

Informix Product Family
Informix
Version 12.10

*IBM Informix Change Data Capture API
Programmer's Guide*



Informix Product Family
Informix
Version 12.10

*IBM Informix Change Data Capture API
Programmer's Guide*



Note

Before using this information and the product it supports, read the information in "Notices" on page B-1.

Edition

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corporation 2008, 2013.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	v
About this publication	v
Types of users	v
Example code conventions.	v
Additional documentation	vi
Compliance with industry standards	vi
Syntax diagrams	vi
How to read a command-line syntax diagram.	vii
Keywords and punctuation	viii
Identifiers and names.	viii
How to provide documentation feedback	ix
Chapter 1. Getting started with the Change Data Capture API	1-1
The Change Data Capture API.	1-1
Change Data Capture API components	1-1
Smart large object read functions	1-2
CDC record sequence numbers	1-4
Data for capture	1-5
Preparing to use the Change Data Capture API	1-5
Writing an application to capture data changes	1-5
Handling errors.	1-6
Handling smart large objects	1-7
Restarting data capture	1-7
Monitoring data capture	1-8
Chapter 2. Change Data Capture functions	2-1
The cdc_activatesess() function.	2-1
The cdc_closesess() function	2-2
The cdc_deactivatesess() function	2-2
The cdc_endcapture() function	2-3
The cdc_errortext() function.	2-4
The cdc_opensess() function	2-5
The cdc_recboundary() function	2-7
The cdc_set_fullrowlogging() function	2-7
The cdc_startcapture() function	2-8
Chapter 3. Change Data Capture records	3-1
Format of CDC records	3-1
The CDC_REC_BEGINTX record	3-2
The CDC_REC_COMMTX record	3-2
The CDC_REC_DELETE record	3-3
The CDC_REC_DISCARD record	3-4
The CDC_REC_ERROR record	3-4
The CDC_REC_INSERT record.	3-5
The CDC_REC_RBTX record	3-6
The CDC_REC_TABSCHEMA record	3-6
The CDC_REC_TIMEOUT record	3-7
The CDC_REC_TRUNCATE record	3-8
The CDC_REC_UPDAFT record	3-8
The CDC_REC_UPDBEF record	3-9
Chapter 4. The syscdc system database	4-1
The syscdccerrcodes table.	4-1
The syscdcrectypes table.	4-1

Chapter 5. Change Data Capture error codes	5-1
Chapter 6. onstat -g cdc	6-1
Chapter 7. Change Data Capture sample program	7-1
Appendix. Accessibility	A-1
Accessibility features for IBM Informix products	A-1
Accessibility features	A-1
Keyboard navigation	A-1
Related accessibility information	A-1
IBM and accessibility	A-1
Dotted decimal syntax diagrams	A-1
Notices	B-1
Trademarks	B-3
Index	X-1

Introduction

About this publication

This publication describes the IBM® Informix® Change Data Capture API and the concepts of capturing changes to data. This publication describes how to use the Change Data Capture API to write an application that captures changed data for external processing.

Types of users

This publication is for database application programmers.

To understand this publication, you must have the following background:

- A working knowledge of your computer, your operating system, and the utilities that your operating system provides
- Some experience working with relational databases or exposure to database concepts
- Some experience with computer programming in the C or Java™ programming language

Example code conventions

Examples of SQL code occur throughout this publication. Except as noted, the code is not specific to any single IBM Informix application development tool.

If only SQL statements are listed in the example, they are not delimited by semicolons. For instance, you might see the code in the following example:

```
CONNECT TO stores_demo
...

DELETE FROM customer
  WHERE customer_num = 121
...

COMMIT WORK
DISCONNECT CURRENT
```

To use this SQL code for a specific product, you must apply the syntax rules for that product. For example, if you are using an SQL API, you must use EXEC SQL at the start of each statement and a semicolon (or other appropriate delimiter) at the end of the statement. If you are using DB–Access, you must delimit multiple statements with semicolons.

Tip: Ellipsis points in a code example indicate that more code would be added in a full application, but it is not necessary to show it to describe the concept being discussed.

For detailed directions on using SQL statements for a particular application development tool or SQL API, see the documentation for your product.

Additional documentation

Documentation about this release of IBM Informix products is available in various formats.

You can access Informix technical information such as information centers, technotes, white papers, and IBM Redbooks® publications online at <http://www.ibm.com/software/data/sw-library/>.

Compliance with industry standards

IBM Informix products are compliant with various standards.

IBM Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992. In addition, many features of IBM Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL Common Applications Environment (CAE) standards.

The IBM Informix Geodetic DataBlade® Module supports a subset of the data types from the *Spatial Data Transfer Standard (SDTS)—Federal Information Processing Standard 173*, as referenced by the document *Content Standard for Geospatial Metadata*, Federal Geographic Data Committee, June 8, 1994 (FGDC Metadata Standard).

Syntax diagrams

Syntax diagrams use special components to describe the syntax for statements and commands.

Table 1. Syntax Diagram Components

Component represented in PDF	Component represented in HTML	Meaning
	>>-----	Statement begins.
	----->	Statement continues on next line.
	>-----	Statement continues from previous line.
	-----<<	Statement ends.
	-----SELECT-----	Required item.
	+-----+ '-----LOCAL-----'	Optional item.
	---+-----ALL-----+--- +---DISTINCT-----+ '---UNIQUE-----'	Required item with choice. Only one item must be present.

Table 1. Syntax Diagram Components (continued)


Component represented in PDF	Component represented in HTML	Meaning
	<pre> -----+----- +--FOR UPDATE-----+ '--FOR READ ONLY--'</pre>	Optional items with choice are shown below the main line, one of which you might specify.
	<pre> .---NEXT----- -----+----- +---PRIOR-----+ '---PREVIOUS-----'</pre>	The values below the main line are optional, one of which you might specify. If you do not specify an item, the value above the line is used by default.
	<pre> -----+----- v +---index_name---+ '---table_name---'</pre>	Optional items. Several items are allowed; a comma must precede each repetition.
	<pre>>>- Table Reference -><</pre>	Reference to a syntax segment.
<p>Table Reference</p>	<pre> Table Reference --+---view-----+-- +---table-----+ '---synonym-----'</pre>	Syntax segment.

How to read a command-line syntax diagram

Command-line syntax diagrams use similar elements to those of other syntax diagrams.

Some of the elements are listed in the table in Syntax Diagrams.

Creating a no-conversion job




Notes:

1 See page Z-1

This diagram has a segment named “Setting the Run Mode,” which according to the diagram footnote is on page Z-1. If this was an actual cross-reference, you

would find this segment on the first page of Appendix Z. Instead, this segment is shown in the following segment diagram. Notice that the diagram uses segment start and end components.

Setting the run mode:



To see how to construct a command correctly, start at the upper left of the main diagram. Follow the diagram to the right, including the elements that you want. The elements in this diagram are case-sensitive because they illustrate utility syntax. Other types of syntax, such as SQL, are not case-sensitive.

The Creating a No-Conversion Job diagram illustrates the following steps:

1. Type **onpladm create job** and then the name of the job.
2. Optionally, type **-p** and then the name of the project.
3. Type the following required elements:
 - **-n**
 - **-d** and the name of the device
 - **-D** and the name of the database
 - **-t** and the name of the table
4. Optionally, you can choose one or more of the following elements and repeat them an arbitrary number of times:
 - **-S** and the server name
 - **-T** and the target server name
 - