

Indicates that the consumer wants to rewind an input table. Valid only for table input arguments. By default, this property is false.

**Usage**

If the UDF requires input table rewind capability, the UDF must set this property during Optimization.

**Returns**

On success, returns sizeof(a\_sql\_byte).

On failure, returns one of the generic describe\_parameter errors, or:

- EXTFNAPIV4\_DESCRIBE\_BUFFER\_SIZE\_MISMATCH – set error returned if the **describe\_buffer** is not the size of a\_sql\_byte.
- EXTFNAPIV4\_DESCRIBE\_INVALID\_STATE – set error returned if the state is not equal to Optimization.

## API Reference for a\_v4\_extfn

- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the UDF attempts to set this attribute on parameter 0.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – set error returned if the UDF attempts to set this attribute on a parameter that is not a table.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF attempts to set this attribute to a value other than 0 or 1.

### Query Processing States

Valid in:

- Optimization state

### Example

In this example, when the function `my_udf_describe` is called during the Optimization state, the call to `describe_parameter_set` informs the producer of the table input parameter 1 that a rewind may be required.

Sample procedure definition:

```
CREATE PROCEDURE my_udf(IN t TABLE(c1 INT))
RESULT (x INT)
EXTERNAL NAME 'my_udf@myudflib';
```

Sample `_describe_extfn` API function code fragment:

```
my_udf_describe(a_v4_extfn_proc_context *cntxt)
{
    if( cntxt->current_state == EXTFNAPIV4_STATE_OPTIMIZATION ) {
        a_sql_byte rewind_required = 1;
        a_sql_int32 ret = 0;

        ret = cntxt->describe_parameter_set( cntxt, 1,
            EXTFNAPIV4_DESCRIBE_PARM_TABLE_REQUEST_REWIND,
            &rewind_required,
            sizeof(a_sql_byte) );

        if( ret <= 0 ) {
            // Handle the error.
        }
    }
}
```

### See also

- `EXTFNAPIV4_DESCRIBE_PARM_TABLE_REQUEST_REWIND` Attribute (Get) on page 254
- `EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND` Attribute (Set) on page 271
- `EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND` Attribute (Get) on page 255
- `_rewind_extfn` on page 318

- *Query Processing States* on page 121

### EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Set)

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND` attribute indicates if the parameter supports rewind. Used in a `describe_parameter_set` scenario.

#### *Data Type*

`a_sql_byte`

#### *Description*

Indicates whether a producer can support rewind. Valid only for table arguments.

You must also provide an implementation of the rewind table callback (`_rewind_extfn()`), if you plan on setting `DESCRIBE_PARM_TABLE_HAS_REWIND` to true. The server cannot execute the UDF if you do not provide the callback method.

#### *Usage*

A UDF sets this property during the Optimization state if it can provide rewind capability for its result table at no cost. If it is expensive for the UDF to provide rewind, do not set this property, or set it to 0. If set to 0, the server provides rewind support.

#### *Returns*

On success, returns `sizeof(a_sql_byte)`.

On failure, returns one of the generic `describe_parameter` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if the **describe\_buffer** is not the size of `a_sql_byte`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not equal to Optimization.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – set error returned if the UDF attempts to set this attribute on a parameter that is not a table.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the specified argument is not the result table.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF attempts to set this attribute to a value other than 0 or 1.

#### *Query Processing States*

Valid in:

- Optimization state

#### **See also**

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Get)* on page 254

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)* on page 269
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Get)* on page 255
- *\_rewind\_extfn* on page 318
- *Query Processing States* on page 121

**EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_UNUSED\_COLUMNS Attribute (Set)**

The `EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS` attribute lists unconsumed columns. Used in a `describe_parameter_set` scenario.

**Data Type**

`a_v4_extfn_column_list`

**Description**

The list of output table columns that are not going to be consumed by the server or the UDF.

For the output table parameter, the UDF normally produces the data for all the columns, and the server consumes all the columns. The same holds true for the input table parameter where the server normally produces the data for all the columns, and the UDF consumes all the columns.

However, in some cases the server, or the UDF, may not consume all the columns. The best practice in such a case is that the UDF performs a **GET** for the output table on the describe attribute **EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_UNUSED\_COLUMNS**. This action queries the server for the list of output table columns which are not going to be consumed by the server. The list can then be used by the UDF when populating the column data for the output table; that is, the UDF skips populating data for unused columns.

In summary, for the output table the UDF polls the list of unused columns. For the input table, the UDF pushes the list of unused columns.

**Usage**

The UDF sets this property during Optimization if it is not going to use certain columns of the input table parameter. The UDF must allocate a `a_v4_extfn_column_list` that includes all the columns of the output table, and then must pass it to the server. The server then marks all the un-projected column ordinals as 1. The server copies the list into its internal data structure.

**Returns**

On success, returns the size of the column list: `sizeof(a_v4_extfn_column_list) + sizeof(a_sql_uint32) * number result columns`.

On failure, returns one of the generic `describe_parameter` errors or:



- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – set error returned if the state is not Optimization.
- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if the UDF attempts to get this attribute on an input table.
- `EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER` – set error returned if the UDF attempts to set this attribute on a parameter that is not a table.

### Query Processing States

Valid in:

- Optimization state

### See also

- *`EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS` Attribute (Get)* on page 256

## \*describe\_udf\_get

The `describe_udf_get` v4 API method gets UDF properties from the server.

### Declaration

```
a_sql_int32 (SQL_CALLBACK *describe_udf_get)(
    a_v4_extfn_proc_context *cntxt,
    a_v4_extfn_describe_udf_type describe_type,
    void *describe_buffer,
    size_t describe_buffer_len );
```

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object for this UDF.
<code>describe_type</code>	A selector indicating what property to retrieve.
<code>describe_buffer</code>	A structure that holds the describe information for the specified property to set on the server. The specific structure or data type is indicated by the <code>describe_type</code> parameter.
<code>describe_buffer_length</code>	The length in bytes of the <code>describe_buffer</code> .

### Returns

On success, returns 0 or the number of bytes written to the `describe_buffer`. A value of 0 indicates that the server was unable to get the attribute but no error condition occurred. If an error occurred, or no property was retrieved, this function returns one of the generic `describe_udf` errors.

### See also

- *\*describe\_udf\_set* on page 275
- *Generic describe\_udf Errors* on page 322

### Attributes for \*describe\_udf\_get

Code showing the attributes for `describe_udf_get`.

```
typedef enum a_v4_extfn_describe_udf_type {  
    EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS,  
    EXTFNAPIV4_DESCRIBE_UDF_LAST  
} a_v4_extGetfn_describe_udf_type;
```

### EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS Attribute (Get)

The `EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS` attribute indicates the number of parameters. Used in a `describe_udf_get` scenario.

### Data Type

`a_sql_uint32`

### Description

The number of parameters supplied to the UDF.

### Usage

Gets the number of parameters as defined in the **CREATE PROCEDURE** statement.

### Returns

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_udf` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – get error returned if the describe buffer is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – get error returned if the state is not greater than Initial.

### Query Processing States

- Annotation state
- Query optimization state
- Plan building state
- Execution state

### See also

- *EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS Attribute (Set)* on page 276
- *Generic describe\_udf Errors* on page 322
- *Query Processing States* on page 121

## **\*describe\_udf\_set**

The `describe_udf_set` v4 API method sets UDF properties on the server.

### *Declaration*

```
a_sql_int32 (SQL_CALLBACK *describe_udf_set) (
    a_v4_extfn_proc_context *cntxt,
    a_v4_extfn_describe_udf_type describe_type,
    const void *describe_buffer,
    size_t describe_buffer_len );
```

### *Parameters*

Parameter	Description
<code>cntxt</code>	The procedure context object for this UDF.
<code>describe_type</code>	A selector indicating what property to set.
<code>describe_buffer</code>	A structure that holds the describe information for the specified property to set on the server. The specific structure or data-type is indicated by the <code>describe_type</code> parameter.
<code>describe_buffer_length</code>	The length, in bytes, of <code>describe_buffer</code> .

### *Returns*

On success, returns the number of bytes written to the `describe_buffer`. If an error occurs, or no property is retrieved, this function returns one of the generic `describe_udf` errors.

If an error occurs, or no property is retrieved, this function returns one of the generic `describe_udf` errors, or:

- `EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER` – set error returned if any of the `cntxt` or `describe_buffer` arguments are NULL or if `describe_buffer_length` is 0.
- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – set error returned if there is a discrepancy between the requested attribute's size and the supplied `describe_buffer_length`.

### **See also**

- *\*describe\_udf\_get* on page 273
- *Generic describe\_udf Errors* on page 322

### **Attributes for \*describe\_udf\_set**

Code showing the attributes for `describe_udf_set`.

```
typedef enum a_v4_extfn_describe_udf_type {
    EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS,
```

```
EXTFNAPIV4_DESCRIBE_UDF_LAST  
} a_v4_extGetfn_describe_udf_type;
```

### *EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS Attribute (Set)*

The `EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS` attribute indicates the number of parameters. Used in a `describe_udf_set` scenario.

#### *Data Type*

`a_sql_uint32`

#### *Description*

The number of parameters supplied to the UDF.

#### *Usage*

If the UDF sets this property, the server compares the value with the number of parameters supplied in the **CREATE PROCEDURE** statement. If the two values do not match, the server returns a SQL error. This allows the UDF to ensure the **CREATE PROCEDURE** statement has the same number of parameters expected by the UDF.

#### *Returns*

On success, returns the `sizeof(a_sql_uint32)`.

On failure, returns one of the generic `describe_udf` errors, or:

- `EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH` – Set error returned if the describe buffer is not the size of `a_sql_uint32`.
- `EXTFNAPIV4_DESCRIBE_INVALID_STATE` – Set error returned if the state is not equal to Annotation.
- `EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE` – set error returned if the UDF tries to reset the parameter datatype.

#### *Query processing states*

- Annotation state

#### **See also**

- *EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARMS Attribute (Get)* on page 274
- *Generic describe\_udf Errors* on page 322
- *Query Processing States* on page 121

## Describe Column Type (`a_v4_extfn_describe_col_type`)

The `a_v4_extfn_describe_col_type` enumerated type selects the column property retrieved or set by the UDF.

### Implementation

```
typedef enum a_v4_extfn_describe_col_type {
    EXTFNAPIV4_DESCRIBE_COL_NAME,
    EXTFNAPIV4_DESCRIBE_COL_TYPE,
    EXTFNAPIV4_DESCRIBE_COL_WIDTH,
    EXTFNAPIV4_DESCRIBE_COL_SCALE,
    EXTFNAPIV4_DESCRIBE_COL_CAN_BE_NULL,
    EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES,
    EXTFNAPIV4_DESCRIBE_COL_IS_UNIQUE,
    EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT,
    EXTFNAPIV4_DESCRIBE_COL_CONSTANT_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER,
    EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE,
    EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT,
    EXTFNAPIV4_DESCRIBE_COL_LAST
} a_v4_extfn_describe_col_type;
```

### Members Summary

Member	Description
<i>EXTFNAPIV4_DESCRIBE_COL_NAME</i>	Column name (valid identifier).
<i>EXTFNAPIV4_DESCRIBE_COL_TYPE</i>	Column data type.
<i>EXTFNAPIV4_DESCRIBE_COL_WIDTH</i>	String width (precision for NUMERIC).
<i>EXTFNAPIV4_DESCRIBE_COL_SCALE</i>	Scale for NUMERIC.
<i>EXTFNAPIV4_DESCRIBE_COL_CAN_BE_NULL</i>	True, if a column can be NULL.
<i>EXTFNAPIV4_DESCRIBE_COL_DISTINCT_VALUES</i>	Estimated number of distinct values in the column.
<i>EXTFNAPIV4_DESCRIBE_COL_IS_UNIQUE</i>	True, if column is unique within the table.
<i>EXTFNAPIV4_DESCRIBE_COL_IS_CONSTANT</i>	True, if column is constant for statement lifetime.
<i>EXTFNAPIV4_DESCRIBE_COL_CONSTANT_VALUE</i>	The value of a parameter, if known at describe time.

Member	Description
<i>EXTFNAPIV4_DESCRIBE_COL_IS_USED_BY_CONSUMER</i>	True, if column is needed by the consumer of the table.
<i>EXTFNAPIV4_DESCRIBE_COL_MINIMUM_VALUE</i>	The minimum value for the column (if known).
<i>EXTFNAPIV4_DESCRIBE_COL_MAXIMUM_VALUE</i>	The maximum value for the column (if known).
<i>EXTFNAPIV4_DESCRIBE_COL_VALUES_SUBSET_OF_INPUT</i>	The result column values are a subset of columns from an input table.
<i>EXTFNAPIV4_DESCRIBE_COL_LAST</i>	First illegal value for v4 API. Out-of-band value.

## Describe Parameter Type (a\_v4\_extfn\_describe\_parm\_type)

The `a_v4_extfn_describe_parm_type` enumerated type selects the parameter property retrieved or set by the UDF.

### Implementation

```
typedef enum a_v4_extfn_describe_parm_type {
    EXTFNAPIV4_DESCRIBE_PARM_NAME,
    EXTFNAPIV4_DESCRIBE_PARM_TYPE,
    EXTFNAPIV4_DESCRIBE_PARM_WIDTH,
    EXTFNAPIV4_DESCRIBE_PARM_SCALE,
    EXTFNAPIV4_DESCRIBE_PARM_CAN_BE_NULL,
    EXTFNAPIV4_DESCRIBE_PARM_DISTINCT_VALUES,
    EXTFNAPIV4_DESCRIBE_PARM_IS_CONSTANT,
    EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE,

    EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_ORDERBY,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_REQUEST_REWIND,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND,
    EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS,

    EXTFNAPIV4_DESCRIBE_PARM_LAST
} a_v4_extfn_describe_parm_type;
```

## Members Summary

Member	Description
<i>EXTFNAPIV4_DESCRIBE_PARM_NAME</i>	Parameter name (valid identifier).
<i>EXTFNAPIV4_DESCRIBE_PARM_TYPE</i>	Data type.
<i>EXTFNAPIV4_DESCRIBE_PARM_WIDTH</i>	String width (precision for NUMERIC).
<i>EXTFNAPIV4_DESCRIBE_PARM_SCALE</i>	Scale for NUMERIC.
<i>EXTFNAPIV4_DESCRIBE_PARM_CAN_BE_NULL</i>	True, if the value can be NULL.
<i>EXTFNAPIV4_DESCRIBE_PARM_DISTINCT_VALUES</i>	Estimated number of distinct values across all invocations.
<i>EXTFNAPIV4_DESCRIBE_PARM_IS_CONSTANT</i>	True, if parameter is a constant for the statement.
<i>EXTFNAPIV4_DESCRIBE_PARM_CONSTANT_VALUE</i>	The value of a parameter, if known at describe time.
These selectors can retrieve or set properties of a table parameter. These enumerator values cannot be used with scalar parameters:	
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_COLUMNS</i>	The number of columns in the table.
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_NUM_ROWS</i>	Estimated number of rows in the table.
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_ORDERBY</i>	The order of rows in a table.
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_PARTITIONBY</i>	The partitioning; use <i>number_of_columns=0</i> for ANY.
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_REQUEST_REWIND</i>	True, if the consumer wants the ability rewind the input table.
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_HAS_REWIND</i>	Return true, if the producer supports rewind.

Member	Description
<i>EXTFNAPIV4_DESCRIBE_PARM_TABLE_UNUSED_COLUMNS</i>	The list of output table columns that are not going to be consumed by the server or the UDF.
<i>EXTFNAPIV4_DESCRIBE_PARM_LAST</i>	First illegal value for v4 API. Out-of-band value.

## Describe Return (*a\_v4\_extfn\_describe\_return*)

The *a\_v4\_extfn\_describe\_return* enumerated type provides a return value, when *a\_v4\_extfn\_proc\_context.describe\_xxx\_get()* or *a\_v4\_extfn\_proc\_context.describe\_xxx\_set()* does not succeed.

### Implementation

```
typedef enum a_v4_extfn_describe_return {
    EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE           = 0, // the specified operation has no
meaning either for this attribute or in
the current context.
    EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH   = -1, // the provided buffer size
does not match the required length or the
length is insufficient.
    EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER      = -2, // the provided parameter number
is invalid
    EXTFNAPIV4_DESCRIBE_INVALID_COLUMN        = -3, // the column number is invalid
for this table parameter
    EXTFNAPIV4_DESCRIBE_INVALID_STATE         = -4, // the describe method call is not
valid in the present state
    EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE     = -5, // the attribute is known but not
appropriate for this object
    EXTFNAPIV4_DESCRIBE_UNKNOWN_ATTRIBUTE     = -6, // the identified attribute is
not known to this server version
    EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER   = -7, // the specified parameter is
not a table parameter (for describe_col_get()
or set())
    EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE = -8, // the specified attribute
value is illegal
    EXTFNAPIV4_DESCRIBE_LAST                  = -9
} a_v4_extfn_describe_return;
```



## Members Summary

Member	Return Value	Description
<i>EXTFNAPIV4_DESCRIBE_NOT_AVAILABLE</i>	0	The specified operation has no meaning either for this attribute or in the current context.
<i>EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH</i>	-1	The provided buffer size does not match the required length, or the length is insufficient.
<i>EXTFNAPIV4_DESCRIBE_INVALID_PARAMETER</i>	-2	The provided parameter number is invalid.
<i>EXTFNAPIV4_DESCRIBE_INVALID_COLUMN</i>	-3	The column number is invalid for this table parameter.
<i>EXTFNAPIV4_DESCRIBE_INVALID_STATE</i>	-4	The describe method call is invalid in the present state.
<i>EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE</i>	-5	The attribute is known but not appropriate for this object.
<i>EXTFNAPIV4_DESCRIBE_UNKNOWN_ATTRIBUTE</i>	-6	The identified attribute is not known to this server version.
<i>EXTFNAPIV4_DESCRIBE_NON_TABLE_PARAMETER</i>	-7	The specified parameter is not a table parameter (for <code>describe_col_get()</code> or <code>describe_col_set()</code> ).
<i>EXTFNAPIV4_DESCRIBE_INVALID_ATTRIBUTE_VALUE</i>	-8	The specified attribute value is illegal.

Member	Return Value	Description
<i>EXTFNAPIV4_DESCRIBE_LAST</i>	-9	First illegal value for v4 API.

### Description

The return value of `a_v4_extfn_proc_context.describe_xxx_get()` and `a_v4_extfn_proc_context.describe_xxx_set()` is a signed integer. If the result is positive, the operation succeeds, and the value is the number of bytes copied. If the return value is less or equal to zero, the operation does not succeed, and the return value is one of the `a_v4_extfn_describe_return` values.

## Describe UDF Type (`a_v4_extfn_describe_udf_type`)

Use the `a_v4_extfn_describe_udf_type` enumerated type to select the logical property the UDF retrieves or sets.

### Implementation

```
typedef enum a_v4_extfn_describe_udf_type {
    EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS,
    EXTFNAPIV4_DESCRIBE_UDF_LAST
} a_v4_extfn_describe_udf_type;
```

### Members Summary

Member	Description
<i>EXTFNAPIV4_DESCRIBE_UDF_NUM_PARMS</i>	The number of parameters supplied to the UDF.
<i>EXTFNAPIV4_DESCRIBE_UDF_LAST</i>	Out-of-band value.

### Description

The `a_v4_extfn_proc_context.describe_udf_get()` method is used by the UDF to retrieve properties, and the `a_v4_extfn_proc_context.describe_udf_set()` method is used by the UDF to set properties about the UDF as a whole. The `a_v4_extfn_describe_udf_type` enumerator selects the logical property the UDF retrieves or sets.

### See also

- *External Procedure Context (`a_v4_extfn_proc_context`)* on page 287

## Execution State (a\_v4\_extfn\_state)

The `a_v4_extfn_state` enumerated type represents the query processing state of a UDF.

### Implementation

```
typedef enum a_v4_extfn_state {
    EXTFNAPIV4_STATE_INITIAL,           // Server initial state,
    not used by UDF
    EXTFNAPIV4_STATE_ANNOTATION,       // Annotating parse
    tree with UDF reference
    EXTFNAPIV4_STATE_OPTIMIZATION,     // Optimizing
    EXTFNAPIV4_STATE_PLAN_BUILDING,    // Building execution
    plan
    EXTFNAPIV4_STATE_EXECUTING,       // Executing UDF and
    fetching results from UDF
    EXTFNAPIV4_STATE_LAST
} a_v4_extfn_state;
```

### Members Summary

Member	Description
<i>EXTFNAPIV4_STATE_INITIAL</i>	Server initial state. The only UDF method that is called during this state is <code>_start_extfn</code> .
<i>EXTFNAPIV4_STATE_ANNOTATION</i>	Annotating parse tree with UDF reference. The UDF is not invocated during this state.
<i>EXTFNAPIV4_STATE_OPTIMIZATION</i>	Optimizing. The server calls the UDF's <code>_start_extfn</code> method, followed by the <code>_describe_extfn</code> function.
<i>EXTFNAPIV4_STATE_PLAN_BUILDING</i>	Building a query execution plan. The server calls the UDF's <code>_describe_extfn</code> function.
<i>EXTFNAPIV4_STATE_EXECUTING</i>	Executing UDF and fetching results from UDF. The server calls the <code>_describe_extfn</code> function before starting to fetch data from the UDF. The server then calls <code>_evaluate_extfn</code> to start the fetch cycle. During the fetch cycle, the server calls the functions defined in <code>a_v4_extfn_table_func</code> . When fetching finishes, the server calls the UDF's <code>_close_extfn</code> function.
<i>EXTFNAPIV4_STATE_LAST</i>	First illegal value for v4 API. Out-of-band value.

*Description*

The `a_v4_extfn_state` enumeration indicates which stage of UDF execution the server is in. When the server makes a transition from one state to the next, the server informs the UDF it is leaving the previous state by calling the UDF's `_leave_state_extfn` function. The server informs the UDF it is entering the new state by calling the UDF's `_enter_state_extfn` function.

The query processing state of a UDF restricts the operations that the UDF can perform. For example, in the Annotation state, the UDF can retrieve the data types only for constant parameters.

**See also**

- *Query Processing States* on page 121
- `_start_extfn` on page 285
- `_evaluate_extfn` on page 286
- `_enter_state_extfn` on page 287
- `_leave_state_extfn` on page 287
- *Table Functions (a\_v4\_extfn\_table\_func)* on page 314

**External Function (a\_v4\_extfn\_proc)**

---

The server uses the `a_v4_extfn_proc` structure to call into the various entry points in the UDF. The server passes an instance of `a_v4_extfn_proc_context` to each of the functions.

*Method Summary*

Method	Description
<code>_start_extfn</code>	Allocates a structure and stores its address in the <code>_user_data</code> field in the <code>a_v4_extfn_proc_context</code> .
<code>_finish_extfn</code>	Deallocates a structure whose address was stored in the <code>user_data</code> field in the <code>a_v4_extfn_proc_context</code> .
<code>_evaluate_extfn</code>	Required function pointer to be called for each invocation of the function on a new set of argument values.
<code>_describe_extfn</code>	See <i>Describe API</i> on page 204.

Method	Description
<code>_enter_state_extfn</code>	The UDF can use this function to allocate structures.
<code>_leave_state_extfn</code>	The UDF can use this function to release memory or resources needed for the state.

### `_start_extfn`

Use the `_start_extfn` v4 API method as an optional pointer to an initializer function, for which the only argument is a pointer to `a_v4_extfn_proc_context` structure.

#### *Declaration*

```
_start_extfn(
a_v4_extfn_proc_context *
)
```

#### *Usage*

Use the `_start_extfn` method to allocate a structure and store its address in the `_user_data` field in the `a_v4_extfn_proc_context`. This function pointer must be set to the null pointer if there is no need for any initialization.

#### *Parameters*

Parameter	Description
<code>cntxt</code>	The procedure context object.

### `_finish_extfn`

Use the `_finish_extfn` v4 API method as an optional pointer to a shutdown function, for which the only argument is a pointer to `a_v4_extfn_proc_context`.

#### *Declaration*

```
_finish_extfn(
    a_v4_extfn_proc_context *cntxt,
)
```

#### *Usage*

The `_finish_extfn` API deallocates a structure for which the address was stored in the `user_data` field in the `a_v4_extfn_proc_context`. This function pointer must be set to the null pointer if there is no need for any cleanup.

*Parameters*

Parameter	Description
cntxt	The procedure context object.

**evaluate\_extfn**

Use the `_evaluate_extfn` v4 API method as a required function pointer that is called for each invocation of the function on a new set of argument values.

*Declaration*

```
_evaluate_extfn(
    a_v4_extfn_proc_context *cntxt,
    void *args_handle
)
```

*Usage*

The `_evaluate_extfn` function must describe to the server how to fetch results by filling in the `a_v4_extfn_table_func` portion of the `a_v4_extfn_table` structure and use the `set_value` method on the context with argument zero to send this information to the server. This function must also inform the server of its output schema by filling in the `a_v4_extfn_value_schema` of the `a_v4_extfn_table` structure before calling `set_value` on argument 0. It can access its input argument values via the `get_value` callback function. Both constant and nonconstant arguments are available to the UDF at this time.

*Parameters*

Parameter	Description
cntxt	The procedure context object.
args_handle	Handle to the arguments in the server.

**describe\_extfn**

`_describe_extfn` is called at the beginning of each state to allow the server to get and set logical properties. The UDF can do this by using the six describe methods (`describe_parameter_get`, `describe_parameter_set`, `describe_column_get`, `describe_column_set`, `describe_udf_get`, and `describe_udf_set`) in the `a_v4_proc_context` object.

See *Describe API* on page 204.

## enter\_state\_extfn

The UDF can implement the `_enter_state_extfn` v4 API method as an optional entry point to be notified whenever the UDF enters a new state.

### *Declaration*

```
_enter_state_extfn(
    a_v4_extfn_proc_context *cntxt,
)
```

### *Usage*

The UDF can use this notification to allocate structures.

### *Parameters*

Parameter	Description
<code>cntxt</code>	The procedure context object.

## leave\_state\_extfn

The `_leave_state_extfn` v4 API method is an optional entry point the UDF can implement to receive a notification when the UDF moves out of a query processing state.

### *Declaration*

```
_leave_state_extfn(
    a_v4_extfn_proc_context *cntxt,
)
```

### *Usage*

The UDF can use this notification to release memory or resources needed for the state.

### *Parameters*

Parameter	Description
<code>cntxt</code>	The procedure context object.

## External Procedure Context (a\_v4\_extfn\_proc\_context)

Use the `a_v4_extfn_proc_context` structure to retain context information from the server and from the UDF.

### *Implementation*

```
typedef struct a_v4_extfn_proc_context {
    .
```

```

.
.
} a_v4_extfn_proc_context;

```

### Method Summary

Re- turn Type	Method	Description
short	<b>get_value</b>	Gets input arguments to the UDF.
short	<b>get_value_is_constant</b>	Allows the UDF to ask whether a given argument is a constant.
short	<b>set_value</b>	Used by the UDF in either the <code>_evaluate_extfn</code> or <code>_describe_extfn</code> functions to describe to the server what its output will look like and to inform the server how to fetch results from the UDF.
a_sql_ uint32	<b>get_is_cancelled</b>	Call the <b>get_is_cancelled</b> callback every second or two to see if the user has interrupted the current statement.
short	<b>set_error</b>	Rolls back the current statement and generates an error.
void	<b>log_message</b>	Writes a message to the message log.
short	<b>convert_value</b>	Converts one data type to another.
short	<b>get_option</b>	Gets the value of a settable option.
void	<b>alloc</b>	Allocates a block of memory of length at least "len".
void	<b>free</b>	Free the memory allocated by <b>alloc()</b> for the specified lifetime.
a_sql_ uint32	<b>describe_column_get</b>	See <i>*describe_column_get</i> on page 204.
a_sql_ uint32	<b>describe_column_set</b>	See <i>*describe_column_set</i> on page 221.
a_sql_ uint32	<b>describe_parameter_get</b>	See <i>*describe_parameter_get</i> on page 238.
a_sql_ uint32	<b>describe_parameter_set</b>	See <i>*describe_parameter_set</i> on page 257.
a_sql_ uint32	<b>describe_udf_get</b>	See <i>*describe_udf_get</i> on page 273.



Re-turn Type	Method	Description
a_sql_uint32	<b>describe_udf_set</b>	See <i>*describe_udf_set</i> on page 275.
short	<b>open_result_set</b>	Opens a result set for a table value.
short	<b>close_result_set</b>	Closes an open result set.
short	<b>get_blob</b>	Retrieves an input parameter that is a blob.
short	<b>set_cannot_be_distributed</b>	Disables distribution at the UDF level even if the library is distributable.

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>_user_data</i>	void *	This data pointer can be filled in by any usage with whatever context data the external routine requires.
<i>_executionMode</i>	a_sql_uint32	Indicates the debug/trace level requested via the <b>External_UDF_Execution_Mode</b> option. This is a read-only field.
<i>current_state</i>	a_sql_uint32	The <i>current_state</i> attribute reflects the current execution mode of the context. This can be queried from functions such as <i>_describe_extfn</i> to determine what course of action to take.

### Description

In addition to retaining context information from the server and the UDF, the structure `a_v4_extfn_proc_context` allows the UDF to call back into the server to perform certain actions. The UDF can store private data in this structure in the `_user_data` member. An instance of this structure gets passed to the functions in the `a_v4_extfn_proc` method by the server. User data is not maintained until after the server reaches the Annotation state.

### get\_value

Use the `get_value` v4 API method to obtain the values of input arguments sent to the UDF in a SQL query.

### Declaration

```
short get_value(
    void *      arg_handle,
    a_sql_uint32 arg_num,
```

```

    an_extfn_value *value
)

```

### Usage

The `get_value` API is used in an evaluation method to retrieve the value of each input argument to the UDF. For narrow argument data types (>32K), a call to `get_value` is sufficient to retrieve the entire argument value.

The `get_value` API can be called from any API that has access to the `arg_handle` pointer. This includes API functions that take `a_v4_table_context` as a parameter. The `a_v4_table_context` has an `args_handle` member variable that can be used for this purpose.

For all fixed-length data types, the data is available in the returned value and no further calls are necessary to obtain all of the data. The producer can decide what the maximum length is that is returned entirely in the call to `get_value` method. All fixed length data types should be guaranteed to fit in a single contiguous buffer. For variable-length data, the limit is producer-dependant.

For nonfixed-length data types, and depending on the length of the data, a blob may need to be created using the `get_blob` method to get the data. You can use the macro **EXTFN\_IS\_INCOMPLETE** on the value returned by `get_value` to determine whether a blob object is required. If **EXTFN\_IS\_INCOMPLETE** evaluates to true, a blob is required.

For input arguments that are tables, the type is **AN\_EXTFN\_TABLE**. For this type of argument, you must create a result set using the `open_result_set` method to read values in from the table.

If a UDF requires the value of an argument prior to the `_evaluate_extfn` API being called, then the UDF should implement the `_describe_extfn` API. From the `_describe_extfn` API, the UDF can obtain the value of constant expressions using the `describe_parameter_get` method.

### Parameters

Parameter	Description
<b>arg_handle</b>	A context pointer provided by the consumer.
<b>arg_num</b>	The index of the argument to get a value for. The argument index starts at 1.
<b>value</b>	The value of the specified argument.

### Returns

1 if successful, 0 otherwise.

### *an\_extfn\_value* Structure

The **an\_extfn\_value** structure represents the value of an input argument returned by the `get_value` API.

This code shows the declaration of the **an\_extfn\_value** structure:

```
short typedef struct an_extfn_value {
    void*      data;
    a_sql_uint32  piece_len,
    an_extfn_value *value {
        a_sql_uint32  total_len;
        a_sql_uint32  remain_len;
    } len;
    a_sql_data_type  type;
} an_extfn_value;
```

This table describes what the returned values of **an\_extfn\_value** object look like after calling the `get_value` method:

Value Returned by <code>get_value</code> API	EXTFN_IS_IN COMPLETE	total_len	piece_len	data
null	FALSE	0	0	null
empty string	FALSE	0	0	non-null
Size < MAX_UINT32	FALSE	actual	actual	non-null
size < MAX_UINT32	TRUE	actual	0	non-null
size >= MAX_UINT32	TRUE	MAX_UINT32	0	non-null

The `type` field of **an\_extfn\_value** contains the data type of the value. For UDFs that have tables as input arguments, the data type of that argument is **DT\_EXTFN\_TABLE**. For v4 Table UDFs, the `remain_len` field is not used.

### See also

- `_evaluate_extfn` on page 286
- *Table Context (a\_v4\_extfn\_table\_context)* on page 306
- `_describe_extfn` on page 286
- *\*describe\_parameter\_get* on page 238

**get\_value\_is\_constant**

Use the `get_value_is_constant` v4 API method to determine whether the specified input argument value is a constant.

*Declaration*

```
short get_value_is_constant(
    void *          arg_handle,
    a_sql_uint32   arg_num,
    an_extfn_value *value_is_constant
)
```

*Usage*

The UDF can ask whether a given argument is a constant. This is useful for optimizing a UDF, for example, where work can be performed once during the first call to the `_evaluate_extfn` function, rather than for every evaluation call.

*Parameters*

Parameter	Description
<code>arg_handle</code>	Handle the arguments in the server.
<code>arg_num</code>	The index value of the input argument being retrieved. Index values are 1..N.
<code>value_is_constant</code>	Out parameter for storing is constant.

*Returns*

1 if successful, 0 otherwise.

**See also**

- `_evaluate_extfn` on page 286

**set\_value**

Use the `set_value` v4 API method to describe to the consumer how many columns the result set has and how data should be read.

*Declaration*

```
short set_value(
    void *          arg_handle,
    a_sql_uint32   arg_num,
    an_extfn_value *value
)
```

### Usage

This method is used by the UDF in the `_evaluate_extfn` API. The UDF must call the `set_value` method to tell the consumer how many columns are in the result set and what set of `a_v4_extfn_table_func` functions the UDF supports.

For the `set_value` API, the UDF provides an appropriate **arg\_handle** pointer via the `_evaluate_extfn` API, or from the **args\_handle** member of `a_v4_extfn_table_context` structure.

The **value** argument for the `set_value` method must be of type `DT_EXTFN_TABLE` for v4 Table UDFs.

### Parameters

Parameter	Description
<b>arg_handle</b>	A context pointer provided by the consumer.
<b>arg_num</b>	The index of the argument to set a value for. The only supported argument is 0.
<b>value</b>	The value of the specified argument.

### Returns

1 if successful, 0 otherwise.

### See also

- `_evaluate_extfn` on page 286
- `Table Functions (a_v4_extfn_table_func)` on page 314
- `Table Context (a_v4_extfn_table_context)` on page 306

## get\_is\_cancelled

Use the `get_is_cancelled` v4 API method to determine whether the statement has been cancelled.

### Declaration

```
short get_is_cancelled(
    a_v4_extfn_proc_context *      cntxt,
)
```

### Usage

If a UDF entry point is performing work for an extended period of time (many seconds), it should, if possible, call the `get_is_cancelled` callback every second or two to see if the user has interrupted the current statement. If the statement has been interrupted, a nonzero value is returned and the UDF entry point should then immediately return. Call the

`_finish_extfn` function to perform necessary cleanup. Do not subsequently call any other UDF entry points.

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object.

### Returns

A nonzero value, if the statement is interrupted.

### See also

- *Scalar and Aggregate UDF Callback Functions* on page 84

## set\_error

Use the `set_error` v4 API method to communicate an error back to the server and eventually to the user.

### Declaration

```
void set_error(
    a_v4_extfn_proc_context *    cntxt,
    a_sql_uint32                 error_number,
    const char                    *error_desc_string
)
```

### Usage

Call the `set_error` API, if a UDF entry point encounters an error that should send an error message to the user and shut down the current statement. When called, `set_error` API rolls back the current statement and the user sees "Error raised by user-defined function: <error\_desc\_string>". The `SQLCODE` is the negated form of the supplied <error\_number>.

To avoid collisions with existing error codes, UDFs should generate error numbers between 17000 and 99999. If a number outside this range is provided, the statement is still rolled back, but the error message is "Invalid error raised by user-defined function: (<error\_number>) <error\_desc\_string>" with a `SQLCODE` of -1577. The maximum length of **error\_desc\_string** is 140 characters.

After a call to `set_error` is made, the UDF entry point should immediately perform a return; eventually the `_finish_extfn` function is called to perform necessary cleanup. Do not subsequently call any other UDF entry points.

*Parameters*

Parameter	Description
<b>cntxt</b>	The procedure context object
<b>error_number</b>	The error number to set
<b>error_desc_string</b>	The message string to use

**See also**

- *Scalar and Aggregate UDF Callback Functions* on page 84

**log\_message**

Use the `log_message` v4 API method to send a message to the server's message log.

*Declaration*

```
short log_message(
    const char      *msg,
    short          msg_length
)
```

*Usage*

The `log_message` method writes a message to the message log. The message string must be a printable text string no longer than 255 bytes; longer messages may be truncated.

*Parameters*

Parameter	Description
<b>msg</b>	The message string to log
<b>msg_length</b>	The length of the message string

**See also**

- *Controlling Error Checking and Call Tracing* on page 27

**convert\_value**

Use the `convert_value` v4 API method to convert data types.

*Declaration*

```
short convert_value(
    an_extfn_value *input,
```

## API Reference for a\_v4\_extfn

```
    an_extfn_value *output
)
```

### Usage

. The primary use of the `convert_value` API is the converting between `DT_DATE`, `DT_TIME`, and `DT_TIMESTAMP`, and `DT_TIMESTAMP_STRUCT`. An input and output `an_extfn_value` is passed to the function.

### Input Parameters

Parameter	Description
<code>an_extfn_value.data</code>	Input data pointer
<code>an_extfn_value.total_len</code>	Length of input data
<code>an_extfn_value.type</code>	<code>DT_</code> datatype of input

### Output Parameters

Parameter	Description
<code>an_extfn_value.data</code>	UDF supplied output data point
<code>an_extfn_value.piece_len</code>	Maximum length of output data.
<code>an_extfn_value.total_len</code>	Server set length of converted
<code>an_extfn_value.type</code>	<code>DT_</code> datatype of desired output

### Returns

1 if successful, 0 otherwise.

### See also

- `get_value` on page 289

## get\_option

The `get_option` v4 API method gets the value of a settable option.

### Declaration

```
short get_option(
a_v4_extfn_proc_context * cntxt,
char *option_name,
an_extfn_value *output
)
```



*Parameters*

Parameter	Description
<b>cntxt</b>	The procedure context object
<b>option_name</b>	Name of the option to get
<b>output</b>	<ul style="list-style-type: none"> <li>• <code>an_extfn_value.data</code> – UDF supplied output data pointer</li> <li>• <code>an_extfn_value.piece_len</code> – maximum length of output data</li> <li>• <code>an_extfn_value.total_len</code> – server set length of converted output</li> <li>• <code>an_extfn_value.type</code> – server set data type of value</li> </ul>

*Returns*

1 if successful, 0 otherwise.

**See also**

- *External Function Prototypes* on page 94
- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287

**alloc**

The `alloc v4` API method allocates a block of memory.

*Declaration*

```
void*alloc(
    a_v4_extfn_proc_context *cntxt,
    size_t len
)
```

*Usage*

Allocates a block of memory of length at least `len`. The returned memory is 8-byte aligned. Sybase recommends using the `alloc()` method as your only means of memory allocation, which allows the server to keep track of how much memory is used by external routines. The server can adapt other memory users, track leaks, and provide improved diagnostics and monitoring.

Memory tracking is enabled only when `external_UDF_execution_mode` is set to a value of 1 or 2 (validation mode or tracing mode).

*Parameters*

Parameter	Description
<b>cntxt</b>	The procedure context object
<b>len</b>	The length, in bytes, to allocate

**See also**

- *free* on page 298
- *Enabling Memory Tracking* on page 135

**free**

The `free` v4 API method frees an allocated block of memory.

*Declaration*

```
void free(
    a_v4_extfn_proc_context *cntxt,
    void *mem
)
```

*Usage*

Frees the memory allocated by `alloc()` for the specified lifetime.

Memory tracking is enabled only when **external\_UDF\_execution\_mode** is set to a value of 1 or 2 (validation mode or tracing mode).

*Parameters*

Parameter	Description
<b>cntxt</b>	The procedure context object
<b>mem</b>	Pointer to the memory allocated using the <code>alloc</code> method

**See also**

- *alloc* on page 297
- *Enabling Memory Tracking* on page 135

**open\_result\_set**

The `open_result_set` v4 API method opens a result set for a table value.

*Declaration*

```
short open_result_set(
    a_v4_extfn_proc_context *cntxt,
```

```
a_v4_extfn_table *table,
a_v4_extfn_table_context **result_set
)
```

### Usage

`open_result_set` opens a result set for a table value. A UDF can open a result set to read rows from an input parameter of type `DT_EXTFN_TABLE`. The server (or another UDF) can open a result set to read rows from the UDF.

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object
<code>table</code>	The table object on which to open a result set
<code>result_set</code>	An output parameter that is set to be an opened result set

### Returns

1 if successful, 0 otherwise.

See the `fetch_block` and `fetch_into` v4 API method descriptions for examples of the use of `open_result_set`.

### See also

- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287
- *fetch\_into* on page 308
- *fetch\_block* on page 311

## close\_result\_set

The `close_result_set` v4 API method closes an open result set.

### Declaration

```
short close_result_set(
    a_v4_extfn_proc_context *cntxt,
    a_v4_extfn_table_context *result_set
)
```

### Usage

You can only use `close_result_set` once per result set.

*Parameters*

Parameter	Description
<b>cntxt</b>	The procedure context object
<b>result_set</b>	The result set to close

*Returns*

1 if successful, 0 otherwise.

**get\_blob**

Use the `get_blob` v4 API method to retrieve an input blob parameter.

*Declaration*

```
short get_blob(
    void          *arg_handle,
    a_sql_uint32  arg_num,
    a_v4_extfn_blob **blob
)
```

*Usage*

Use `get_blob` to retrieve a blob input parameter after calling `get_value()`. Use the macro `EXTFN_IS_INCOMPLETE` to determine if a blob object is required to read the data for the value returned from `get_value()`, if `piece_len < total_len`. The blob object is returned as an output parameter and is owned by the caller.

`get_blob` obtains a blob handle that can be used to read the contents of the blob. Call this method only on columns that contain blob objects.

*Parameters*

Parameter	Description
<b>arg_handle</b>	Handle to the arguments in the server
<b>arg_num</b>	The argument is a number 1...N
<b>blob</b>	Output argument containing the blob object

*Returns*

1 if successful, 0 otherwise.

**See also**

- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287
- *get\_value* on page 289

## set\_cannot\_be\_distributed

The `set_cannot_be_distributed` v4 API method disables distributions at the UDF level, even if the distribution criteria are met at the library level.

### *Declaration*

```
void set_cannot_be_distributed( a_v4_extfn_proc_context *cntxt )
```

### *Usage*

In the default behavior, if the library is distributable, then the UDF is distributable. Use `set_cannot_be_distributed` in the UDF to push the decision to disable distribution to the server.

## License Information (a\_v4\_extfn\_license\_info)

If you are a Sybase design partner, use the `a_v4_extfn_license_info` structure to define library-level license validations for your UDFs, including your company name, library version information, and a Sybase-supplied license key.

### *Implementation*

```
typedef struct an_extfn_license_info {
    short    version;
} an_extfn_license_info;

typedef struct a_v4_extfn_license_info {
    an_extfn_license_info version;

    const char    name[255];
    const char    info[255];
    void *        key;
} a_v4_extfn_license_info;
```

### *Data Member Summary*

Data Member	Description
<b>version</b>	Internal use only. Must be set to 1.
<b>name</b>	Value the UDF sets as your company name.
<b>info</b>	Value the UDF sets for additional library information such as library version and build numbers.
<b>key</b>	(Sybase partners only) A Sybase-supplied license key. The key is a 26-character array.

## Optimizer Estimate (a\_v4\_extfn\_estimate)

Use the `a_v4_extfn_estimate` structure to describe an estimate, which includes a value and a confidence level.

### Implementation

```
typedef struct a_v4_extfn_estimate {
    double    value;
    double    confidence;
} a_v4_extfn_estimate;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>value</i>	double	The value for the estimate.
<i>confidence</i>	double	The confidence level associated with the estimate. The confidence varies from 0.0 to 1.0, with 0.0 meaning the estimate is invalid and 1.0 meaning the estimate is known to be true.

## Order By List (a\_v4\_extfn\_orderby\_list)

Use the `a_v4_extfn_orderby_list` structure to describe the ORDER BY property of a table.

### Implementation

```
typedef struct a_v4_extfn_orderby_list {
    a_sql_uint32    number_of_elements;
    a_v4_extfn_order_el order_elements[1];    // there are
number_of_elements entries
} a_v4_extfn_orderby_list;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>number_of_elements</i>	a_sql_uint32	The number of entries
<i>order_elements[1]</i>	a_v4_extfn_order_el	The order of the elements

*Description*

There are *number\_of\_elements* entries, each with a flag indicating whether the element is ascending or descending, and a column index indicating the appropriate column in the associated table.

**See also**

- *Column Order (a\_v4\_extfn\_order\_el)* on page 202

## Partition By Column Number (a\_v4\_extfn\_partitionby\_col\_num)

---

The `a_v4_extfn_partitionby_col_num` enumerated type represents the column number to allow the UDF to express **PARTITION BY** support similar to that of SQL support.

*Implementation*

```
typedef enum a_v4_extfn_partitionby_col_num {
    EXTFNAPIV4_PARTITION_BY_COLUMN_NONE = -1,          // NO PARTITION
    BY
    EXTFNAPIV4_PARTITION_BY_COLUMN_ANY = 0,           // PARTITION BY
    ANY
                                                    // + INTEGER representing a specific
column ordinal
} a_v4_extfn_partitionby_col_num;
```

*Members Summary*

Member of a_v4_extfn_partitionby_col_num Enumerated Type	Value	Description
<i>EXTFNAPIV4_PARTITION_BY_COLUMN_NONE</i>	-1	NO PARTITION BY
<i>EXTFNAPIV4_PARTITION_BY_COLUMN_ANY</i>	0	PARTITION BY ANY positive integer representing a specific column ordinal
<i>Column Ordinal Number</i>	$N > 0$	Ordinal for the table column number to partition on

*Description*

This structure allows the UDF to programmatically describe the partitioning and the column to partition on.

Use this enumeration when populating the `a_v4_extfn_column_list` `number_of_columns` field. When describing partition by support to the server, the UDF

## API Reference for a\_v4\_extfn

sets the `number_of_columns` to one of the enumerated values, or to a positive integer representing the number of column ordinals listed. For example, to describe to the server that no partitioning is supported, create the structure as:

```
a_v4_extfn_column_list nopby = {
EXTFNAPIV4_PARTITION_BY_COLUMN_NONE,
0
};
```

The `EXTFNAPIV4_PARTITION_BY_COLUMN_ANY` member informs the server that the UDF supports any form of partitioning.

To describe a set of ordinals to partition on, create the structure as:

```
a_v4_extfn_column_list nopby = {
2,
3, 4
};
```

This describes a partition by over 2 columns whose ordinals are 3 and 4.

---

**Note:** This example is for illustrative purposes only and is not legal code. The caller must allocate the structure accordingly with room for 3 integers.

---

## Row (a\_v4\_extfn\_row)

---

Use the `a_v4_extfn_row` structure to represent the data in a single row.

### Implementation

```
/* a_v4_extfn_row - */
typedef struct a_v4_extfn_row {
    a_sql_uint32          *row_status;
    a_v4_extfn_column_data *column_data;
} a_v4_extfn_row;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>row_status</i>	<code>a_sql_uint32 *</code>	The status of the row. Set to 1 for existing rows and 0 otherwise.
<i>column_data</i>	<code>a_v4_extfn_column_data *</code>	An array of column data for the row.

### Description

The row structure contains information for a specific row of columns. This structure defines the status of an individual row and includes a pointer to the individual columns within the row. The row status is a flag that indicates the existence of a row. The row status flag can be altered by nested fetch calls without requiring manipulation of the row block structure.



The *row\_status* flag set as 1 indicates that the row is available and can be included in the result set. The *row\_status* set as 0 means the row should be ignored. This is useful when the TPF is acting as a filter because TPF may pass through rows of an input table to the result set, but it may also want to skip certain rows, which it can do by setting a status of 0 for those rows.

### See also

- *Column Data (a\_v4\_extfn\_column\_data)* on page 200

## Row Block (a\_v4\_extfn\_row\_block)

---

Use the `a_v4_extfn_row_block` structure to represent the data in a block of rows.

### Implementation

```
/* a_v4_extfn_row_block - */
typedef struct a_v4_extfn_row_block {
    a_sql_uint32      max_rows;
    a_sql_uint32      num_rows;
    a_v4_extfn_row    *row_data;
} a_v4_extfn_row_block;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>max_rows</i>	a_sql_uint32	The maximum number of rows this row block can handle
<i>num_rows</i>	a_sql_uint32	Must be less than or equal to the maximum of rows the row block contains
<i>row_data</i>	a_v4_extfn_row *	The row data vector

### Description

The row block structure is utilized by the `fetch_into` and `fetch_block` methods to allow the production and consumption of data. The allocator sets the maximum number of rows. The producer incorrectly sets the number of rows. The data consumer should not attempt to read more than number of rows produced.

The owner of the `row_block` structure determines the value of *max\_rows* data member. For example, when a table UDF is implementing `fetch_into`, the value of *max\_rows* is determined by the server as the number of rows that can fit into 128K of memory. However, when a table UDF is implementing `fetch_block`, the table UDF itself determines the value of *max\_rows*.

### Restrictions and Limitations

The value for the both the *num\_rows* and *max\_rows* is > 0. The *num\_rows* must be <= *max\_rows*. The *row\_data* field should not be NULL for a valid row block.

## Table (a\_v4\_extfn\_table)

Use the `a_v4_extfn_table` structure to represent how data is stored in a table and how the consumer fetches that data.

### Implementation

```
typedef struct a_v4_extfn_table {
    a_v4_extfn_table_func *func;
    a_sql_uint32          number_of_columns;
} a_v4_extfn_table;
```

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>func</i>	<code>a_v4_extfn_table_func *</code>	This member holds a set of function pointers that the consumer uses to fetch result data
<i>number_of_columns</i>	<code>a_sql_uint32 *</code>	The number of columns in the table

## Table Context (a\_v4\_extfn\_table\_context)

The `a_v4_extfn_table_context` structure represents an open result set over a table.

### Implementation

```
typedef struct a_v4_extfn_table_context {
//    size_t struct_size;

    /* fetch_into() - fetch into a specified row_block. This entry point
       is used when the consumer has a transfer area with a specific format.
       The fetch_into() function will write the fetched rows into the provided row block.
    */
    short (UDF_CALLBACK *fetch_into)(a_v4_extfn_table_context *cntxt,
a_v4_extfn_row_block *);

    /* fetch_block() - fetch a block of rows. This entry point is used
       when the consumer does not need the data in a particular format. For example,
       if the consumer is reading a result set and formatting it as HTML, the consumer
       does not care how the transfer area is layed out. The fetch_block() entry point is
       more efficient if the consumer does not need a specific layout.

       The row_block parameter is in/out. The first call should point to a NULL row
       block.
       The fetch_block() call sets row_block to a block that can be consumed, and this
       block
       should be passed on the next fetch_block() call.
    */
    short (UDF_CALLBACK *fetch_block)(a_v4_extfn_table_context *cntxt,
a_v4_extfn_row_block **row_block);

    /* rewind() - this is an optional entry point. If NULL, rewind is not supported.
       Otherwise,
       the rewind() entry point restarts the result set at the beginning of the table.
    */
};
```

```

*/
short (UDF_CALLBACK *rewind)(a_v4_extfn_table_context *);

/* get_blob() - If the specified column has a blob object, return it. The blob
   is returned as an out parameter and is owned by the caller. This method should
   only be called on a column that contains a blob. The helper macro
EXTFN_COL_IS_BLOB can
   be used to determine whether a column contains a blob.
*/
short (UDF_CALLBACK *get_blob)(a_v4_extfn_table_context *cntxt,
                               a_v4_extfn_column_data *col,
                               a_v4_extfn_blob **blob);

/* The following fields are reserved for future use and must be initialized to NULL.
*/
void *reserved1_must_be_null;
void *reserved2_must_be_null;
void *reserved3_must_be_null;
void *reserved4_must_be_null;
void *reserved5_must_be_null;

a_v4_extfn_proc_context *proc_context;
void *args_handle; // use in
a_v4_extfn_proc_context::get_value() etc.
a_v4_extfn_table *table;
void *user_data;
void *server_internal_use;

/* The following fields are reserved for future use and must be initialized to NULL.
*/
void *reserved6_must_be_null;
void *reserved7_must_be_null;
void *reserved8_must_be_null;
void *reserved9_must_be_null;
void *reserved10_must_be_null;
} a_v4_extfn_table_context;

```

### Method Summary

Data Type	Method	Description
short	<b>fetch_into</b>	Fetch into a specified row_block
short	<b>fetch_block</b>	Fetch a block of rows
short	<b>rewind</b>	Restarts the result set at the beginning of the table
short	<b>get_blob</b>	Return a blob object, if the specified column has a blob object

### Data Members and Data Types Summary

Data Member	Data Type	Description
<i>proc_context</i>	a_v4_extfn_proc_c ontext *	A pointer to the procedure context object. The UDF can use this to set errors, log messages, cancel, and so on.
<i>args_handle</i>	void *	A handle to the arguments provided by the server.

Data Member	Data Type	Description
<i>table</i>	a_v4_extfn_table *	Points to the open result set table. This is populated after a_v4_extfn_proc_context open_result_set has been called.
<i>user_data</i>	void *	This data pointer can be filled in by any usage with whatever context data the external routine requires.
<i>server_internal_use</i>	void *	Internal use only.

### Description

The a\_v4\_extfn\_table\_context structure acts as a middle layer between the producer and the consumer to help manage the data, when the consumer and producer require separate formats.

A UDF can read rows from an input table parameter using a\_v4\_extfn\_table\_context. The server or another UDF can read rows from the result table of a UDF using a\_v4\_extfn\_table\_context.

The server implements the methods of a\_v4\_extfn\_table\_context, which gives the server an opportunity to resolve impedance mismatches.

### See also

- *fetch\_into* on page 308
- *fetch\_block* on page 311
- *rewind* on page 313

## fetch\_into

The `fetch_into` v4 API method fetches data into a specified row block.

### Declaration

```
short fetch_into(
a_v4_extfn_table_context *cntxt,
a_v4_extfn_row_block *)
```

### Usage

The `fetch_into` method is useful when the producer does not know how data should be arranged in memory. This method is used as an entry point when the consumer has a transfer area with a specific format. The `fetch_into()` function writes the fetched rows into the provided row block. This method is part of the a\_v4\_extfn\_table\_context structure.

Use `fetch_into` when the consumer owns the memory for the data transfer area and requests that the producer use this area. You use the `fetch_into` method when the

consumer cares about how the data transfer area is set up and it is up to the producer to perform the necessary data copying into this area.

### Parameters

Parameter	Description
<b>cntxt</b>	The table context object obtained from the <code>open_result_set</code> API
<b>row_block</b>	The row block object to fetch into

### Returns

1 if successful, 0 otherwise.

If the UDF returns 1, the consumer knows that there are more rows left and the `fetch_into` method should be called again. However, a UDF returning a value of 0 indicates that there are no more rows and a call to the `fetch_into` method is unnecessary.

Consider the following procedure definition, which is an example of a TPF function that consumes an input parameter table and produces it as a result table. Both are instances of SQL values that are obtained and returned through the `get_value` and `set_value` v4 API methods, respectively.

```
CREATE PROCEDURE FETCH_EX( IN a INT, INT b TABLE( c1 INT ) )
  RESULT SET ( rc INT )
```

This procedure definition contains two table objects:

- The input table parameter named `b`
- The return result set table

The following example shows how output tables are fetched from by the caller, in this case, the server. The server might decide to use the `fetch_into` method. Input tables are fetched from by the called entity, in this case the TPF. The TPF decides which fetch API to use.

```
SELECT rc from FETCH_EX( 1, TABLE( SELECT c1 from TABLE ) )
```

The example shows that prior to fetching/consuming from an input table, a table context must be established via the `open_result_set` API on the `a_v4_extfn_proc` structure. The `open_result_set` requires a table object, which can be obtained through the `get_value` API.

```
an_extfn_value    arg;
ctx->get_value( args_handle, 3, &arg );

if( arg.type != DT_EXTFN_TABLE ) {
    // handle error
}

a_v4_extfn_table_context    *rs = NULL;
```

```
a_v4_extfn_table      *inTable = arg.data;
ctx->open_result_set( ctx, inTable, &rs );
```

After the table context is created, the `rs` structure executes the `fetch_into` API and fetches the rows.

```
a_v4_extfn_row_block  *rb = // get a row block to hold a series of
INT values.
rs->fetch_into( rs, &rb ) // fetch the rows.
```

Prior to producing rows to a result table, a table object must be created and returned to the caller via the `set_value` API on the `a_v4_extfn_proc_context` structure.

This example shows that a table UDF must create an instance of the `a_v4_extfn_table` structure. Each invocation of the table UDF should return a separate instance of the `a_v4_extfn_table` structure. The table contains the state fields to keep track of the current row and the number of rows to generate. State for a table can be stored as a field of the instance.

```
typedef struct rg_table : a_v4_extfn_table {
    a_sql_uint32      rows_to_generate;
    a_sql_uint32      current_row;
} my_table;
```

In the following example, each time a row is produced, **current\_row** is incremented until the number of rows to be generated is reached, when `fetch_into` returns false to indicate end-of-file. The consumer executes the `fetch_into` API implemented by the table UDF. As part of the call to the `fetch_into` method, the consumer provides the table context, as well as the row block to fetch into.

```
rs->fetch_into( rs, &rb )

short UDF_CALLBACK my_table_func_fetch_into(
    a_v4_extfn_table_context *tctx,
    a_v4_extfn_row_block *rb)
/*****/
{
    my_table *myTable = tctx->table;

    if( rgTable->current_row < rgTable->rows_to_generate ) {
        // Produce the row...
        rgTable->current_row++;
        return 1;
    }

    return 0;
}
```

### See also

- *The fetch\_into Method* on page 130
- *Table Context (a\_v4\_extfn\_table\_context)* on page 306
- *Row Block (a\_v4\_extfn\_row\_block)* on page 305

- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287
- *get\_value* on page 289
- *set\_value* on page 292
- *Table (a\_v4\_extfn\_table)* on page 306

## **fetch\_block**

The `fetch_block` v4 API method fetches a block of rows.

### *Declaration*

```
short fetch_block(
  a_v4_extfn_table_context *cntxt,
  a_v4_extfn_row_block **row_block)
```

### *Usage*

The `fetch_block` method is used as an entry point when the consumer does not need the data in a particular format. `fetch_block` requests that the producer create a data transfer area and provide a pointer to that area. The consumer owns the memory and takes responsibility for copying data from this area.

The `fetch_block` is more efficient if the consumer does not require a specific layout. The `fetch_block` call sets a `fetch_block` to a block that can be consumed, and this block should be passed on the next `fetch_block` call. This method is part of the `a_v4_extfn_table_context` structure.

### *Parameters*

Parameter	Description
<code>cntxt</code>	The table context object.
<code>row_block</code>	An in/out parameter. The first call should always point to a NULL <code>row_block</code> .

When `fetch_block` is called and `row_block` points to NULL, the UDF must allocate a `a_v4_extfn_row_block` structure.

### *Returns*

1 if successful, 0 otherwise.

If the UDF returns 1, the consumer knows that there are more rows left and calls the `fetch_block` method again. However, a UDF returning a value of 0 indicates that there are no more rows and a call to the `fetch_block` method is unnecessary.

Consider the following procedure definition, which is an example of a TPF function that consumes an input parameter table and produces it as a result table. Both are instances of SQL values that are obtained and returned through the `get_value` and `set_value` v4 API methods, respectively.

## API Reference for a\_v4\_extfn

```
CREATE PROCEDURE FETCH_EX( IN a INT, INT b TABLE( c1 INT ) )
    RESULT SET ( rc INT )
```

This procedure definition contains two table objects:

- The input table parameter named b
- The return result set table

The following example shows how output tables are fetched from by the caller, in this case, the server. The server might decide to use the `fetch_block` method. Input tables are fetched from by the called entity, in this case the TPF, which decides which fetch API to use.

```
SELECT rc from FETCH_EX( 1, TABLE( SELECT c1 from TABLE ) )
```

The example shows that prior to fetching/consuming from an input table, a table context must be established via the `open_result_set` API on the `a_v4_extfn_proc` structure. The `open_result_set` requires a table object, which can be obtained through the `get_value` API.

```
an_extfn_value    arg;
ctx->get_value( args_handle, 3, &arg );

if( arg.type != DT_EXTFN_TABLE ) {
    // handle error
}

a_v4_extfn_table_context    *rs = NULL;
a_v4_extfn_table            *inTable = arg.data;
ctx->open_result_set( ctx, inTable, &rs );
```

After the table context is created, the `rs` structure executes the `fetch_block` API and fetches the rows.

```
a_v4_extfn_row_block    *rb = // get a row block to hold a series of
INT values.
rs->fetch_block( rs, &rb ) // fetch the rows.
```

Prior to producing rows to a result table, a table object must be created and returned to the caller via the `set_value` API on the `a_v4_extfn_proc_context` structure.

This example shows that a table UDF must create an instance of the `a_v4_extfn_table` structure. Each invocation of the table UDF should return a separate instance of the `a_v4_extfn_table` structure. The table contains the state fields to keep track of the current row and the number of rows to generate. State for a table can be stored as a field of the instance.

```
typedef struct rg_table : a_v4_extfn_table {
    a_sql_uint32    rows_to_generate;
    a_sql_uint32    current_row;
} my_table;
```

### See also

- *The fetch\_block Method* on page 130



- *Table Context (a\_v4\_extfn\_table\_context)* on page 306
- *Row Block (a\_v4\_extfn\_row\_block)* on page 305
- *External Procedure Context (a\_v4\_extfn\_proc\_context)* on page 287
- *get\_value* on page 289
- *set\_value* on page 292
- *open\_result\_set* on page 298
- *Table (a\_v4\_extfn\_table)* on page 306

## **rewind**

Use the `rewind v4` API method to restart a result set at the beginning of the table.

### *Declaration*

```
short rewind(
    a_v4_extfn_table_context      *cntxt,
)
```

### *Usage*

Call the `rewind` method on an open result set to rewind the table to the beginning. If the UDF intends to rewind an input table, it must inform the producer during the state

**EXTFNAPIV4\_STATE\_OPTIMIZATION** using the

**EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND** parameter.

`rewind()` is an optional entry point. If `NULL`, `rewind` is not supported. Otherwise, the `rewind()` entry point restarts the result set at the beginning of the table.

### *Parameters*

Parameter	Description
<code>cntxt</code>	The table context object

### *Returns*

1 if successful, 0 otherwise.

### **See also**

- *Query Optimization State* on page 124
- *Execution State* on page 128
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)* on page 269

## **get\_blob**

Use the `get_blob` v4 API method to return a blob object from a specified column.

### *Declaration*

```
short get_blob(
a_v4_extfn_table_context *cntxt,
a_v4_extfn_column_data *col,
a_v4_extfn_blob **blob
)
```

### *Usage*

The blob is returned as an output parameter and is owned by the caller. Call this method only on a column that contains a blob.

Use the helper macro `EXTFN_COL_IS_BLOB` to determine whether a column contains a blob. This is the declaration of `EXTFN_COL_IS_BLOB` in the header file `extfnapi_v4.h`:

```
#define EXTFN_COL_IS_BLOB(c, n) (c[n].blob_handle != NULL)
```

### *Parameters*

Parameter	Description
<code>cntxt</code>	The table context object
<code>col</code>	The column data pointer for which to get the blob
<code>blob</code>	On success, contains the blob object associated with the column

### *Returns*

1 if successful, 0 otherwise.

### **See also**

- *Table Context (a\_v4\_extfn\_table\_context)* on page 306

## **Table Functions (a\_v4\_extfn\_table\_func)**

The consumer uses the `a_v4_extfn_table_func` structure to retrieve results from the producer.

### *Implementation*

```
typedef struct a_v4_extfn_table_func {
//     size_t struct_size;
}
```

```

/* Open a result set. The UDF can allocate any resources needed
for the result set.
*/
short (UDF_CALLBACK *_open_extfn)(a_v4_extfn_table_context *);

/* Fetch rows into a provided row block. The UDF should implement
this method if it does
not have a preferred layout for its transfer area.
*/
short (UDF_CALLBACK *_fetch_into_extfn)(a_v4_extfn_table_context
*, a_v4_extfn_row_block
*row_block);

/* Fetch a block that is allocated and configured by the UDF. The
UDF should implement this
method if it has a preferred layout of the transfer area.
*/
short (UDF_CALLBACK *_fetch_block_extfn)
(a_v4_extfn_table_context *, a_v4_extfn_row_block
**row_block);

/* Restart a result set at the beginning of the table. This is an
optional entry point.
*/
short (UDF_CALLBACK *_rewind_extfn)(a_v4_extfn_table_context *);

/* Close a result set. The UDF can release any resources
allocated for the result set.
*/
short (UDF_CALLBACK *_close_extfn)(a_v4_extfn_table_context *);

/* The following fields are reserved for future use and must be
initialized to NULL. */
void *_reserved1_must_be_null;
void *_reserved2_must_be_null;
} a_v4_extfn_table_func;

```

### Method Summary

Method	Data Type	Description
<code>_open_extfn</code>	void	Called by the server to initiate row fetching by opening a result set. The UDF can allocate any resources needed for the result set.
<code>_fetch_into_extfn</code>	short	Fetch rows into a provided row block. The UDF implements this method, if it does not have a preferred layout for its transfer area.

Method	Data Type	Description
<code>_fetch_block_extfn</code>	short	Fetch a block that is allocated and configured by the UDF. The UDF implements this method, if it has a preferred layout of the transfer area.
<code>_rewind_extfn</code>	void	Optional function called by the server to restart the fetching from the beginning of the table.
<code>_close_extfn</code>	void	Called by the server to terminate row fetching by closing the result set. The UDF can release any resources allocated for the result set.
<code>_reserved1_must_be_null</code>	void	Reserved for future use. Must be initialized to NULL.
<code>_reserved1_must_be_null</code>	void	Reserved for future use. Must be initialized to NULL.

**Description**

The `a_v4_extfn_table_func` structure defines the methods used to fetch results from a table.

**See also**

- *Table (a\_v4\_extfn\_table)* on page 306
- *Table Context (a\_v4\_extfn\_table\_context)* on page 306
- *\_open\_extfn* on page 316
- *\_fetch\_into\_extfn* on page 317
- *\_fetch\_block\_extfn* on page 318
- *\_rewind\_extfn* on page 318
- *\_close\_extfn* on page 319

**open\_extfn**

The server calls the `_open_extfn v4` API method to initiate fetching of rows.

**Declaration**

```
void _open_extfn(
    a_v4_extfn_table_context *cntxt,
)
```

### Usage

The UDF uses this method to open a result set and allocate any resources (for example, streams) needed for sending results to the server.

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object

### See also

- *Table Context (a\_v4\_extfn\_table\_context)* on page 306

## \_\_fetch\_into\_extfn

The `__fetch_into_extfn` v4 API method fetches rows into a provided row block.

### Declaration

```
short __fetch_into_extfn(
    a_v4_extfn_table_context *cntxt,
    a_v4_extfn_row_block *row_block
)
```

### Usage

The UDF should implement this method, if it does not have a preferred layout for its transfer area.

### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object
<code>row_block</code>	The row block object to fetch into.

### Returns

1 if successful, 0 otherwise.

### See also

- *Table Context (a\_v4\_extfn\_table\_context)* on page 306
- *Row Block (a\_v4\_extfn\_row\_block)* on page 305

## fetch\_block\_extfn

The `_fetch_block_extfn` v4 API method fetches a block that is allocated and configured by the UDF.

### *Declaration*

```
short _fetch_block_extfn(
a_v4_extfn_table_context *cntxt,
a_v4_extfn_row_block **
)
```

### *Usage*

The UDF should implement this method, if it has a preferred layout for its transfer area.

### *Parameters*

Parameter	Description
<code>cntxt</code>	The procedure context object
<code>row_block</code>	The row block object to fetch into

### *Returns*

1 if successful, 0 otherwise.

### **See also**

- *Table Context (a\_v4\_extfn\_table\_context)* on page 306
- *Row Block (a\_v4\_extfn\_row\_block)* on page 305

## rewind\_extfn

The `_rewind_extfn` v4 API method restarts a result set at the beginning of the table.

### *Declaration*

```
void _rewind_extfn(
a_v4_extfn_table_context *cntxt,
)
```

### *Usage*

This function is an optional entry point. The UDF implements the `_rewind_extfn` method when the result table is rewound to the beginning. The UDF should consider implementing this method only if it can provide the rewind functionality in an efficient and cost-effective manner.

If a UDF chooses to implement the `_rewind_extfn` method, it should tell the consumer during the state **EXTFNAPIV4\_STATE\_OPTIMIZATION** by setting the **EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND** parameter for argument 0.

The UDF may decide not to provide the rewind functionality, in which case the server compensates and provides the functionality.

**Note:** The server can choose not to call the `_rewind_extfn` method to perform the rewind.

#### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object

#### Returns

No return value.

#### See also

- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Get)* on page 254
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND Attribute (Set)* on page 269
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Set)* on page 271
- *EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HAS\_REWIND Attribute (Get)* on page 255
- *Query Processing States* on page 121
- *Execution State (a\_v4\_extfn\_state)* on page 283
- *Table Context (a\_v4\_extfn\_table\_context)* on page 306

### close\_extfn

The server calls the `_close_extfn v4` API method to terminate fetching of rows.

#### Declaration

```
void _close_extfn(
    a_v4_extfn_table_context *cntxt,
)
```

#### Usage

The UDF uses this method when fetching is complete to close a result set and release any resources allocated for the result set.

#### Parameters

Parameter	Description
<code>cntxt</code>	The procedure context object

**See also**

- *Table Context (a\_v4\_extfn\_table\_context)* on page 306



## API Troubleshooting for a\_v4\_extfn

The `describe_column`, `describe_parameter`, and `describe_udf` v4 API methods can return generic error messages. Executing a UDF that does not exist on the server returns a `Could not execute statement error`.

### Generic `describe_column` Errors

Common error returns for `describe_column` get and set calls.

Get	Set
EXTFNAPIV4_DESCRIBE_BUF- FER_SIZE_MISMATCH – get error returned if the <code>cntxt</code> or <code>describe_buffer</code> are NULL, or if the <code>describe_buffer_length</code> is 0.	EXTFNAPIV4_DESCRIBE_BUF- FER_SIZE_MISMATCH – set error returned if the <code>cntxt</code> or <code>describe_buffer</code> are NULL, or if the <code>describe_buffer_length</code> is 0.
EXTFNAPIV4_DESCRIBE_INVA- LID_STATE – get error returned if the <code>cntxt</code> parameter is not a valid context.	EXTFNAPIV4_DESCRIBE_INVA- LID_STATE – set error returned if the <code>cntxt</code> parameter is not a valid context.
EXTFNAPIV4_DESCRIBE_INVA- LID_PARAMETER – get error returned if the provided parameter number is outside legal range for the procedure: <code>&lt; 0</code> or <code>&gt;</code> the number of parameters for the procedure.	EXTFNAPIV4_DESCRIBE_INVA- LID_PARAMETER – set error returned if the provided parameter number is outside legal range for the procedure: <code>&lt; 0</code> or <code>&gt;</code> the number of parameters for the procedure.
EXTFNAPIV4_DESCRIBE_NON_TA- BLE_PARAMETER – get error returned if <code>arg_num</code> is not a table parameter.	EXTFNAPIV4_DESCRIBE_NON_TA- BLE_PARAMETER – set error returned if <code>arg_num</code> is not a table parameter.
EXTFNAPIV4_DESCRIBE_INVA- LID_COLUMN – get error returned if the column number is not valid for the table parameter.	EXTFNAPIV4_DESCRIBE_INVA- LID_COLUMN – set error returned if the column number is not valid for the table parameter.
EXTFNAPIV4_DESCRIBE_UN- KNOWN_ATTRIBUTE – get error returned if the value of <code>describe_type</code> is not one of the valid describe types from <code>a_v4_extfn_de-</code> <code>scribe_parm_type</code> .	EXTFNAPIV4_DESCRIBE_UN- KNOWN_ATTRIBUTE – set error returned if the value of <code>describe_type</code> is not one of the valid describe types from <code>a_v4_extfn_de-</code> <code>scribe_parm_type</code> .

## Generic describe\_udf Errors

Common error returns for `describe_udf` get and set calls.

Get	Set
EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – get error returned if any of the <code>cntxt</code> or <code>describe_buffer</code> arguments are NULL or if <code>describe_buffer_length</code> is 0.	EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – set error returned if any of the <code>cntxt</code> or <code>describe_buffer</code> arguments are NULL or if <code>describe_buffer_length</code> is 0.
EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – get error returned if the <code>cntxt</code> parameter is an invalid context.	EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – set error returned if the <code>cntxt</code> parameter is an invalid context.
EXTFNAPIV4_DESCRIBE_UNKNOWN_ATTRIBUTE – get error returned if the value of <code>describe_type</code> is not one of the <code>a_v4_extfn_describe_udf_type</code> describe types.	EXTFNAPIV4_DESCRIBE_UNKNOWN_ATTRIBUTE – set error returned if the value of <code>describe_type</code> is not one of the <code>a_v4_extfn_describe_udf_type</code> describe types.

## Generic describe\_parameter Errors

Common error returns for `describe_parameter` get and set calls.

Get	Set
EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – get error returned if the <code>cntxt</code> or <code>describe_buffer</code> is NULL, or if <code>describe_buffer_length</code> is 0.	EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – set error returned if the <code>cntxt</code> or <code>describe_buffer</code> is NULL, or if <code>describe_buffer_length</code> is 0.
EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – get error returned if the <code>cntxt</code> parameter is invalid.	EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – set error returned if the <code>cntxt</code> parameter is invalid.
EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – get error returned if the provided parameter number is outside the legal range for the procedure; that is, if the parameter number is < 0 or > the number of parameters for the procedure.	EXTFNAPIV4_DESCRIBE_INVA-LID_PARAMETER – set error returned if the provided parameter number is outside the legal range for the procedure; that is, if the parameter number is < 0 or > the number of parameters for the procedure.

Get	Set
EXTFNAPIV4_DESCRIBE_UNKNOWN_ATTRIBUTE – get error returned if the value of <b>describe_type</b> is an invalid <code>a_v4_extfn_describe_parm_type</code> describe type.	EXTFNAPIV4_DESCRIBE_UNKNOWN_ATTRIBUTE – set error returned if the value of <b>describe_type</b> is an invalid <code>a_v4_extfn_describe_parm_type</code> describe type.
EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH – get error returned if there is a discrepancy between the requested attribute size and the supplied <code>describe_buffer_length</code> . For fixed-size attributes, such as an <code>a_sql_byte</code> data type, the sizes must match. For variable-length attribute data types, such as <code>char [ ]</code> , the supplied buffer needs to be at least large enough to hold the value of the requested attribute.	EXTFNAPIV4_DESCRIBE_BUFFER_SIZE_MISMATCH – set error returned if there is a discrepancy between the size of the requested attribute and the supplied <b>describe_buffer_length</b> . For fixed-size attributes, such as an <code>a_sql_byte</code> data type, the sizes must match.

**See also**

- *Query Processing States* on page 121

## Missing UDF Returns an Error

---

An attempt to execute a UDF that does not exist on the server returns an error.

If you attempt to execute a query similar to:

```
select my_sum1(n_tabkey) from tabudf()
```

where:

- `tabudf()` is a table UDF, and
- the UDF `my_sum1()` does not exist on the server,

this error is returned:

```
Could not execute statement.
External procedures or functions are not allowed across server types.
SQLCODE=-1579, ODBC 3 State="HY000"
Line 1, column 1
```



# External Environment for UDFs

An improperly defined UDF may cause memory violations or may lead to a database server failure. Running a UDF outside the database server, in an external environment, eliminates this risk to the server.

If a runtime exception occurs in the external environment, the server process is unaffected. The server issues an error to the UDF caller, and any subsequent calls to the UDF result in a restart of the external environment.

---

**Note:** The external runtime environments do not require the IQ\_UDF or IQ\_IDA license. The external runtime environments do not require the `a_v3_extfn` or `a_v4_extfn` APIs.

---

The database server includes support for these external runtime environments for UDFs:

- CLR (Microsoft .NET Common Language Runtime)
- ESQ and ODBC (C/C++ Embedded SQL or ODBC server-side requests)
- Java
- Perl
- PHP

Each environment has its own set of APIs for processing arguments and returning values back to the server. The Java external environment, for example, uses the JDBC API.

## System Tables

The system table `SYSEXTERNENV` stores the information needed to identify and launch each of the external environments.

The system table `SYSEXTERNENVOBJECT` stores non-Java external objects.

## SQL Statements

The following SQL syntax allows you to set or modify the location of external environments in the `SYSEXTERNENV` table.

```
ALTER EXTERNAL ENVIRONMENT environment-name
  [ LOCATION location-string ]
```

Once an external environment is set up to be used on the database server, you can then install objects into the database and create stored procedures and functions that make use of these objects within the external environment. Installing, creating, and using these objects, stored procedures, and stored functions is similar to installing Java classes and creating and using Java stored procedures and functions.

To add a comment for an external environment, you can execute:

```
COMMENT ON EXTERNAL ENVIRONMENT environment-name
  IS comment-string
```

## External Environment for UDFs

To install a Perl or PHP external object (for example, a Perl script) from a file or an expression into the database, execute an **INSTALL EXTERNAL OBJECT** statement similar to:

```
INSTALL EXTERNAL OBJECT object-name-string
  [ update-mode ]
  FROM { FILE file-path | VALUE expression }
  ENVIRONMENT environment-name
```

To add a comment for an installed Perl or PHP external object, you can execute:

```
COMMENT ON EXTERNAL [ENVIRONMENT] OBJECT object-name-string
  IS comment-string
```

To remove an installed Perl or PHP external object from the database, use a **REMOVE EXTERNAL OBJECT** statement:

```
REMOVE EXTERNAL OBJECT object-name-string
```

Once external objects are installed in the database, you can use them in external stored procedure and function definitions (similar to the current mechanism for creating Java stored procedures and functions).

```
CREATE PROCEDURE procedure-name(...)
  EXTERNAL NAME '...'
  LANGUAGE environment-name

CREATE FUNCTION function-name(...)
  RETURNS ...
  EXTERNAL NAME '...'
  LANGUAGE environment-name
```

Once these stored procedures and functions are created, you can use them like any other stored procedure or function in the database. The database server, when encountering an external environment stored procedure or function, automatically launches the external environment (if it has not already been started), and sends the necessary information to get the external environment to fetch the external object from the database and execute it. Any result sets or return values resulting from the execution are returned as needed.

To start or stop an external environment on demand, use the **START EXTERNAL ENVIRONMENT** and **STOP EXTERNAL ENVIRONMENT** statements:

```
START EXTERNAL ENVIRONMENT environment-name
STOP EXTERNAL ENVIRONMENT environment-name
```

## Executing UDFs from an External Environment

---

Execute UDFs in the CLR, ESQL and ODBC, Java, Perl, or PHP external environments.

### Prerequisites

There are no licensing prerequisites. The external runtime environments do not require the IQ\_IDA license. The external runtime environments do not require the `a_v3_extfn` or `a_v4_extfn` APIs.

## Task

1. Set up the external environment to be used on the database server.

```
ALTER EXTERNAL ENVIRONMENT environment-name
 [ LOCATION location-string ]
```

2. Install external objects (CLR, ESQL and ODBC, Java, Perl, or PHP) into the database.
3. Use **CREATE PROCEDURE** and **CREATE FUNCTION** statements to create stored procedures and functions that make use of these objects within the external environment.
4. Reference the stored procedure or function. Reference stored procedures in the **FROM** clause of your query.

## See also

- *The CLR External Environment* on page 327
- *The ESQL and ODBC External Environments* on page 330
- *The Java External Environment* on page 339
- *The PERL External Environment* on page 362
- *The PHP External Environment* on page 366
- *CREATE PROCEDURE Statement (Java UDF)* on page 357
- *CREATE FUNCTION Statement (Java UDF)* on page 359

## External Environment Restrictions

---

Restrictions apply to all external environments for UDFs.

- The **NO RESULT SET** option is not supported.
- Only **IN** parameters are supported: **INOUT/OUT** are not supported.
- Functions with **LONG VARCHAR** or **LONG BINARY** result values are not permitted.

## See also

- *Java External Environment Restrictions* on page 346

## The CLR External Environment

---

The database server includes support for CLR stored procedures and functions. A CLR stored procedure or function behaves the same as a SQL stored procedure or function except that the code for the procedure or function is written in a .NET language such as C# or Visual Basic, and the execution of the procedure or function takes place outside the database server (that is, within a separate .NET executable).

There is only one instance of this .NET executable per database. All connections executing CLR functions and stored procedures use the same .NET executable instance, but the

## External Environment for UDFs

namespaces for each connection are separate. Statics persist for the duration of the connection, but are not shareable across connections. Only .NET version 2.0 is supported.

To call an external CLR function or procedure, you define a corresponding stored procedure or function with an EXTERNAL NAME string defining which DLL to load and which function within the assembly to call. You must also specify LANGUAGE CLR when defining the stored procedure or function. An example declaration follows:

```
CREATE PROCEDURE clr_stored_proc(  
    IN p1 INT,  
    IN p2 UNSIGNED SMALLINT,  
    OUT p3 LONG VARCHAR)  
EXTERNAL NAME 'MyCLRTest.dll::MyCLRTest.Run( int, ushort, out  
string )'  
LANGUAGE CLR;
```

In this example, the stored procedure called **clr\_stored\_proc**, when executed, loads the DLL `MyCLRTest.dll` and calls the function `MyCLRTest.Run`. The **clr\_stored\_proc** procedure takes three SQL parameters, two IN parameters, one of type INT and one of type UNSIGNED SMALLINT, and one OUT parameter of type LONG VARCHAR. On the .NET side, these three parameters translate to input arguments of type int and ushort and an output argument of type string. In addition to out arguments, the CLR function can also have ref arguments. A user must declare a ref CLR argument if the corresponding stored procedure has an INOUT parameter.

The following table lists the various CLR argument types and the corresponding suggested SQL data type:

CLR type	Recommended SQL data type
bool	bit
byte	tinyint
short	smallint
ushort	unsigned smallint
int	int
uint	unsigned int
long	bigint
ulong	unsigned bigint
decimal	numeric
float	real
double	double
DateTime	timestamp



CLR type	Recommended SQL data type
string	long varchar
byte[]	long binary

The declaration of the DLL can be either a relative or absolute path. If the specified path is relative, then the external .NET executable searches the path, and other locations, for the DLL. The executable does not search the Global Assembly Cache (GAC) for the DLL.

Like the existing Java stored procedures and functions, CLR stored procedures and functions can make server-side requests back to the database, and they can return result sets. Also, like Java, any information output to Console.Out and Console.Error is automatically redirected to the database server messages window.

To use CLR in the database, make sure the database server is able to locate and start the CLR executable. You can verify if the database server is able to locate and start the CLR executable by executing the following statement:

```
START EXTERNAL ENVIRONMENT CLR;
```

If the database server fails to start CLR, then the database server is likely not able to locate the CLR executable. The CLR executable is dbextclr12.exe.

Note that the `START EXTERNAL ENVIRONMENT CLR` statement is not necessary other than to verify that the database server can launch CLR executables. In general, making a CLR stored procedure or function call starts CLR automatically.

Similarly, the `STOP EXTERNAL ENVIRONMENT CLR` statement is not necessary to stop an instance of CLR since the instance automatically goes away when the connection terminates. However, if you are completely done with CLR and you want to free up some resources, then the `STOP EXTERNAL ENVIRONMENT CLR` statement releases the CLR instance for your connection.

Unlike the Perl, PHP, and Java external environments, the CLR environment does not require the installation of anything in the database. As a result, you do not need to execute any `INSTALL` statements before using the CLR external environment.

Here is an example of a function written in C# that can be run within an external environment.

```
public class StaticTest
{
    private static int val = 0;

    public static int GetValue() {
        val += 1;
        return val;
    }
}
```

When compiled into a dynamic link library, this function can be called from an external environment. An executable image called `dbextclr12.exe` is started by the database server and it loads the dynamic link library for you.

## External Environment for UDFs

To build this application into a dynamic link library using the Microsoft C# compiler, use a command like the following. The source code for the above example is assumed to reside in a file called `StaticTest.cs`.

```
csc /target:library /out:clrtest.dll StaticTest.cs
```

This command places the compiled code in a DLL called **clrtest.dll**. To call the compiled C# function, `GetValue`, a wrapper is defined as follows using Interactive SQL:

```
CREATE FUNCTION stc_get_value()  
RETURNS INT  
EXTERNAL NAME 'clrtest.dll::StaticTest.GetValue() int '  
LANGUAGE CLR;
```

For CLR, the `EXTERNAL NAME` string is specified in a single line of SQL. You may be required to include the path to the DLL as part of the `EXTERNAL NAME` string so that it can be located. For dependent assemblies (for example, if `myLib.dll` has code that calls functions in, or in some way depends on, `myOtherLib.dll`) then it is up to the .NET Framework to load the dependencies. The CLR External Environment will take care of loading the specified assembly, but extra steps might be required to ensure that dependent assemblies are loaded. One solution is to register all dependencies in the Global Assembly Cache (GAC) by using the Microsoft `gacutil` utility installed with the .NET Framework. For custom-developed libraries, `gacutil` requires that these be signed with a strong name key before they can be registered in the GAC.

To execute the sample compiled C# function, execute the following statement.

```
SELECT stc_get_value();
```

Each time the C# function is called, a new integer result is produced. The sequence of values returned is 1, 2, 3, and so on.

## The ESQL and ODBC External Environments

---

To run a compiled native C function in an external environment instead of within the database server, the stored procedure or function is defined with the `EXTERNAL NAME` clause followed by the `LANGUAGE` attribute specifying one of `C_ESQL32`, `C_ESQL64`, `C_ODBC32`, or `C_ODBC64`.

Unlike the Perl, PHP, and Java external environments, you do not install any source code or compiled objects in the database. As a result, you do not need to execute any `INSTALL` statements before using the ESQL and ODBC external environments.

Here is an example of a function written in C++ that can be run within the database server or in an external environment.

```
#include <windows.h>  
#include <stdio.h>  
#include <stdlib.h>  
#include <string.h>
```

```

#include "extfnapi.h"

BOOL WINAPI DllMain( HMODULE hModule,
    DWORD ul_reason_for_call,
    LPVOID lpReserved
    )
{
    return TRUE;
}

// Note: extfn_use_new_api used only for
// execution in the database server

extern "C" __declspec( dllexport )
a_sql_uint32 extfn_use_new_api( void )
{
    return( EXTFN_API_VERSION );
}

extern "C" __declspec( dllexport )
void SimpleCFunction(
    an_extfn_api *api,
    void *arg_handle )
{
    short          result;
    an_extfn_value arg;
    an_extfn_value retval;
    int *          intptr;
    int            i, j, k;

    j = 1000;
    k = 0;
    for( i = 1; i <= 4; i++ )
    {
        result = api->get_value( arg_handle, i, &arg );
        if( result == 0 || arg.data == NULL ) break;
        if( arg.type & DT_TYPES != DT_INT ) break;
        intptr = (int *) arg.data;
        k += *intptr * j;
        j = j / 10;
    }
    retval.type = DT_INT;
    retval.data = (void*)&k;
    retval.piece_len = retval.len.total_len =
        (a_sql_uint32) sizeof( int );
    api->set_value( arg_handle, 0, &retval, 0 );
    return;
}

```

When compiled into a dynamic link library or shared object, this function can be called from an external environment. An executable image called `dbexternc12` is started by the database server and this executable image loads the dynamic link library or shared object for you.

Note that 32-bit or 64-bit versions of the database server can be used and either version can start 32-bit or 64-bit versions of `dbexternc12`. This is one of the advantages of using the

external environment. Note that once `dbxternc12` is started by the database server, it does not terminate until the connection has been terminated or a `STOP EXTERNAL ENVIRONMENT` statement (with the correct environment name) is executed. Each connection that does an external environment call will get its own copy of `dbxternc12`.

To call the compiled native function, `SimpleCFunction`, a wrapper is defined as follows:

```
CREATE FUNCTION SimpleCDemo(  
    IN arg1 INT,  
    IN arg2 INT,  
    IN arg3 INT,  
    IN arg4 INT )  
RETURNS INT  
EXTERNAL NAME 'SimpleCFunction@c:\\c\\extdemo.dll'  
LANGUAGE C_ODBC32;
```

This is almost identical to the way a compiled native function is described when it is to be loaded into the database server's address space. The one difference is the use of the `LANGUAGE C_ODBC32` clause. This clause indicates that `SimpleCDemo` is a function running in an external environment and that it is using 32-bit ODBC calls. The language specification of `C_ESQL32`, `C_ESQL64`, `C_ODBC32`, or `C_ODBC64` tells the database server whether the external C function issues 32-bit or 64-bit ODBC, ESQL, or `a_v4_extfn` API calls when making server-side requests.

When the native function uses none of the ODBC, ESQL, or SQL Anywhere C API calls to make server-side requests, then either `C_ODBC32` or `C_ESQL32` can be used for 32-bit applications and either `C_ODBC64` or `C_ESQL64` can be used for 64-bit applications. This is the case in the external C function shown above. It does not use any of these APIs.

To execute the sample compiled native function, execute the following statement.

```
SELECT SimpleCDemo(1,2,3,4);
```

To use server-side ODBC, the C/C++ code must use the default database connection. To get a handle to the database connection, call `get_value` with an `EXTFN_CONNECTION_HANDLE_ARG_NUM` argument. The argument tells the database server to return the current external environment connection rather than opening a new one.

```
#include <windows.h>  
#include <stdio.h>  
#include "odbc.h"  
#include "extfnapi.h"  
  
BOOL APIENTRY DllMain( HMODULE hModule,  
    DWORD ul_reason_for_call,  
    LPVOID lpReserved  
    )  
{  
    return TRUE;  
}  
  
extern "C" __declspec( dllexport )
```

```

void ServerSideFunction( an_extfn_api *api, void *arg_handle )
{
    short                result;
    an_extfn_value       arg;
    an_extfn_value       retval;
    SQLRETURN            ret;

    ret = -1;
    // set up the return value struct
    retval.type = DT_INT;
    retval.data = (void*) &ret;
    retval.pieces_len = retval.len.total_len =
        (a_sql_uint32) sizeof( int );

    result = api->get_value( arg_handle,
                            EXTFN_CONNECTION_HANDLE_ARG_NUM,
                            &arg );
    if( result == 0 || arg.data == NULL )
    {
        api->set_value( arg_handle, 0, &retval, 0 );
        return;
    }

    HDBC dbc = (HDBC)arg.data;
    HSTMT stmt = SQL_NULL_HSTMT;
    ret = SQLAllocHandle( SQL_HANDLE_STMT, dbc, &stmt );
    if( ret != SQL_SUCCESS ) return;
    ret = SQLExecDirect( stmt,
                        (SQLCHAR *) "INSERT INTO odbcTab "
                        "SELECT table_id, table_name "
                        "FROM SYS.SYSTAB", SQL_NTS );
    if( ret == SQL_SUCCESS )
    {
        SQLExecDirect( stmt,
                        (SQLCHAR *) "COMMIT", SQL_NTS );
    }
    SQLFreeHandle( SQL_HANDLE_STMT, stmt );

    api->set_value( arg_handle, 0, &retval, 0 );
    return;
}

```

If the above ODBC code is stored in the file `extodbc.cpp`, it can be built for Windows using the following commands.

```
cl extodbc.cpp /LD /Ic:\sa12\sdk\include odbc32.lib
```

The following example creates a table, defines the stored procedure wrapper to call the compiled native function, and then calls the native function to populate the table.

```

CREATE TABLE odbcTab(c1 int, c2 char(128));

CREATE FUNCTION ServerSideODBC( )
RETURNS INT
EXTERNAL NAME 'ServerSideFunction@extodbc.dll'
LANGUAGE C_ODBC32;

```

## External Environment for UDFs

```
SELECT ServerSideODBC();

// The following statement should return two identical rows
SELECT COUNT(*) FROM odbcTab
UNION ALL
SELECT COUNT(*) FROM SYS.SYSTAB;
```

Similarly, to use server-side ESQL, the C/C++ code must use the default database connection. To get a handle to the database connection, call `get_value` with an `EXTFN_CONNECTION_HANDLE_ARG_NUM` argument. The argument tells the database server to return the current external environment connection rather than opening a new one.

```
#include <windows.h>
#include <stdio.h>

#include "sqlca.h"
#include "sqlda.h"
#include "extfnapi.h"

BOOL WINAPI DllMain( HMODULE hModule,
    DWORD ul_reason_for_call,
    LPVOID lpReserved
)
{
    return TRUE;
}

EXEC SQL INCLUDE SQLCA;
static SQLCA *_sqlc;
EXEC SQL SET SQLCA "_sqlc";
EXEC SQL WHENEVER SQLERROR { ret = _sqlc->sqlcode; };

extern "C" __declspec( dllexport )
void ServerSideFunction( an_extfn_api *api, void *arg_handle )
{
    short                result;
    an_extfn_value       arg;
    an_extfn_value       retval;

    EXEC SQL BEGIN DECLARE SECTION;
    char *stmt_text =
        "INSERT INTO esqlTab "
        "SELECT table_id, table_name "
        "FROM SYS.SYSTAB";
    char *stmt_commit =
        "COMMIT";
    EXEC SQL END DECLARE SECTION;

    int ret = -1;

    // set up the return value struct
    retval.type = DT_INT;
    retval.data = (void*) &ret;
```

```

retval.piece_len = retval.len.total_len =
    (a_sql_uint32) sizeof( int );

result = api->get_value( arg_handle,
                        EXTFN_CONNECTION_HANDLE_ARG_NUM,
                        &arg );
if( result == 0 || arg.data == NULL )
{
    api->set_value( arg_handle, 0, &retval, 0 );
    return;
}
ret = 0;
_sqlc = (SQLCA *)arg.data;

EXEC SQL EXECUTE IMMEDIATE :stmt_text;
EXEC SQL EXECUTE IMMEDIATE :stmt_commit;

api->set_value( arg_handle, 0, &retval, 0 );
}

```

If the above embedded SQL statements are stored in the file `extesql.sql`, it can be built for Windows using the following commands.

```

sqlpp extesql.sql extesql.cpp
cl extesql.cpp /LD /Ic:\sa12\sdk\include c:\sa12\sdk\lib
\x86\dblibtm.lib

```

The following example creates a table, defines the stored procedure wrapper to call the compiled native function, and then calls the native function to populate the table.

```

CREATE TABLE esqlTab(c1 int, c2 char(128));

CREATE FUNCTION ServerSideESQL( )
RETURNS INT
EXTERNAL NAME 'ServerSideFunction@extesql.dll'
LANGUAGE C_ESQL32;

SELECT ServerSideESQL();

// The following statement should return two identical rows
SELECT COUNT(*) FROM esqlTab
UNION ALL
SELECT COUNT(*) FROM SYS.SYSTAB;

```

As in the previous examples, to use server-side SQL Anywhere C API calls, the C/C++ code must use the default database connection. To get a handle to the database connection, call `get_value` with an `EXTFN_CONNECTION_HANDLE_ARG_NUM` argument. The argument tells the database server to return the current external environment connection rather than opening a new one. The following example shows the framework for obtaining the connection handle, initializing the C API environment, and transforming the connection handle into a connection object (`a_sqlany_connection`) that can be used with the SQL Anywhere C API.

```

include <windows.h>
#include "sacapidll.h"

```

## External Environment for UDFs

```
#include "extfnapi.h"

BOOL WINAPI DllMain( HMODULE hModule,
    DWORD ul_reason_for_call,
    LPVOID lpReserved
)
{
    return TRUE;
}

extern "C" __declspec( dllexport )
void ServerSideFunction( an_extfn_api *extapi, void *arg_handle )
{
    short          result;
    an_extfn_value arg;
    an_extfn_value retval;
    unsigned       offset;
    char           *cmd;

    SQLAnywhereInterface capi;
    a_sqlany_connection * sqlany_conn;
    unsigned int      max_api_ver;

    result = extapi->get_value( arg_handle,
        EXTFN_CONNECTION_HANDLE_ARG_NUM,
        &arg );

    if( result == 0 || arg.data == NULL )
    {
        return;
    }
    if( !sqlany_initialize_interface( &capi, NULL ) )
    {
        return;
    }
    if( !capi.sqlany_init( "MyApp",
        SQLANY_CURRENT_API_VERSION,
        &max_api_ver ) )
    {
        sqlany_finalize_interface( &capi );
        return;
    }
    sqlany_conn = sqlany_make_connection( arg.data );

    // processing code goes here

    capi.sqlany_fini();

    sqlany_finalize_interface( &capi );
    return;
}
```

If the above C code is stored in the file `extcapi.c`, it can be built for Windows using the following commands.



```
cl /LD /Tp extcapi.c /Tp c:\sa12\SDK\C\sacapidll.c
/Ic:\sa12\SDK\Include c:\sa12\SDK\Lib\X86\dbcapi.lib
```

The following example defines the stored procedure wrapper to call the compiled native function, and then calls the native function.

```
CREATE FUNCTION ServerSideC()
RETURNS INT
EXTERNAL NAME 'ServerSideFunction@extcapi.dll'
LANGUAGE C_ESQL32;

SELECT ServerSideC();
```

The LANGUAGE attribute in the above example specifies C\_ESQL32. For 64-bit applications, you would use C\_ESQL64. You must use the embedded SQL language attribute since the SQL Anywhere C API is built on the same layer (library) as ESQL.

As mentioned earlier, each connection that does an external environment call will start its own copy of dbexternc12. This executable application is loaded automatically by the server the first time an external environment call is made. However, you can use the START EXTERNAL ENVIRONMENT statement to preload dbexternc12. This is useful if you want to avoid the slight delay that is incurred when an external environment call is executed for the first time. Here is an example of the statement.

```
START EXTERNAL ENVIRONMENT C_ESQL32
```

Another case where preloading dbexternc12 is useful is when you want to debug your external function. You can use the debugger to attach to the running dbexternc12 process and set breakpoints in your external function.

The STOP EXTERNAL ENVIRONMENT statement is useful when updating a dynamic link library or shared object. It will terminate the native library loader, dbexternc12, for the current connection thereby releasing access to the dynamic link library or shared object. If multiple connections are using the same dynamic link library or shared object then each of their copies of dbexternc12 must be terminated. The appropriate external environment name must be specified in the STOP EXTERNAL ENVIRONMENT statement. Here is an example of the statement.

```
STOP EXTERNAL ENVIRONMENT C_ESQL32
```

To return a result set from an external function, the compiled native function must use the native function call interface.

The following code fragment shows how to set up a result set information structure. It contains a column count, a pointer to an array of column information structures, and a pointer to an array of column data value structures. The example also uses the SQL Anywhere C API.

```
an_extfn_result_set_info    rs_info;

int columns = capi.sqlany_num_cols( sqlany_stmt );

an_extfn_result_set_column_info *col_info =
    (an_extfn_result_set_column_info *)
```

## External Environment for UDFs

```
    malloc( columns * sizeof(an_extfn_result_set_column_info) );

an_extfn_result_set_column_data *col_data =
    (an_extfn_result_set_column_data *)
    malloc( columns * sizeof(an_extfn_result_set_column_data) );

rs_info.number_of_columns    = columns;
rs_info.column_infos         = col_info;
rs_info.column_data_values   = col_data;
```

The following code fragment shows how to describe the result set. It uses the SQL Anywhere C API to obtain column information for a SQL query that was executed previously by the C API. The information that is obtained from the SQL Anywhere C API for each column is transformed into a column name, type, width, index, and null value indicator that will be used to describe the result set.

```
a_sqlany_column_info      info;
for( int i = 0; i < columns; i++ )
{
    if( sqlany_get_column_info( sqlany_stmt, i, &info ) )
    {
        // set up a column description
        col_info[i].column_name = info.name;
        col_info[i].column_type = info.native_type;
        switch( info.native_type )
        {
            case DT_DATE:          // DATE is converted to string by C API
            case DT_TIME:          // TIME is converted to string by C API
            case DT_TIMESTAMP:     // TIMESTAMP is converted to string by
C API
            case DT_DECIMAL:      // DECIMAL is converted to string by C
API
                col_info[i].column_type = DT_FIXCHAR;
                break;
            case DT_FLOAT:         // FLOAT is converted to double by C API
                col_info[i].column_type = DT_DOUBLE;
                break;
            case DT_BIT:           // BIT is converted to tinyint by C API
                col_info[i].column_type = DT_TINYINT;
                break;
        }
        col_info[i].column_width = info.max_size;
        col_info[i].column_index = i + 1; // column indices are origin
1
        col_info[i].column_can_be_null = info.nullable;
    }
}
// send the result set description
if( extapi->set_value( arg_handle,
                    EXTFN_RESULT_SET_ARG_NUM,
                    (an_extfn_value *)&rs_info,
                    EXTFN_RESULT_SET_DESCRIBE ) == 0 )
{
    // failed
    free( col_info );
}
```

```

    free( col_data );
    return;
}

```

Once the result set has been described, the result set rows can be returned. The following code fragment shows how to return the rows of the result set. It uses the SQL Anywhere C API to fetch the rows for a SQL query that was executed previously by the C API. The rows returned by the SQL Anywhere C API are sent back, one at a time, to the calling environment. The array of column data value structures must be filled in before returning each row. The column data value structure consists of a column index, a pointer to a data value, a data length, and an append flag.

```

a_sqlany_data_value *value = (a_sqlany_data_value *)
    malloc( columns * sizeof(a_sqlany_data_value) );

while( capi.sqlany_fetch_next( sqlany_stmt ) )
{
    for( int i = 0; i < columns; i++ )
    {
        if( capi.sqlany_get_column( sqlany_stmt, i, &value[i] ) )
        {
            col_data[i].column_index = i + 1;
            col_data[i].column_data = value[i].buffer;
            col_data[i].data_length =
(a_sql_uint32)*(value[i].length);
            col_data[i].append = 0;
            if( *(value[i].is_null) )
            {
                // Received a NULL value
                col_data[i].column_data = NULL;
            }
        }
    }
    if( extapi->set_value( arg_handle,
        EXTFN_RESULT_SET_ARG_NUM,
        (an_extfn_value *)&rs_info,
        EXTFN_RESULT_SET_NEW_ROW_FLUSH ) == 0 )
    {
        // failed
        free( value );
        free( col_data );
        free( col_data );
        extapi->set_value( arg_handle, 0, &retval, 0 );
        return;
    }
}

```

## The Java External Environment

---

The database server includes support for Java stored procedures and functions. A Java stored procedure or function behaves the same as a SQL stored procedure or function except that the

## External Environment for UDFs

code for the procedure or function is written in Java and the execution of the procedure or function takes place outside the database server (that is, within a Java VM environment).

It should be noted that there is one instance of the Java VM for each database rather than one instance per connection. Java stored procedures can return result sets.

There are a few prerequisites to using Java in the database support:

1. A copy of the Java Runtime Environment must be installed on the database server computer.
2. The database server must be able to locate the Java executable (the Java VM).

To use Java in the database, make sure that the database server is able to locate and start the Java executable. Verify that this can be done by executing:

```
START EXTERNAL ENVIRONMENT JAVA;
```

If the database server fails to start Java then the problem probably occurs because the database server is not able to locate the Java executable. In this case, you should execute an **ALTER EXTERNAL ENVIRONMENT** statement to explicitly set the location of the Java executable. Make sure to include the executable file name.

```
ALTER EXTERNAL ENVIRONMENT JAVA  
  LOCATION 'java-path';
```

For example:

```
ALTER EXTERNAL ENVIRONMENT JAVA  
  LOCATION 'c:\\jdk1.6.0\\jre\\bin\\java.exe';
```

You can query the location of the Java VM that the database server will use by executing the following SQL query:

```
SELECT db_property('JAVAVM');
```

Note that the **START EXTERNAL ENVIRONMENT JAVA** statement is not necessary other than to verify that the database server can start the Java VM. In general, making a Java stored procedure or function call starts the Java VM automatically.

Similarly, the **STOP EXTERNAL ENVIRONMENT JAVA** statement is not necessary to stop an instance of Java since the instance automatically goes away when the all connections to the database have terminated. However, if you are completely done with Java and you want to make it possible to free up some resources, then the **STOP EXTERNAL ENVIRONMENT JAVA** statement decrements the usage count for the Java VM.

Once you have verified that the database server can start the Java VM executable, the next thing to do is to install the necessary Java class code into the database. Do this by using the **INSTALL JAVA** statement. For example, you can execute the following statement to install a Java class from a file into the database.

```
INSTALL JAVA  
NEW  
FROM FILE 'java-class-file';
```

You can also install a Java JAR file into the database.

```
INSTALL JAVA
NEW
JAR 'jar-name'
FROM FILE 'jar-file';
```

Java classes can be installed from a variable, as follows:

```
CREATE VARIABLE JavaClass LONG VARCHAR;
SET JavaClass = xp_read_file('java-class-file')
INSTALL JAVA
NEW
FROM JavaClass;
```

To remove a Java class from the database, use the **REMOVE JAVA** statement, as follows:

```
REMOVE JAVA CLASS java-class
```

To remove a Java JAR from the database, use the **REMOVE JAVA** statement, as follows:

```
REMOVE JAVA JAR 'jar-name'
```

To modify existing Java classes, you can use the **UPDATE** clause of the **INSTALL JAVA** statement, as follows:

```
INSTALL JAVA
UPDATE
FROM FILE 'java-class-file'
```

You can also update existing Java JAR files in the database.

```
INSTALL JAVA
UPDATE
JAR 'jar-name'
FROM FILE 'jar-file';
```

Java classes can be updated from a variable, as follows:

```
CREATE VARIABLE JavaClass LONG VARCHAR;
SET JavaClass = xp_read_file('java-class-file')
INSTALL JAVA
UPDATE
FROM JavaClass;
```

Once the Java class is installed in the database, you can then create stored procedures and functions to interface to the Java methods. The **EXTERNAL NAME** string contains the information needed to call the Java method and to return OUT parameters and return values. The **LANGUAGE** attribute of the **EXTERNAL NAME** clause must specify **JAVA**. The format of the **EXTERNAL NAME** clause is:

```
EXTERNAL NAME 'java-call' LANGUAGE JAVA
```

## External Environment for UDFs

java-call :

```
[package-name.]class-name.method-name method-signature
```

method-signature :

```
( [ field-descriptor, ... ] ) return-descriptor
```

field-descriptor and return-descriptor :

- Z
- |B
- |S
- |I
- |J
- |F
- |D
- |C
- |V
- |[descriptor
- |Lclass-name;

A Java method signature is a compact character representation of the types of the parameters and the type of the return value. If the number of parameters is less than the number indicated in the method-signature, then the difference must equal the number specified in **DYNAMIC RESULT SETS**, and each parameter in the method signature that is more than those in the procedure parameter list must have a method signature of **[Ljava/SQL/ResultSet;**

For Java UDFs, you do not need to set **DYNAMIC RESULT SETS**; **DYNAMIC RESULT SETS** equal to 1 is implied.

The *field-descriptor* and *return-descriptor* have the following meanings:

Field type	Java data type
B	byte
C	char
D	double
F	float
I	int
J	long
L class-name;	an instance of the class class-name. The class name must be fully qualified, and any dot in the name must be replaced by a /. For example, java/lang/String

Field type	Java data type
S	short
V	void
Z	Boolean
[	use one for each dimension of an array

For example,

```
double some_method(
    boolean a,
    int b,
    java.math.BigDecimal c,
    byte [][] d,
    java.sql.ResultSet[] rs ) {
}
```

would have the following signature:

```
'(ZILjava/math/BigDecimal;[[B[Ljava/SQL/ResultSet;)D'
```

The following procedure creates an interface to a Java method. The Java method does not return any value (V).

```
CREATE PROCEDURE insertfix()
EXTERNAL NAME 'JDBCExample.InsertFixed()'V'
LANGUAGE JAVA;
```

The following procedure creates an interface to a Java method that has a String ([Ljava/lang/String;) input argument. The Java method does not return any value (V).

```
CREATE PROCEDURE InvoiceMain( IN arg1 CHAR(50) )
EXTERNAL NAME 'Invoice.main([Ljava/lang/String;)V'
LANGUAGE JAVA;
```

The following procedure creates an interface to a Java method `Invoice.init` which takes a string argument (Ljava/lang/String;), a double (D), another string argument (Ljava/lang/String;), and another double (D), and returns no value (V).

```
CREATE PROCEDURE init( IN arg1 CHAR(50),
                      IN arg2 DOUBLE,
                      IN arg3 CHAR(50),
                      IN arg4 DOUBLE)
EXTERNAL NAME 'Invoice.init(Ljava/lang/String;DLjava/lang/
String;)V'
LANGUAGE JAVA
```

The following Java example contains the function `main` which takes a string argument and writes it to the database server messages window. It also contains the function `whare` that returns a Java String.

```
import java.io.*;

public class Hello
{
```

## External Environment for UDFs

```
public static void main( String[] args )
{
    System.out.print( "Hello" );
        for ( int i = 0; i < args.length; i++ )
            System.out.print( " " + args[i] );
    System.out.println();
}
public static String whare()
{
    return( "I am SQL Anywhere." );
}
}
```

The Java code above is placed in the file `Hello.java` and compiled using the Java compiler. The class file that results is loaded into the database as follows.

```
INSTALL JAVA
NEW
FROM FILE 'Hello.class';
```

Using Interactive SQL, the stored procedure that will interface to the method `main` in the class `Hello` is created as follows:

```
CREATE PROCEDURE HelloDemo( IN name LONG VARCHAR )
EXTERNAL NAME 'Hello.main([Ljava/lang/String;)V'
LANGUAGE JAVA;
```

Note that the argument to `main` is described as an array of `java.lang.String`. Using Interactive SQL, test the interface by executing the following SQL statement.

```
CALL HelloDemo('SQL Anywhere');
```

If you check the database server messages window, you will find the message written there. All output to `System.out` is redirected to the server messages window.

Using Interactive SQL, the function that will interface to the method `whare` in the class `Hello` is created as follows:

```
CREATE FUNCTION Whare()
RETURNS LONG VARCHAR
EXTERNAL NAME 'Hello.whoAreYou(V)Ljava/lang/String;'
LANGUAGE JAVA;
```

Note that the function `whare` is described as returning a `java.lang.String`. Using Interactive SQL, test the interface by executing the following SQL statement.

```
SELECT Whare();
```

You should see the response in the Interactive SQL Results window.

In attempting to troubleshoot why a Java external environment did not start, that is, if the application gets a "main thread not found" error when a Java call is made, the DBA should check the following:

- If the Java VM is a different bitness than the database server, then ensure that the client libraries with the same bitness as the VM are installed on the database server machine.



- Ensure that the `sajdbc.jar` and `dbjdbc12/libdbjdbc12` shared objects are from the same software build.
- If more than one `sajdbc.jar` are on the database server machine, make sure they are all synchronized to the same software version.
- If the database server machine is very busy, then there is a chance the error is being reported due to a timeout.

**See also**

- *INSTALL JAVA Statement* on page 355
- *CREATE PROCEDURE Statement (Java UDF)* on page 357
- *CREATE FUNCTION Statement (Java UDF)* on page 359
- *REMOVE Statement* on page 360
- *START JAVA Statement* on page 361
- *STOP JAVA Statement* on page 362

## **Java External Environment in a Multiplex**

Before you can use Java external environment UDFs in a multiplex configuration, install the Java class file or JAR files on each node of the multiplex that requires the UDF.

Use Sybase Control Center, Sybase Central, or the Interactive SQL **INSTALL JAVA** statement to install the Java class file and JAR.

**See also**

- *INSTALL JAVA Statement* on page 355

### **Installing a Class Using Sybase Central**

To make your Java class available within the database, you install the class into the database from Sybase Central. You must know the path and file name of the class you want to install.

You require DBA authority to install a class.

To install a class (Sybase Central):

1. Connect to a database as a DBA user.
2. Open the **External Environments** folder.
3. Under this folder, open the **Java** folder.
4. Right-click the right pane and choose **New » Java Class**.
5. Follow the instructions in the wizard.

### Installing a Class Using Interactive SQL

To make your Java class available within the database, you install the class into the database using the `INSTALL JAVA` statement from Interactive SQL. You must know the path and file name of the class you want to install.

1. Connect to the database as a DBA user.
2. Execute the following statement:

```
INSTALL JAVA NEW  
FROM FILE 'path\\ClassName.class';
```

`path` is the directory where the class file is located, and `ClassName.class` is the name of the class file.

The double backslash ensures that the backslash is not treated as an escape character.

For example, to install a class in a file named `Utility.class`, held in the directory `c:\source`, you would execute the following statement:

```
INSTALL JAVA NEW  
FROM FILE 'c:\\source\\Utility.class';
```

If you use a relative path, it must be relative to the current working directory of the database server.

### Java External Environment Restrictions

Before developing Java UDFs and Java table UDFs, familiarize yourself with the restrictions specific to the Java external environment for UDFs.

- Aggregate Java functions are not supported.
- Query fragments involving Java UDFs are not eligible for DQP or SMP processing.
- You cannot **DROP** tables involved in the current query from within the Java external environment.
- You cannot **ALTER** tables involved in the current query from within the Java external environment.
- `UNSIGNED SMALLINT` datatype is not supported.
- Numeric functions are limited to a precision of 255 or less.
- Only one result set is permitted for Java table UDFs.

#### **See also**

- *External Environment Restrictions* on page 327

### Java VM Memory Options

Use the `java_vm_options` option to specify any additional command line options that are required to start the Java virtual machine (VM).

Use this syntax:

```
SET OPTION PUBLIC.java_vm_options='java-options' ;
```

In the following example, you use **java\_vm\_options** to set the maximum heap size of the Java VM to 512 megabytes:

```
SET OPTION PUBLIC.java_vm_options='-Xmx512m' ;
```

In the following example, you set the initial heap size of the Java VM to 32 megabytes:

```
SET OPTION PUBLIC.java_vm_options='-Xms32m' ;
```

## SQL Data Type Conversions for Java UDFs

SQL-to-Java and Java-to-SQL data type conversions are carried out according to the JDBC standard. LOB data types LONG VARCHAR and LONG BINARY are supported for input values but not for return values.

### SQL to Java Data Type Conversion

The data type conversions used by the input values of Java Scalar UDFs and Java Table UDFs.

SQL type	Java type
BIGINT	long
BINARY	byte[ ]
BIT	boolean
CHAR	String
DATE	java.sql.Date
DECIMAL	java.math.BigDecimal
DOUBLE	double
IMAGE	byte[ ]
INTEGER	int
LONG BINARY	byte[ ] <b>Note:</b> Large object data support requires a separately licensed Sybase IQ option.
LONG VARCHAR	String <b>Note:</b> Large object data support requires a separately licensed Sybase IQ option.
MONEY	java.math.BigDecimal

## External Environment for UDFs

SQL type	Java type
NUMERIC	java.math.BigDecimal
REAL	float
SMALLINT	short
SMALLMONEY	java.math.BigDecimal
TEXT	String
TIME	java.sql.Time
TIMESTAMP	java.sql.Timestamp
TINYINT	byte
UNSIGNED BIGINT	java.math.BigDecimal (with a precision of 20 and scale of 0)
UNSIGNED INT	java long
VARBINARY	byte[ ]
VARCHAR	String

### **Java to SQL Data Type Conversion**

The return-value data types of Java scalar UDFs and Java Table UDFs.

Java Type	SQL Type
String	CHAR
String	VARCHAR
String	TEXT
java.math.BigDecimal	NUMERIC
java.math.BigDecimal	MONEY
java.math.BigDecimal	SMALLMONEY
boolean	BIT
byte	TINYINT
short	SMALLINT

Java Type	SQL Type
int	INTEGER
long	BIGINT
float	REAL
double	DOUBLE
byte[ ]	VARBINARY
byte[ ]	IMAGE
java.sql.Date	DATE
java.sql.Time	TIME
java.sql.Timestamp	DATETIME/TIMESTAMP
java.lang.Double	DOUBLE
java.lang.Float	REAL
java.lang.Integer	INTEGER
java.lang.Long	BIGINT

## Creating a Java Scalar UDF

Create and compile a Java class, install the class file onto the server, and create the function definition.

### Prerequisites

- You are familiar with Java and can compile a .java file. You know where the resulting .class file will reside on the file system.
- You are familiar with Interactive SQL. You can connect to the iqdemo database from Interactive SQL, and can issue the **START EXTERNAL ENVIRONMENT JAVA** command from Interactive SQL.

### Task

Use this task as template when creating your own Java UDFs.

1. Place this Java code in a file named HelloJavaUDF.java:

```
public class HelloJavaUDF
{
    public static String helloJava( String name )
    {
        // Simply return Hello and the name passed in.
    }
}
```

## External Environment for UDFs

```
        return "Hello " + name;
    }
}
```

This creates the Java class `HelloJavaUDF` with a static method `helloJava`. The method takes a single string argument and returns a string.

2. Compile `HelloJavaUDF.java`:  
`javac <pathtojavafile>/HelloJavaUDF.java`
3. In Interactive SQL, connect to the `iqdemo` database.
4. In Interactive SQL, install the class file onto the server:

Options	Description
<b>Using absolute path</b>	<code>INSTALL JAVA NEW FROM FILE</code> '<absolutepathtofile>/HelloJavaUDF.class'  Example: <code>INSTALL JAVA NEW FROM FILE 'd:/mydirectory/HelloJavaUDF.class'</code>
<b>Using relative path</b>	<code>INSTALL JAVA NEW FROM FILE</code> '<pathrelativetocwd>/HelloJavaUDF.class'  Example: <code>INSTALL JAVA NEW FROM FILE 'myreldir/HelloJavaUDF.class'</code>

5. In Interactive SQL, create the function definition.  
Provide the following information:
  - The Java package, class, and method names
  - The Java data types of your function arguments, and their corresponding SQL data types
  - The SQL name assign to the Java UDF

```
CREATE FUNCTION my_helloJava(IN name VARCHAR(249) )
RETURNS VARCHAR(255)
EXTERNAL NAME 'example>HelloJavaUDF.helloJava(Ljava/lang/
String;)Ljava/lang/String;'
LANGUAGE JAVA
```

6. In Interactive SQL, use the Java UDF in a query against the `iqdemo` database:  
`SELECT my_helloJava( GivenName ) FROM Customers WHERE ID < 110`

### See also

- *SQL to Java Data Type Conversion* on page 347
- *Java to SQL Data Type Conversion* on page 348

## Creating a Java Scalar UDF Version of the SQL substr Function

Create a Java UDF deployment where the SQL function passes multiple arguments to the Java UDF.

### Prerequisites

- You are familiar with Java and can compile a .java file. You know where the resulting .class file will reside on the file system.
- You are familiar with Interactive SQL. You can connect to the iqdemo database from Interactive SQL, and can issue the **START EXTERNAL ENVIRONMENT JAVA** command from Interactive SQL.

### Task

1. Place this Java code in a file named MyJavaSubstr:

```
public class MyJavaSubstr
{
    public static String my_java_substr( String in, int start, int
length )
    {
        String rc = null;

        if ( start < 1 )
        {
            start = 1;
        }

        // Convert the SQL start, length to Java start, end.
        start --; // Java is 0 based, but SQL is one based.
        int endIndex = start+length;

        try {
            if ( in != null )
            {
                rc = in.substring( start, endIndex );
            }
        } catch ( IndexOutOfBoundsException ex )
        {
            System.out.println("ScalarTestFunctions:
my_java_substr("+in+", "+start+", "+length+"
failed");
            System.out.println(ex);
        }
        return rc;
    }
}
```

2. In Interactive SQL, connect to the iqdemo database.
3. In Interactive SQL, install the class file onto the server:

## External Environment for UDFs

```
INSTALL JAVA NEW FROM FILE '<pathToFile>/  
MyJavaSubstr.class'
```

### 4. In Interactive SQL, create the function definition:

```
CREATE or REPLACE FUNCTION java_substr(IN a VARCHAR(255), IN b  
INT, IN c INT )  
  RETURNS VARCHAR(255)  
  EXTERNAL NAME  
    'example.MyJavaSubstr.my_java_substr(Ljava/lang/  
String;II)Ljava/lang/String;'  
  LANGUAGE JAVA
```

Notice the code snippet **Ljava/lang/String;II** indicating parameter types `String`, `int`, `int`.

### 5. In Interactive SQL, use the Java UDF in a query against the `iqdemo` database:

```
select GivenName, java_substr(Surname,1,1) from Customers where  
lcase(java_substr(Surname,1,1)) = 'a';
```

## Creating a Java Table UDF

Create, compile, and install a Java row generator and create the Java table UDF function definition.

### Prerequisites

- You are familiar with Java and can compile a `.java` file. You know where the resulting `.class` file will reside on the file system.
- You are familiar with Interactive SQL. You can connect to the `iqdemo` database from Interactive SQL, and can issue the **START EXTERNAL ENVIRONMENT JAVA** command from Interactive SQL.

### Task

This example executes a Java row generator (`RowGenerator`) that takes a single integer input and returns that number of rows in a result set. The result set has two columns: one `INTEGER` and one `VARCHAR`. The `RowGenerator` relies on two utility classes:

- `example.ResultSetImpl`
- `example.ResultSetMetaDataImpl`

These are simple implementations of the `java.sql.ResultSet` interface and `java.sql.ResultSetMetaData` interface.

#### 1. Place this code in a file named `RowGenerator.java`:

```
package example;  
  
import java.sql.*;  
  
public class RowGenerator {
```



```

public static void rowGenerator( int numRows, ResultSet rset[] ) {
    // Create the meta data needed for the result set
    ResultSetMetaDataImpl rsmd = new ResultSetMetaDataImpl(2);

    //The first column is the SQL type INTEGER.
    rsmd.setColumnType(1, Types.INTEGER);
    rsmd.setColumnNme(1,"c1");
    rsmd.setColumnLabel(1,"c1");
    rsmd.setTableName(1,"MyTable");

    // The second column is the SQL type VARCHAR length 255
    rsmd.setColumnType(2, Types.VARCHAR);
    rsmd.setColumnNme(2,"c2");
    rsmd.setColumnLabel(2,"c2");
    rsmd.setColumnDisplaySize(2, 255);
    rsmd.setTableName(2,"MyTable");

    // Create result set using the ResultSetMetaData
    ResultSetImpl rs = null;
    try {
        rs = new ResultSetImpl( (ResultSetMetaData)rsmd );
        rs.beforeFirst(); // Make sure we are at the beginning.
    } catch( Exception e ) {
        System.out.println( "Error: couldn't create result set." );
        System.out.println( e.toString() );
    }

    // Add the rows to the result set and populate them
    for( int i = 0; i < numRows; i++ ) {
        try {
            rs.insertRow(); // insert a new row.
            rs.updateInt( 1, i ); // put the integer value in the first column
            rs.updateString( 2, ("Str" + i) ); // put the VARCHAR/String value
in the second column
        } catch( Exception e ) {
            System.out.println( "Error: couldn't insert row/data on row " +
i );
            System.out.println( e.toString() );
        }
    }

    try {
        rs.beforeFirst(); // rewind the result set so that the server gets
it from the beginning.
    } catch( Exception e ) {
        System.out.println( e.toString() );
    }
    rset[0] = rs; // assign the result set to the 1st of the passed in
array.
}
}

```

2. Compile RowGenerator.java, ResultSetImpl.java, and ResultSetMetaData.java. The Windows directory %ALLUSERSPROFILE%

## External Environment for UDFs

\samples\java (\$IQDIR15/samples/java on UNIX) contains  
ResultSetImpl.java and ResultSetMetaData.java.

```
javac <pathtojavafile>/ResultSetMetaDataImpl.java
```

```
javac <pathtojavafile>/ResultSetImpl.java
```

```
javac <pathtojavafile>/RowGenerator.java
```

3. In Interactive SQL, connect to the iqdemo database.

4. In Interactive SQL, install the three class files:

```
INSTALL JAVA NEW FROM FILE '<pathtofile>/  
ResultSetMataDataImpl.class'
```

```
INSTALL JAVA NEW FROM FILE '<pathtofile>/  
ResultSetImpl.class'
```

```
INSTALL JAVA NEW FROM FILE '<pathtofile>/  
RowGenerator.class'
```

5. In Interactive SQL, create the Java Table function definition.

Be ready to provide this information:

- The Java package, class, and method names
- The Java data types of your function arguments, and their corresponding SQL data types
- The SQL name to assign to the Java UDF

```
CREATE or REPLACE PROCEDURE rowgenerator( IN numRows INTEGER )  
  RESULT ( c1 INTEGER , c2 VARCHAR(255) )  
  EXTERNAL NAME  
    'example.RowGenerator.rowGenerator([Ljava/sql/  
ResultSet;])V'  
  LANGUAGE JAVA
```

---

**Note:** The RESULT set has two columns; one INTEGER and the other VARCHAR ( 255 ). The Java prototype has two arguments; one INT ( I ) and the other an array of java.sql.ResultSets ( [Ljava/sql/ResultSet; ). The Java prototype shows the function returning Void ( V ).

---

6. In Interactive SQL, use the Java table UDF in a query against the iqdemo database:

```
SELECT * from rowGenerator(5);
```

The query returns five rows of two columns.

### See also

- *SQL to Java Data Type Conversion* on page 347
- *Java to SQL Data Type Conversion* on page 348

## Java External Environment SQL Statement Reference

Use these SQL statements when developing Java stored procedures and functions.

### INSTALL JAVA Statement

Makes Java classes available for use within a database.

### Syntax

```
INSTALL JAVA [ install-mode ] [ JAR jar-name ] FROM source
```

### Parameters

- **install-mode:** – { **NEW** | **UPDATE** }
- **source:** – { **FILE** *filename* | **URL** *url-value* }

### Examples

- **Example 1** – Install the user-created Java class named “Demo” by providing the file name and location of the class:

```
INSTALL JAVA NEW
FROM FILE 'D:\JavaClass\Demo.class'
```

After installation, the class is referenced using its name. Its original file path location is no longer used. For example, this statement uses the class installed in the previous statement:

```
CREATE VARIABLE d Demo
```

If the Demo class was a member of the package `sybase.work`, the fully qualified name of the class must be used:

```
CREATE VARIABLE d sybase.work.Demo
```

- **Example 2** – Install all the classes contained in a zip file and associate them within the database with a JAR file name:

```
INSTALL JAVA
JAR 'Widgets'
FROM FILE 'C:\Jars\Widget.zip'
```

The location of the zip file is not retained and classes must be referenced using the fully qualified class name (package name and class name).

### Usage

- **Install mode** – Specifying an install mode of **NEW** requires that the referenced Java classes be new classes, rather than updates of currently installed classes. An error occurs if a class with the same name exists in the database and the **NEW** install mode is used.

An install mode of **UPDATE** specifies that the referenced Java classes may include replacements for Java classes already installed in the given database.

## External Environment for UDFs

To update a class or JAR, you must have DBA authority and a newer version of the compiled class file or JAR file available in a file on disk.

Only new connections established after installing the class, or that use the class for the first time after installing the class, use the new definition. Once the Java VM loads a class definition, it stays in memory until the connection closes.

If you have been using a Java class or objects based on a class in the current connection, you need to disconnect and reconnect to use the new class definition.

If install mode is omitted, the default is **NEW**.

- **JAR** – Specifies that the *file-name* or *text-pointer* must designate a JAR file or a column containing a JAR. JAR files typically have extensions of `.jar` or `.zip`.

Installed JAR and zip files can be compressed or uncompressed. However, JAR files produced by the Sun JDK `jar` utility are not supported. Files produced by other zip utilities are supported.

If the JAR option is specified, then the JAR is retained as a JAR after the classes that it contains have been installed. That JAR is the associated JAR of each of those classes. The set of JARs installed in a database with the **JAR** option are called the retained JARs of the database.

Retained JARs are referenced in **INSTALL** and **REMOVE** statements. Retained JARs have no effect on other uses of Java-SQL classes. Retained JARs are used by the SQL system for requests by other systems for the class associated with given data. If a requested class has an associated JAR, the SQL system can supply that JAR, rather than the individual class.

*jar-name* is a character string value of length up to 255 bytes. *jar-name* is used to identify the retained JAR in subsequent **INSTALL**, **UPDATE**, and **REMOVE** statements.

- **source** – Specifies the location of the Java classes to be installed.

The formats supported for *file-name* include fully qualified file names, such as `'c:\libs\jarname.jar'` and `'/usr/u/libs/jarname.jar'`, and relative file names, which are relative to the current working directory of the database server.

The *filename* must identify either a class file or a JAR file.

The class definition for each class is loaded by the VM of each connection the first time that class is used. When you **INSTALL** a class, the VM on your connection is implicitly restarted. Therefore, you have immediate access to the new class, whether the **INSTALL** has an *install-mode* of **NEW** or **UPDATE**.

For other connections, the new class is loaded the next time a VM accesses the class for the first time. If the class is already loaded by a VM, that connection does not see the new class until the VM is restarted for that connection (for example, with a **STOP JAVA** and **START JAVA**).

## Standards

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

## Permissions

- Requires DBA permissions to execute the **INSTALL** statement.
- All installed classes can be referenced in any way by any user.

## CREATE PROCEDURE Statement (Java UDF)

Creates an interface to an external Java table UDF.

**CREATE PROCEDURE** Statement reference information for external procedures is located in a separate topic. **CREATE PROCEDURE** Statement reference information for table UDFs is located in a separate topic.

If your query references Sybase IQ tables, note that different syntax and parameters apply compared to a query that references only catalog store tables.

Java table UDFs are only supported in the **FROM** clause.

## Syntax

For a query referencing at least one Sybase IQ table:

```
CREATE[ OR REPLACE ] PROCEDURE
[ owner.]procedure-name ( [ parameter, ...] )
[ RESULT (result-column, ...)]
[ SQL SECURITY { INVOKER | DEFINER } ]
EXTERNAL NAME 'java-call' [ LANGUAGE Java ] }
```

For a query referencing catalog store tables only:

```
CREATE[ OR REPLACE ] PROCEDURE
[ owner.]procedure-name ( [ parameter, ...] )
[ RESULT (result-column, ...)]
| NO RESULT SET
[ DYNAMIC RESULT SETS integer-expression ]
[ SQL SECURITY { INVOKER | DEFINER } ]
EXTERNAL NAME 'java-call' [ LANGUAGE Java ] }
```

## Parameters

- **parameter** – For a query referencing at least one Sybase IQ table:  
[ IN *parameter\_mode* *parameter-name* *data-type* [ DEFAULT *expression* ]

For a query referencing catalog store tables only:

```
[ IN | OUT | INOUT ] parameter_mode parameter-name data-type [ DEFAULT expression ]
```

- **result-column** – *column-name* *data-type*

- **JAVA** – [ **ALLOW** | **DISALLOW SERVER SIDE REQUESTS** ]
- **java-call** – '[ package-name.]class-name.method-name method-signature'

### Usage

For Java table functions, exactly one result set is allowed. If the Java table functions are joined with a Sybase IQ table or if a column from a Sybase IQ table is an argument to the Java table function then only one result set is supported.

If the Java table function is the only item in the **FROM** clause then *N* number of result sets are allowed.

**JAVA:** [ **ALLOW** | **DISALLOW SERVER SIDE REQUESTS** ]:

**DISALLOW** is the default.

**ALLOW** indicates that server-side connections are allowed.

---

**Note:** Do not specify **ALLOW** unless necessary. A setting of **ALLOW** slows down certain types of Sybase IQ table joins. If you change a procedure definition from **ALLOW** to **DISALLOW**, or vice-versa, the change will not be recognized until you make a new connection.

Do not use UDFs with both **ALLOW SERVER SIDE REQUESTS** and **DISALLOW SERVER SIDE REQUESTS** in the same query.

---

### Standards

- SQL—ISO/ANSI SQL compliant.
- Sybase—The Transact-SQL **CREATE PROCEDURE** statement is different.
- SQLJ—The syntax extensions for Java result sets are as specified in the proposed SQLJ1 standard.

### Permissions

Must have RESOURCE authority, unless creating a temporary procedure. Users with DBA authority can create UDFs for other users by specifying an owner. Users must have DBA authority to create an external UDFs or to create an external UDF for another user

### Referencing Temporary Tables Within Procedures

Sharing a temporary table between procedures can cause problems if the table definitions are inconsistent.

For example, suppose you have two procedures `procA` and `procB`, both of which define a temporary table, `temp_table`, and call another procedure called `sharedProc`. Neither `procA` nor `procB` has been called yet, so the temporary table does not yet exist.

Now, suppose that the `procA` definition for `temp_table` is slightly different than the definition in `procB`—while both used the same column names and types, the column order is different.

When you call `procA`, it returns the expected result. However, when you call `procB`, it returns a different result.

This is because when `procA` was called, it created `temp_table`, and then called `sharedProc`. When `sharedProc` was called, the **SELECT** statement inside of it was parsed and validated, and then a parsed representation of the statement is cached so that it can be used again when another **SELECT** statement is executed. The cached version reflects the column ordering from the table definition in `procA`.

Calling `procB` causes the `temp_table` to be recreated, but with different column ordering. When `procB` calls `sharedProc`, the database server uses the cached representation of the **SELECT** statement. So, the results are different.

You can avoid this from happening by doing one of the following:

- ensure that temporary tables used in this way are defined consistently
- consider using a global temporary table instead

### **CREATE FUNCTION Statement (Java UDF)**

Creates a new external Java table UDF function in the database.

#### **Syntax**

```
CREATE [ OR REPLACE | TEMPORARY ] FUNCTION [ owner. ]function-name
( [ parameter, ... ] )
RETURNS data-type routine-characteristics
[ SQL SECURITY { INVOKER | DEFINER } ]
{ compound-statement
| AS tsql-compound-statement
| EXTERNAL NAME 'java-call' LANGUAGE JAVA }
```

#### **Parameters**

- **parameter** – IN *parameter-name data-type* [ **DEFAULT** *expression* ]
- **routine-characteristics** – **ON EXCEPTION RESUME** | [ **NOT** ] **DETERMINISTIC**
- **tsql-compound-statement** – *sql-statement sql-statement ...*
- **environment-name** – **JAVA** [ **ALLOW** | **DISALLOW SERVER SIDE REQUESTS** ]
- **java-call** – '[ *package-name.class-name.method-name method-signature* ]'
- **method-signature** – ( [ *field-descriptor, ...* ] ) *return-descriptor*
- **field-descriptor and return-descriptor** – **Z** | **B** | **S** | **I** | **J** | **F** | **D** | **C** | **V** | [ *descriptor* | **L** *class-name* ;

#### **Examples**

- **Example 1** – Create an external function written in Java:

```
CREATE FUNCTION dba.encrypt( IN name char(254) )
RETURNS VARCHAR
EXTERNAL NAME
```

```
'Scramble.encrypt (Ljava/lang/String;)Ljava/lang/String;'
LANGUAGE JAVA
```

### Usage

LONG BINARY and LONG VARCHAR are not permitted as return-value data types.

EXTERNAL NAME LANGUAGE JAVA—A function that uses **EXTERNAL NAME** with a **LANGUAGE JAVA** clause is a wrapper around a Java method.

**iq-environment-name: JAVA [ ALLOW / DISALLOW SERVER SIDE REQUESTS ]:**

**DISALLOW** is the default.

**ALLOW** indicates that server-side connections are allowed.

---

**Note:** Do not specify **ALLOW** unless necessary. A setting of **ALLOW** slows down certain types of Sybase IQ table joins.

Do not use UDFs with both **ALLOW SERVER SIDE REQUESTS** and **DISALLOW SERVER SIDE REQUESTS** in the same query.

---

### Standards

- SQL—ISO/ANSI SQL compliant.
- Sybase—Not supported by Adaptive Server Enterprise.

### Permissions

Must have RESOURCE authority.

External Java functions must have DBA authority.

### REMOVE Statement

Removes a class, a package, or a JAR file from a database. Removed classes are no longer available for use as a variable type.

### Syntax

```
REMOVE JAVA classes_to_remove
```

### Parameters

- **classes\_to\_remove:** – { **CLASS** *java\_class\_name* [, *java\_class\_name*]... | **PACKAGE** *java\_package\_name* [, *java\_package\_name*]... | **JAR** *jar\_name* [, *jar\_name*]... [ **RETAIN CLASSES** ] }
- **jar\_name:** – *character\_string\_expression*



## Examples

- **Example 1** – Remove a Java class named “Demo” from the current database:

```
REMOVE JAVA CLASS Demo
```

## Usage

Any class, package, or JAR to be removed must already be installed.

*java\_class\_name*—The name of one or more Java classes to be removed. Those classes must be installed classes in the current database.

*java\_package\_name*—The name of one or more Java packages to be removed. Those packages must be the name of packages in the current database.

*jar\_name*—A character string value of maximum length 255.

Each *jar\_name* must be equal to the *jar\_name* of a retained JAR in the current database. Equality of *jar\_name* is determined by the character string comparison rules of the SQL system.

If **JAR...RETAIN CLASSES** is specified, the specified JARs are no longer retained in the database, and the retained classes have no associated JAR. If **RETAIN CLASSES** is specified, this is the only action of the **REMOVE** statement.

## Standards

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise. A similar feature is available in an Adaptive Server Enterprise-compatible manner using nested transactions.

## Permissions

Must have DBA authority or must own the object.

## START JAVA Statement

Starts the Java VM.

## Syntax

```
START EXTERNAL ENVIRONMENT JAVA
```

## Examples

- **Example 1** – Start the Java VM:

```
START EXTERNAL ENVIRONMENT JAVA
```

### **Usage**

Use **START EXTERNAL ENVIRONMENT JAVA** to load the VM at a convenient time, so that when the user starts to use Java functionality there is no initial pause while the VM is loaded.

### **Standards**

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not applicable.

### **Permissions**

Must have DBA authority.

### **STOP JAVA Statement**

Releases resources associated with the Java VM.

### **Syntax**

```
STOP EXTERNAL ENVIRONMENT JAVA
```

### **Usage**

The main use of **STOP EXTERNAL ENVIRONMENT JAVA** is to economize on the use of system resources.

### **Standards**

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not applicable.

### **Permissions**

DBA authority

## **The PERL External Environment**

---

A Perl stored procedure or function behaves the same as a SQL stored procedure or function except that the code for the procedure or function is written in Perl and the execution of the procedure or function takes place outside the database server (that is, within a Perl executable instance).

It should be noted that there is a separate instance of the Perl executable for each connection that uses Perl stored procedures and functions. This behavior is different from Java stored procedures and functions. For Java, there is one instance of the Java VM for each database rather than one instance per connection. The other major difference between Perl and Java is

that Perl stored procedures do not return result sets, whereas Java stored procedures can return result sets.

There are a few prerequisites to using Perl in the database support:

1. Perl must be installed on the database server computer and the database server must be able to locate the Perl executable.
1. The DBD::SQLAnywhere driver must be installed on the database server computer.
2. On Windows, Microsoft Visual Studio must also be installed. This is a prerequisite since it is necessary for installing the DBD::SQLAnywhere driver.

In addition to the above prerequisites, the database administrator must also install the Perl External Environment module.

To install the external environment module (Windows):

- Run the following commands from the SDK\PerlEnv subdirectory:

```
perl Makefile.PL
nmake
nmake install
```

To install the external environment module (UNIX):

- Run the following commands from the sdk/perlenv subdirectory:

```
perl Makefile.PL
make
make install
```

Once the Perl external environment module has been built and installed, the Perl in the database support can be used.

To use Perl in the database, make sure that the database server is able to locate and start the Perl executable. Verify that this can be done by executing:

```
START EXTERNAL ENVIRONMENT PERL;
```

If the database server fails to start Perl, then the problem probably occurs because the database server is not able to locate the Perl executable. In this case, you should execute an **ALTER EXTERNAL ENVIRONMENT** statement to explicitly set the location of the Perl executable. Make sure to include the executable file name.

```
ALTER EXTERNAL ENVIRONMENT PERL
LOCATION 'perl-path';
```

For example:

```
ALTER EXTERNAL ENVIRONMENT PERL
LOCATION 'c:\\Perl\\bin\\perl.exe';
```

Note that the **START EXTERNAL ENVIRONMENT PERL** statement is not necessary other than to verify that the database server can start Perl. In general, making a Perl stored procedure or function call starts Perl automatically.

Similarly, the **STOP EXTERNAL ENVIRONMENT PERL** statement is not necessary to stop an instance of Perl since the instance automatically goes away when the connection terminates.

## External Environment for UDFs

However, if you are completely done with Perl and you want to free up some resources, then the **STOP EXTERNAL ENVIRONMENT PERL** statement releases the Perl instance for your connection.

Once you have verified that the database server can start the Perl executable, the next thing to do is to install the necessary Perl code into the database. Do this by using the **INSTALL** statement. For example, you can execute the following statement to install a Perl script from a file into the database.

```
INSTALL EXTERNAL OBJECT 'perl-script'  
  NEW  
  FROM FILE 'perl-file'  
  ENVIRONMENT PERL;
```

Perl code also can be built and installed from an expression, as follows:

```
INSTALL EXTERNAL OBJECT 'perl-script'  
  NEW  
  FROM VALUE 'perl-statements'  
  ENVIRONMENT PERL;
```

Perl code also can be built and installed from a variable, as follows:

```
CREATE VARIABLE PerlVariable LONG VARCHAR;  
SET PerlVariable = 'perl-statements';  
INSTALL EXTERNAL OBJECT 'perl-script'  
  NEW  
  FROM VALUE PerlVariable  
  ENVIRONMENT PERL;
```

To remove Perl code from the database, use the **REMOVE** statement, as follows:

```
REMOVE EXTERNAL OBJECT 'perl-script'
```

To modify existing Perl code, you can use the **UPDATE** clause of the **INSTALL EXTERNAL OBJECT** statement, as follows:

```
INSTALL EXTERNAL OBJECT 'perl-script'  
  UPDATE  
  FROM FILE 'perl-file'  
  ENVIRONMENT PERL  
INSTALL EXTERNAL OBJECT 'perl-script'  
  UPDATE  
  FROM VALUE 'perl-statements'  
  ENVIRONMENT PERL  
SET PerlVariable = 'perl-statements';  
INSTALL EXTERNAL OBJECT 'perl-script'  
  UPDATE  
  FROM VALUE PerlVariable  
  ENVIRONMENT PERL
```

Once the Perl code is installed in the database, you can then create the necessary Perl stored procedures and functions. When creating Perl stored procedures and functions, the **LANGUAGE** is always **PERL** and the **EXTERNAL NAME** string contains the information

needed to call the Perl subroutines and to return OUT parameters and return values. The following global variables are available to the Perl code on each call:

- **`$sa_perl_return`** – This is used to set the return value for a function call.
- **`$sa_perl_argN`** – where N is a positive integer [0 .. n]. This is used for passing the SQL arguments down to the Perl code. For example, `$sa_perl_arg0` refers to argument 0, `$sa_perl_arg1` refers to argument 1, and so on.
- **`$sa_perl_default_connection`** – This is used for making server-side Perl calls.
- **`$sa_output_handle`** – This is used for sending output from the Perl code to the database server messages window.

A Perl stored procedure can be created with any set of data types for input and output arguments, and for the return value. However, all non-binary data types are mapped to strings when making the Perl call while binary data is mapped to an array of numbers. A simple Perl example follows:

```
INSTALL EXTERNAL OBJECT 'SimplePerlExample'
NEW
FROM VALUE 'sub SimplePerlSub{
    return( ($_[0] * 1000) +
            ($_[1] * 100) +
            ($_[2] * 10) +
            $_[3] );
}'
ENVIRONMENT PERL;

CREATE FUNCTION SimplePerlDemo(
    IN thousands INT,
    IN hundreds INT,
    IN tens INT,
    IN ones INT)
RETURNS INT
EXTERNAL NAME '<file=SimplePerlExample>'
    $sa_perl_return = SimplePerlSub(
        $sa_perl_arg0,
        $sa_perl_arg1,
        $sa_perl_arg2,
        $sa_perl_arg3)'
LANGUAGE PERL;

// The number 1234 should appear
SELECT SimplePerlDemo(1,2,3,4);
```

The following Perl example takes a string and writes it to the database server messages window:

```
INSTALL EXTERNAL OBJECT 'PerlConsoleExample'
NEW
FROM VALUE 'sub WriteToServerConsole { print $sa_output_handle
$_[0]; }'
ENVIRONMENT PERL;

CREATE PROCEDURE PerlWriteToConsole( IN str LONG VARCHAR)
```

## External Environment for UDFs

```
EXTERNAL NAME '<file=PerlConsoleExample>
  WriteToServerConsole( $sa_perl_arg0 )'
LANGUAGE PERL;

// 'Hello world' should appear in the database server messages window
CALL PerlWriteToConsole( 'Hello world' );
```

To use server-side Perl, the Perl code must use the *\$sa\_perl\_default\_connection* variable. The following example creates a table and then calls a Perl stored procedure to populate the table:

```
CREATE TABLE perlTab(c1 int, c2 char(128));

INSTALL EXTERNAL OBJECT 'ServerSidePerlExample'
  NEW
  FROM VALUE 'sub ServerSidePerlSub
    { $sa_perl_default_connection->do(
      "INSERT INTO perlTab SELECT table_id, table_name FROM
SYS.SYSTAB" );
      $sa_perl_default_connection->do(
        "COMMIT" );
    }'
  ENVIRONMENT PERL;

CREATE PROCEDURE PerlPopulateTable()
  EXTERNAL NAME '<file=ServerSidePerlExample> ServerSidePerlSub()'
  LANGUAGE PERL;

CALL PerlPopulateTable();

// The following should return 2 identical rows
SELECT count(*) FROM perlTab
UNION ALL
SELECT count(*) FROM SYS.SYSTAB;
```

## The PHP External Environment

---

A PHP stored procedure or function behaves the same as a SQL stored procedure or function except that the code for the procedure or function is written in PHP and the execution of the procedure or function takes place outside the database server (that is, within a PHP executable instance).

There is a separate instance of the PHP executable for each connection that uses PHP stored procedures and functions. This behavior is quite different from Java stored procedures and functions. For Java, there is one instance of the Java VM for each database rather than one instance per connection. The other major difference between PHP and Java is that PHP stored procedures do not return result sets, whereas Java stored procedures can return result sets. PHP only returns an object of type LONG VARCHAR, which is the output of the PHP script.

There are two prerequisites to using PHP in the database support:

1. A copy of PHP must be installed on the database server computer and the database server must be able to locate the PHP executable.

2. The PHP extension must be installed on the database server computer.

In addition to the above two prerequisites, the database administrator must also install the PHP External Environment module. Prebuilt modules for several versions of PHP are included. To install prebuilt modules, copy the appropriate driver module to your PHP extensions directory (which can be found in `php.ini`). On UNIX, you can also use a symbolic link.

To install the external environment module (Windows):

1. Locate the `php.ini` file for your PHP installation, and open it in a text editor. Locate the line that specifies the location of the `extension_dir` directory. If `extension_dir` is not set to any specific directory, it is a good idea to set it to point to an isolated directory for better system security.
2. Copy the desired external environment PHP module from the installation directory to your PHP extensions directory. Change the `x.y` to reflect the version you have selected.

```
copy "%SQLANY12%\Bin32\php-5.x.y_sqlanywhere_extenv12.dll"
php-dir\ext
```

3. Add the following line to the Dynamic Extensions section of the `php.ini` file to load the external environment PHP module automatically. Change the `x.y` to reflect the version you have selected.

```
extension=php-5.x.y_sqlanywhere_extenv12.dll
```

Save and close `php.ini`.

4. Make sure that you have also installed the PHP driver from the installation directory into your PHP extensions directory. This file name follows the pattern `php-5.x.y_sqlanywhere.dll` where `x` and `y` are the version numbers. It should match the version numbers of the file that you copied in step 2.

To install the external environment module (UNIX):

1. Locate the `php.ini` file for your PHP installation, and open it in a text editor. Locate the line that specifies the location of the `extension_dir` directory. If `extension_dir` is not set to any specific directory, it is a good idea to set it to point to an isolated directory for better system security.
2. Copy the desired external environment PHP module from the installation directory to your PHP installation directory. Change the `x.y` to reflect the version you have selected.

```
cp $SQLANY12/bin32/php-5.x.y_sqlanywhere_extenv12.so
php-dir/ext
```

3. Add the following line to the Dynamic Extensions section of the `php.ini` file to load the external environment PHP module automatically. Change the `x.y` to reflect the version you have selected.

```
extension=php-5.x.y_sqlanywhere_extenv12.so
```

Save and close `php.ini`.

4. Make sure that you have also installed the PHP driver from the installation directory into your PHP extensions directory. This file name follows the pattern

## External Environment for UDFs

`php-5.x.y_sqlanywhere.so` where x and y are the version numbers. It should match the version numbers of the file that you copied in step 2.

To use PHP in the database, the database server must be able to locate and start the PHP executable. You can verify if the database server is able to locate and start the PHP executable by executing the following statement:

```
START EXTERNAL ENVIRONMENT PHP;
```

If you see a message that states that 'external executable' could not be found, then the problem is that the database server is not able to locate the PHP executable. In this case, you should execute an **ALTER EXTERNAL ENVIRONMENT** statement to explicitly set the location of the PHP executable including the executable name or you should ensure that the **PATH** environment variable includes the directory containing the PHP executable.

```
ALTER EXTERNAL ENVIRONMENT PHP  
LOCATION 'php-path';
```

For example:

```
ALTER EXTERNAL ENVIRONMENT PHP  
LOCATION 'c:\\php\\php-5.2.6-win32\\php.exe';
```

To restore the default setting, execute the following statement:

```
ALTER EXTERNAL ENVIRONMENT PHP  
LOCATION 'php';
```

The **START EXTERNAL ENVIRONMENT PHP** statement is not necessary other than to verify that the database server can start PHP. In general, making a PHP stored procedure or function call starts PHP automatically.

Similarly, the **STOP EXTERNAL ENVIRONMENT PHP** statement is not necessary to stop an instance of PHP since the instance automatically goes away when the connection terminates. However, if you are completely done with PHP and you want to free up some resources, then the **STOP EXTERNAL ENVIRONMENT PHP** statement releases the PHP instance for your connection.

Once you have verified that the database server can start the PHP executable, the next thing to do is to install the necessary PHP code into the database. Do this by using the **INSTALL** statement. For example, you can execute the following statement to install a particular PHP script into the database.

```
INSTALL EXTERNAL OBJECT 'php-script'  
NEW  
FROM FILE 'php-file'  
ENVIRONMENT PHP;
```

PHP code can also be built and installed from an expression as follows:

```
INSTALL EXTERNAL OBJECT 'php-script'  
NEW  
FROM VALUE 'php-statements'  
ENVIRONMENT PHP;
```

PHP code can also be built and installed from a variable as follows:



```
CREATE VARIABLE PHPVariable LONG VARCHAR;
SET PHPVariable = 'php-statements';
INSTALL EXTERNAL OBJECT 'php-script'
  NEW
  FROM VALUE PHPVariable
  ENVIRONMENT PHP;
```

To remove PHP code from the database, use the **REMOVE** statement as follows:

```
REMOVE EXTERNAL OBJECT 'php-script';
```

To modify existing PHP code, you can use the **UPDATE** clause of the **INSTALL** statement as follows:

```
INSTALL EXTERNAL OBJECT 'php-script'
  UPDATE
  FROM FILE 'php-file'
  ENVIRONMENT PHP;
INSTALL EXTERNAL OBJECT 'php-script'
  UPDATE
  FROM VALUE 'php-statements'
  ENVIRONMENT PHP;
SET PHPVariable = 'php-statements';
INSTALL EXTERNAL OBJECT 'php-script'
  UPDATE
  FROM VALUE PHPVariable
  ENVIRONMENT PHP;
```

Once the PHP code is installed in the database, you can then go ahead and create the necessary PHP stored procedures and functions. When creating PHP stored procedures and functions, the LANGUAGE is always PHP and the EXTERNAL NAME string contains the information needed to call the PHP subroutines and for returning OUT parameters.

The arguments are passed to the PHP script in the \$argv array, similar to the way PHP would take arguments from the command line (that is, \$argv[1] is the first argument). To set an output parameter, assign it to the appropriate \$argv element. The return value is always the output from the script (as a LONG VARCHAR).

A PHP stored procedure can be created with any set of data types for input or output arguments. However, the parameters are converted to and from a boolean, integer, double, or string for use inside the PHP script. The return value is always an object of type LONG VARCHAR. A simple PHP example follows:

```
INSTALL EXTERNAL OBJECT 'SimplePHPExample'
  NEW
  FROM VALUE '<?php function SimplePHPFunction(
    $arg1, $arg2, $arg3, $arg4 )
  { return ($arg1 * 1000) +
    ($arg2 * 100) +
    ($arg3 * 10) +
    $arg4;
  } ?>'
  ENVIRONMENT PHP;

CREATE FUNCTION SimplePHPDemo(
```

## External Environment for UDFs

```
    IN thousands INT,  
    IN hundreds INT,  
    IN tens INT,  
    IN ones INT)  
RETURNS LONG VARCHAR  
EXTERNAL NAME '<file=SimplePHPExample> print SimplePHPFunction(  
    $argv[1], $argv[2], $argv[3], $argv[4]);'  
LANGUAGE PHP;  
  
// The number 1234 should appear  
SELECT SimplePHPDemo(1,2,3,4);
```

For PHP, the **EXTERNAL NAME** string is specified in a single line of SQL.

To use server-side PHP, the PHP code can use the default database connection. To get a handle to the database connection, call `sasql_pconnect` with an empty string argument (" or ""). The empty string argument tells the PHP driver to return the current external environment connection rather than opening a new one. The following example creates a table and then calls a PHP stored procedure to populate the table:

```
CREATE TABLE phpTab(c1 int, c2 char(128));  
  
INSTALL EXTERNAL OBJECT 'ServerSidePHPExample'  
NEW  
FROM VALUE '<?php function ServerSidePHPSub() {  
    $conn = sasql_pconnect( '' );  
    sasql_query( $conn,  
        "INSERT INTO phpTab  
        SELECT table_id, table_name FROM SYS.SYSTAB" );  
    sasql_commit( $conn );  
} ?>'  
ENVIRONMENT PHP;  
  
CREATE PROCEDURE PHPPopulateTable()  
EXTERNAL NAME '<file=ServerSidePHPExample> ServerSidePHPSub()'  
LANGUAGE PHP;  
  
CALL PHPPopulateTable();  
  
// The following should return 2 identical rows  
SELECT count(*) FROM phpTab  
UNION ALL  
SELECT count(*) FROM SYS.SYSTAB;
```

For PHP, the **EXTERNAL NAME** string is specified in a single line of SQL. In the above example, note that the single quotes are doubled-up because of the way quotes are parsed in SQL. If the PHP source code was in a file, then the single quotes would not be doubled-up.

To return an error back to the database server, throw a PHP exception. The following example shows how to do this.

```
CREATE TABLE phpTab(c1 int, c2 char(128));  
  
INSTALL EXTERNAL OBJECT 'ServerSidePHPExample'  
NEW
```

```

FROM VALUE '<?php function ServerSidePHPSub() {
    $conn = sasql_pconnect( '' );
    if( !sasql_query( $conn,
        "INSERT INTO phpTabNoExist
            SELECT table_id, table_name FROM SYS.SYSTAB" )
    ) throw new Exception(
        sasql_error( $conn ),
        sasql_errorcode( $conn )
    );
    sasql_commit( $conn );
} ?>'
ENVIRONMENT PHP;

CREATE PROCEDURE PHPPopulateTable()
EXTERNAL NAME
    '<file=ServerSidePHPExample> ServerSidePHPSub()'
LANGUAGE PHP;

CALL PHPPopulateTable();

```

The above example should terminate with error `SQL_UNHANDLED_EXTENV_EXCEPTION` indicating that the table `phpTabNoExist` could not be found.



# Index

- \_close\_extfn
  - v4 API method 319
- \_describe\_extfn 204, 286
- \_enter\_state\_extfn 287
- \_fetch\_block\_extfn
  - v4 API method 318
- \_fetch\_into\_extfn
  - v4 API method 317
- \_finish\_extfn 285
- \_leave\_state\_extfn 287
- \_open\_extfn
  - v4 API method 316
- \_rewind\_extfn
  - v4 API method 318
- \_start\_extfn 285
- .NET external environment 325, 327

## A

- a\_v3\_extfn API
  - Upgrading to a\_v4\_extfn API 16
- a\_v4\_extfn API
  - Upgrading from a\_v3\_extfn API 16
- a\_v4\_extfn\_blob
  - blob 195
  - blob\_length 196
  - close\_istream 197
  - open\_istream 197
  - release 198
  - structure 195
- a\_v4\_extfn\_blob\_istream
  - blob input stream 199
  - get 199
  - structure 199
- a\_v4\_extfn\_col\_subset\_of\_input
  - column values subset 203
  - structure 203
- a\_v4\_extfn\_column\_data
  - column data 200
  - structure 200
- a\_v4\_extfn\_column\_list
  - column list 202
  - structure 202
- a\_v4\_extfn\_describe\_col\_type enumerator 277
- a\_v4\_extfn\_describe\_parm\_type enumerator 278

- a\_v4\_extfn\_describe\_return enumerator 280
- a\_v4\_extfn\_describe\_udf\_type enumerator 282
- a\_v4\_extfn\_estimate
  - optimizer estimate 302
  - structure 302
- a\_v4\_extfn\_license\_info 301
- a\_v4\_extfn\_order\_el
  - column order 202
  - structure 202
- a\_v4\_extfn\_orderby\_list
  - order by list 302
  - structure 302
- a\_v4\_extfn\_partitionby\_col\_num enumerator 303
- a\_v4\_extfn\_proc 97
  - external function 284
  - structure 284
- a\_v4\_extfn\_proc\_context
  - convert\_value method 295
  - external procedure context 287
  - get\_blob method 300
  - get\_is\_cancelled method 293
  - get\_value method 289
  - get\_value\_is\_constant method 292
  - log\_message method 295
  - set\_error method 294
  - set\_value method 292
  - structure 287
- a\_v4\_extfn\_row 304
- a\_v4\_extfn\_row\_block 305
- a\_v4\_extfn\_state enumerator 283
- a\_v4\_extfn\_table
  - structure 306
  - table 306
- a\_v4\_extfn\_table\_context
  - get\_blob method 314
  - structure 306
  - table context 306
- a\_v4\_extfn\_table\_func
  - structure 314
  - table functions 314
- aCC
  - HP-UX 22
  - Itanium 22
- aggregate
  - calculation context 61

## Index

- context structure 62
  - creating user-defined function 53
  - descriptor structure 57
  - aggregate functions
    - declaring 47
    - defining 54
    - my\_bit\_or example 51, 73
    - my\_bit\_xor example 51, 70
    - my\_interpolate example 52, 76
    - my\_sum example 50, 65
  - AIX
    - PowerPC 21
    - xIC 21
  - aliases
    - for columns 189
    - in SELECT statement 188, 189
  - ALL
    - keyword in SELECT statement 188
  - alloc 135
    - v4 API method 297, 298
  - ALTER EXTERNAL ENVIRONMENT JAVA 339
  - ALTER PROCEDURE statement
    - syntax 167
  - and data types 183
  - annotation state 122
  - API
    - declaring version 94
    - external functions 94
  - B**
  - BIGINT data type 9
  - BINARY (<n>) data type 9
  - BIT data type 14
  - blob
    - a\_v4\_extfn\_blob 195
  - BLOB data type 9, 14
  - blob input stream
    - a\_v4\_extfn\_blob\_istream 199
  - build.bat 21
  - build.sh 21
  - building
    - shared libraries 19, 21–24
  - C**
  - C/C++
    - new operator 31
    - restrictions 31
  - C/C++ external environment 325, 330
  - calculation
    - aggregate context 61
  - call tracing
    - configuring 27
  - calling pattern
    - aggregate 85
    - aggregate with unbounded window 87
    - optimized cumulative moving window aggregate 90
    - optimized cumulative window aggregate 88
    - optimized moving window following aggregate 92
    - optimized moving window without current 93
  - scalar syntax 85
  - simple aggregate grouped 86
  - simple aggregate ungrouped 86
  - unoptimized cumulative moving window aggregate 89
  - unoptimized cumulative window aggregate 88
  - unoptimized moving window following aggregate 91
  - unoptimized moving window without current 92
- catalog store 185, 191
  - CHAR(<n>) data type 9
  - classes
    - installing 355
    - removing 360
  - CLOB data type 9, 14
  - close\_result\_set
    - v4 API method 299
  - CLR external environment 325, 327
  - column data
    - a\_v4\_extfn\_column\_data 200
  - column list
    - a\_v4\_extfn\_column\_list 202
  - column number
    - partition by 303
  - column order
    - a\_v4\_extfn\_order\_el 202
  - column subset
    - a\_v4\_extfn\_col\_subset\_of\_input 203
  - columns
    - aliases 189
  - compile
    - switches 19, 21–24
  - consumer 101
  - contains-expression 183

- context
  - aggregate structure 62
  - scalar structure 40
- context area 82
- context variables 82
- convert\_value method
  - a\_v4\_extfn\_proc\_context 295
- CREATE AGGREGATE FUNCTION statement
  - 98
  - syntax 47
- CREATE FUNCTION statement 98
  - external environment 359
  - Java 359
  - syntax 33, 83, 172
  - UDF 359
- CREATE PROCEDURE statement for external
  - procedures
  - syntax 169, 357
- creating
  - aggregate user-defined function 53
  - external stored procedures 169, 357
  - scalar user-defined function 38
  - user-defined functions 31, 32
- CUBE operator 192
  - SELECT statement 192
- cumulative window aggregate
  - OLAP-style optimized calling pattern 88
  - OLAP-style unoptimized calling pattern 88
- D**
- data type conversion 347
  - Java to SQL 348
  - SQL to Java 347
- data types 183
  - LONG BINARY 37, 45
  - supported 9
  - unsupported 14
- debug environment
  - Microsoft Visual Studio 28
- DECIMAL(<precision>, <scale>) data type 14
- declaration
  - aggregate 47
  - aggregate my\_bit\_or example 51
  - aggregate my\_bit\_xor example 51
  - aggregate my\_interpolate example 52
  - aggregate my\_sum example 50
  - scalar 33
  - scalar my\_byte\_length example 37
  - scalar my\_plus example 34
  - scalar my\_plus\_counter example 36
- declaring
  - API version 94
- DEFAULT\_TABLE\_UDF\_ROW\_COUNT option
  - 178
- definition
  - aggregate functions 54
  - aggregate my\_bit\_or example 73
  - aggregate my\_bit\_xor example 70
  - aggregate my\_interpolate example 76
  - aggregate my\_sum example 65
  - scalar functions 38
  - scalar my\_byte\_length example 45
  - scalar my\_plus example 42
  - scalar my\_plus\_counter example 43
- describe
  - return value 280
- describe\_column
  - errors, generic 321
- describe\_column\_get 204
  - attributes 205
- describe\_column\_set 221
  - attributes 222
- describe\_parameter
  - errors, generic 322
- describe\_parameter\_get 141, 238
- describe\_parameter\_set 141, 257
- describe\_udf
  - errors, generic 322
- describe\_udf\_get 273
  - attributes 274
- describe\_udf\_set 275
- description
  - aggregate structure 57
  - scalar structure 39
- disable
  - user-defined functions 25
- disjunction of subquery predicates 191
- DISTINCT keyword 188
- DOUBLE data type 9
- DQP 346
- dropping
  - user-defined functions 29
- DUMMY 185
- dummy IQ table 185
- dynamic library interface
  - configuring 15

**E**

## enabling

- user-defined functions 3, 25

## enumerated type

- a\_v4\_extfn\_describe\_col\_type 277
- a\_v4\_extfn\_describe\_parm\_type 278
- a\_v4\_extfn\_describe\_return 280
- a\_v4\_extfn\_describe\_udf\_type 282
- a\_v4\_extfn\_partitionby\_col\_num 303
- a\_v4\_extfn\_state 283

## error checking

- configuring 27
- UDF does not exist 323

## ESQL external environment 330

## evaluate\_extfn 286

## evaluating statements 27

## execute permissions

- granting 29
- revoking 29

## executing partition state 128

## execution state 128

- a\_v4\_extfn\_state enumerator 283

## exporting data

- SELECT statement 185

## External environment 325

- restrictions 327, 346

## external function

- a\_v4\_extfn\_proc 284
- prototypes 94

## external library

- unloading 26

## EXTERNAL NAME clause 33

## external procedure context

- a\_v4\_extfn\_proc\_context 287
- alloc method 297, 298
- close\_result\_set method 299
- get\_option method 296
- open\_result\_set method 298
- set\_cannot\_be\_distributed 301

## external procedures

- creating 169, 357

## external stored procedures

- creating 169, 357

## external\_udf\_execution\_mode option 27

## extfn\_get\_library\_version

- method 17

## extfn\_get\_license\_info 18, 19

## extfn\_use\_new\_api 97

## EXTFNAPIV4\_DESCRIBE\_COL\_CAN\_BE\_NULL

- get 209
- set 227

## EXTFNAPIV4\_DESCRIBE\_COL\_CONSTANT\_VALUE

- get 214
- set 230

## EXTFNAPIV4\_DESCRIBE\_COL\_DISTINCT\_VALUES

- set 210, 228

## EXTFNAPIV4\_DESCRIBE\_COL\_IS\_CONSTANT

- get 213
- set 230

## EXTFNAPIV4\_DESCRIBE\_COL\_IS\_UNIQUE

- get 212
- set 229

## EXTFNAPIV4\_DESCRIBE\_COL\_IS\_USED\_BY\_CONSUMER

- get 215
- set 231

## EXTFNAPIV4\_DESCRIBE\_COL\_MAXIMUM\_VALUE

- get 218
- set 235

## EXTFNAPIV4\_DESCRIBE\_COL\_MINIMUM\_VALUE

- get 216
- set 233

## EXTFNAPIV4\_DESCRIBE\_COL\_NAME

- set 206, 222

## EXTFNAPIV4\_DESCRIBE\_COL\_SCALE

- get 208
- set 226

## EXTFNAPIV4\_DESCRIBE\_COL\_TYPE

- get 207
- set 224

## EXTFNAPIV4\_DESCRIBE\_COL\_VALUES\_SUBSET\_OF\_INPUT

- get 220
- set 236

## EXTFNAPIV4\_DESCRIBE\_COL\_WIDTH

- set 207, 225

## EXTFNAPIV4\_DESCRIBE\_PARM\_CAN\_BE\_NULL

- get 244
- set 262



- EXTFNAPIV4\_DESCRIBE\_PARM\_CONSTANT\_VALUE
    - get 248
    - set 264
  - EXTFNAPIV4\_DESCRIBE\_PARM\_DISTINCT\_VALUES
    - get 246
    - set 263
  - EXTFNAPIV4\_DESCRIBE\_PARM\_IS\_CONSTANT
    - get 247
    - set 263
  - EXTFNAPIV4\_DESCRIBE\_PARM\_NAME
    - get 239
    - set 258
  - EXTFNAPIV4\_DESCRIBE\_PARM\_SCALE
    - get 242
    - set 261
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_HASS\_REWIND
    - get 255
    - set 271
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_COLUMNS
    - get 249
    - set 264
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_NUM\_ROWS
    - get 250
    - set 265
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_OR\_DERBY
    - get 251
    - set 266
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY 140, 141
    - get 252
    - set 268
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_PARTITIONBY UDF 143
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_REQUEST\_REWIND
    - get 254
    - set 269
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TABLE\_UNUSED\_COLUMNS
    - get 256
    - set 272
  - EXTFNAPIV4\_DESCRIBE\_PARM\_TYPE
    - get 240
    - set 259
  - EXTFNAPIV4\_DESCRIBE\_PARM\_WIDTH
    - get 241
    - set 260
  - EXTFNAPIV4\_DESCRIBE\_UDF\_NUM\_PARAMS
    - get 274
    - set 276
  - extfnapiv4.h 97
- F**
- fetch\_block
    - producing data 132
    - v4 API method 129, 130, 311
  - fetch\_into
    - producing data 131
    - v4 API method 129, 130, 308
  - FIRST
    - to return one row 189
  - FLOAT data type 9
  - free 135
  - FROM clause 183, 185, 191
    - SELECT statement 190
  - selects from stored procedure result sets 188
  - syntax 179
  - FROM contains-expression 183
  - from procedures 170
  - functions
    - callback 84
    - creating 172
    - external, prototypes 94
    - get\_piece 96
    - get\_value 96
    - GETUID 36
    - NUMBER 36
    - prototypes 95
    - user-defined 3
- G**
- g++
    - Linux 22
    - x86 22
  - get\_blob method
    - a\_v4\_extfn\_proc\_context 300
    - a\_v4\_extfn\_table\_context 314

## Index

get\_is\_cancelled method  
    a\_v4\_extfn\_proc\_context 293  
get\_option  
    v4 API method 296  
get\_value method  
    a\_v4\_extfn\_proc\_context 289  
get\_value\_is\_constant method  
    a\_v4\_extfn\_proc\_context 292  
GETUID function 36  
granting  
    execute permissions 29  
GROUP BY clause 36  
    SELECT statement 191

## H

HAVING clause 36  
heading name 189  
HP-UX  
    aCC 22  
    Itanium 22

## I

IGNORE NULL VALUES 34, 36  
impact of FROM clause 185  
initial state 121  
input argument  
    LONG BINARY 37, 45  
INSTALL JAVA statement  
    syntax 355  
installing the Java class code 339  
INT data type 9  
interface  
    dynamic library 15  
INTO clause  
    SELECT statement 189  
iq\_dummy 185  
iq\_dummy table 185  
IQ\_UDF license 3  
Itanium  
    aCC 22  
    HP-UX 22

## J

jar files  
    installing 355  
    removing 360

Java  
    installing classes 355  
    removing classes 360  
Java class  
    in multiplex 345  
    installing 339, 345, 346  
    removing 339  
Java external environment 325, 346, 349, 351, 352  
Java External Environment 339  
Java JAR  
    in multiplex 345  
    installing 339  
    removing 339  
Java method  
    calling 339  
Java table UDF 357  
    Creating 352  
Java UDF  
    creating 349, 351  
Java VM  
    setting location of 339  
    starting 339, 361  
    stopping 362  
JDBC API 325  
join columns 183  
joins  
    FROM clause syntax 179  
    SELECT statement 190

## L

library  
    dynamic interface 15  
    external 26  
    interface style 15  
library version  
    extfn\_get\_library\_version 17  
libv4apiex dynamic library 111, 114, 118, 160, 162,  
    166  
license  
    IQ\_UDF 3  
LIMIT keyword  
    SELECT statement 188  
link  
    switches 19, 21–24  
Linux  
    g++ 4.1.1 22  
    PowerPC 22  
    X86 22  
    xIC 22

LOB data type 9, 14  
 log files 28  
 log\_message method  
   a\_v4\_extfn\_proc\_context 295  
 LONG BINARY  
   input argument 37, 45  
 LONG BINARY data type 14  
 LONG BINARY(<n>) data type 9  
 LONG VARCHAR data type 14  
 LONG VARCHAR(<n>) data type 9

## M

memory tracking 135  
 moving window aggregate  
   OLAP-style optimized calling pattern 90  
   OLAP-style unoptimized calling pattern 89  
 moving window following aggregate  
   OLAP-style optimized calling pattern 92  
   OLAP-style unoptimized calling pattern 91  
 moving window without current  
   OLAP-style optimized calling pattern 93  
   OLAP-style unoptimized calling pattern 92  
 my\_bit\_or example  
   declaration 51  
   definition 73  
 my\_bit\_xor example  
   declaration 51  
   definition 70  
 my\_byte\_length example 37  
   declaration 37  
   definition 45  
 my\_interpolate example  
   declaration 52  
   definition 76  
 my\_plus example  
   declaration 34  
   definition 42  
 my\_plus\_counter example  
   declaration 36  
   definition 43  
 my\_sum example  
   declaration 50  
   definition 65

## N

naming conventions 8  
 new operator  
   C/C++ 31

NULL 34, 36, 43, 96  
 NUMBER function 36  
 NUMERIC(<precision>, <scale>) data type 14

## O

ODBC external environment 330  
 OLAP-style calling pattern  
   aggregate with unbounded window 87  
   optimized cumulative moving window  
     aggregate 90  
   optimized cumulative window aggregate 88  
   optimized moving window following  
     aggregate 92  
   optimized moving window without current 93  
   unoptimized cumulative moving window  
     aggregate 89  
   unoptimized cumulative window aggregate 88  
   unoptimized moving window following  
     aggregate 91  
   unoptimized moving window without current  
     92  
 ON clause 36  
 open\_result\_set  
   v4 API method 298  
 optimized calling pattern  
   OLAP-style cumulative window aggregate 88  
   OLAP-style moving window aggregate 90  
   OLAP-style moving window following  
     aggregate 92  
   OLAP-style moving window without current  
     93  
 optimizer estimate  
   a\_v4\_extfn\_estimate 302  
 options 185  
   unexpected behavior 191  
 order by 140  
 ORDER BY clause 47, 53, 193  
 order by list  
   a\_v4\_extfn\_orderby\_list 302  
 OVER clause 47, 53

## P

packages  
   installing 355  
   removing 360  
 Parallel TPF 140

## Index

- parameter type
  - a\_v4\_extfn\_describe\_parm\_type 278
- partition by
  - column number 303
- pattern
  - calling, aggregate 85
  - calling, scalar 85
- performance 185
- performance for joins 183
- Perl external environment 325
- PERL external environment 362
- permissions
  - granting 29
  - revoking 29
  - user-defined functions 29
- phases
  - query processing 121
- PHP external environment 325, 366
- plan building state 127
- PowerPC
  - AIX 21
  - Linux 22
  - xIC 22
  - xLC 21
- predicates
  - disjunction of 191
- procedures 170
  - replicating 167
  - select from result sets 188
- processing by SQL Anywhere 185
- processing queries without 185, 191
- producer 101
- projections
  - SELECT statement 188
- prototypes
  - external function 94

## Q

- queries 185
  - processing by SQL Anywhere 191
  - SELECT statement 185
- query optimization state 124
- query processing 121, 122, 124, 127, 128
- query processing states
  - annotation 283
  - execution 283
  - optimization 283
  - plan building 283
- querying tables 185, 191

## R

- REAL data type 9
- REMOVE JAVA 339
- REMOVE statement
  - syntax 360
- replication
  - of procedures 167
- RESPECT NULL VALUES 34, 36
- restrictions
  - C/C++ 31
- result sets 170
  - SELECT from 188
- return value
  - describe 280
- revoking
  - execute permissions 29
- rewind
  - v4 API method 313
- ROLLUP operator 192
  - SELECT statement 192
- row block 305
  - allocating 132
- row blocks
  - about 128
  - fetch methods for 129
  - producing data 131

## S

- samples
  - table UDF 99
  - TPF 99
- scalar functions
  - callback functions 84
  - context structure 40
  - creating user-defined function 38
  - declaring 33
  - defining 38
  - descriptor structure 39
  - my\_byte\_length example 37, 45
  - my\_plus example 34, 42
  - my\_plus\_counter example 36, 43
- security
  - user-defined functions 25
- SELECT INTO
  - returning results in a base table 188
  - returning results in a host variable 188

- returning results in a temporary table 188
- select list
  - SELECT statement 189
- SELECT statement
  - FIRST 189
  - FROM clause syntax 179
  - LIMIT 188
  - syntax 185
  - TOP 189
- server
  - disabling UDFs 25
  - enabling UDFs 25
- SET clause 36
- set\_cannot\_be\_distributed
  - v4 API method 301
- set\_error method
  - a\_v4\_extfn\_proc\_context 294
- set\_value method
  - a\_v4\_extfn\_proc\_context 292
- shared libraries
  - building 19, 21–24
- simple aggregate grouped
  - calling pattern 86
- simple aggregate ungrouped
  - calling pattern 86
- Solaris
  - SPARC 23
  - Sun Studio 12 23
  - X86 23
- SPARC
  - Solaris 23
  - Sun Studio 12 23
- SQL statements 166
- START EXTERNAL ENVIRONMENT JAVA 339
- START JAVA statement
  - syntax 361
- starting
  - Java VM 361
- states
  - annotation 122
  - execution 128
  - initial 121
  - plan building 127
  - query optimization 124
  - query processing 121
- STOP JAVA statement
  - syntax 362
- stopping
  - Java VM 362
- stored procedures
  - selecting into result sets 188
- structure
  - a\_v4\_extfn\_blob 195
  - a\_v4\_extfn\_blob\_istream 199
  - a\_v4\_extfn\_col\_subset\_of\_input 203
  - a\_v4\_extfn\_column\_data 200
  - a\_v4\_extfn\_column\_list 202
  - a\_v4\_extfn\_estimate 302
  - a\_v4\_extfn\_order\_el 202
  - a\_v4\_extfn\_orderby\_list 302
  - a\_v4\_extfn\_proc 284
  - a\_v4\_extfn\_proc\_context 287
  - a\_v4\_extfn\_table 306
  - a\_v4\_extfn\_table\_context 306
  - a\_v4\_extfn\_table\_func 314
  - aggregate context 62
  - aggregate descriptor 57
  - scalar context 40
  - scalar descriptor 39
- Studio 12
  - See Sun Studio 12
- subqueries
  - disjunction of 191
- Sun Studio 12
  - Solaris 23
  - SPARC 23
  - x86 23
- switches
  - compile 19, 21–24
  - link 19, 21–24
- Sybase IQ
  - description 1
- syntax
  - aggregate context 62
  - aggregate declaration 47
  - aggregate definition 54
  - aggregate description 57
  - API version 94
  - calculation context 61
  - calling user-defined functions 83
  - CREATE FUNCTION statement 83
  - disabling user-defined functions 25
  - dropping user-defined functions 29
  - dynamic library interface 15
  - enabling user-defined functions 25
  - function prototypes 95
  - scalar context 40
  - scalar declaration 33

## Index

- scalar definition 38
- scalar description 39
- SYSTEM dbspace 185, 191
- system tables 185

## T

### table

- a\_v4\_extfn\_table 306
- temporary 358

### table context

- a\_v4\_extfn\_table\_context 306
- fetch\_block method 130, 311
- fetch\_into method 130, 308
- rewind method 313

### TABLE data type 9

### table functions

- \_close\_extfn method 319
- \_fetch\_block\_extfn method 318
- \_fetch\_into\_extfn method 317
- \_open\_extfn method 316
- \_rewind\_extfn method 318
- a\_v4\_extfn\_table\_func 314

### Table parameterized function

- definition 136

### table UDF

- creation steps 103
- definition 97
- developing 97, 103
- example udf\_rg\_2 114
- examples 105
- restrictions 99
- sample udf\_rg\_1 106
- sample udf\_rg\_2 111
- sample udf\_rg\_3 115
- samples directory udf\_rg\_1.cxx 106
- samples directory udf\_rg\_2.cxx 111, 114
- samples directory udf\_rg\_3.cxx 115
- users 97, 98

### Table UDF

- example udf\_rg\_1 111
- example udf\_rg\_3 118
- samples directory udf\_rg\_1.cxx 111
- samples directory udf\_rg\_3.cxx 118

### TABLE\_UDF\_ROW\_BLOCK\_SIZE\_KB Option 179

### tables 185

### temporary table 358

### temporary tables

- populating 190

### testing 25

### text search 183

### TIME data type 9

### TINYINT data type 9

### TOP

- specify number of rows 189

### TPF

- definition 136
- developing 97, 136
- example tpf\_blob 166
- example tpf\_rg\_1 160
- example tpf\_rg\_2 162
- restrictions 99
- samples directory tpf\_blob.cxx 166
- samples directory tpf\_rg\_1.cxx 160
- samples directory tpf\_rg\_2.cxx 162
- users 98

### tpf\_blob.cxx

- running the TPF 166

### tpf\_rg\_1.cxx

- running the TPF 160
- TPF samples 160

### tpf\_rg\_2.cxx

- running the TPF 162
- TPF samples 162

### tptf\_blob.cxx

- TPF samples 166

## U

### UDF

- See user-defined functions

### udf\_proc\_describe 97

### udf\_proc\_evaluate 97

### udf\_proc\_version 97

### udf\_rg\_1.cxx

- running the Table UDF 111
- table UDF sample 1 106
- Table UDF samples 111

### udf\_rg\_2.cxx

- running the table UDF 114
- table UDF sample 2 111
- Table UDF samples 114

### udf\_rg\_3.cxx

- running the Table UDF 118
- table UDF sample 3 115
- Table UDF samples 118

### unbounded window

- OLAP-style aggregate calling pattern 87

### unexpected behavior 185

- unloading
    - external library 26
  - unoptimized calling pattern
    - OLAP-style cumulative window aggregate 88
    - OLAP-style moving window aggregate 89
    - OLAP-style moving window following aggregate 91
    - OLAP-style moving window without current 92
  - UNSIGNED data type 9
  - UNSIGNED INT data type 9
  - UPDATE statement 36
  - user-defined functions 25, 37
    - aggregate, creating 53
    - callback functions 84
    - calling 83
    - calling non-existent UDF 323
    - calling pattern, aggregate 85
    - calling pattern, scalar 85
    - creating 31, 32
    - debugging 28
    - disabling 25
    - dropping 29
    - enabling 3, 25
    - error 323
    - execution permissions 29
    - my\_bit\_or example 51, 73
    - my\_bit\_xor example 51, 70
    - my\_byte\_length example 45
    - my\_interpolate example 52, 76
    - my\_plus example 34, 42
    - my\_plus\_counter example 36, 43
    - my\_sum example 50, 65
    - scalar, creating 38
    - security 25
    - using 3
- V**
- v4 API
    - \_close\_extfn method 319
    - \_fetch\_block\_extfn method 318
    - \_fetch\_into\_extfn method 317
    - \_open\_extfn method 316
    - \_rewind\_extfn method 318
    - alloc method 297, 298
    - backward-compatibility 16
    - close\_result\_set method 299
    - fetch\_block method 130, 311
    - fetch\_into method 130, 308
    - get\_option method 296
    - open\_result\_set method 298
    - rewind method 313
    - set\_cannot\_be\_distributed method 301
  - v4\_extfn\_partitionby\_col\_num 141
  - VARBINARY(<n>) data type 9
  - VARCHAR(<n>) data type 9
  - variable 170
  - variable result sets 170
  - variables
    - select into 190
  - version
    - declaring for API 94
  - Visual Studio
    - debugging UDFs 28
  - Visual Studio 2009
    - Windows 24
    - x86 24
- W**
- WHERE clause 36
    - SELECT statement 191
  - WINDOW FRAME clause 53
  - Windows
    - Visual Studio 2009 24
    - X86 24
- X**
- x86
    - g++ 22
    - Linux 22
    - Solaris 23
    - Sun Studio 12 23
    - Visual Studio 2009 24
    - Windows 24
  - xIC
    - Linux 22
    - PowerPC 22
  - xIC
    - AIX 21
    - PowerPC 21

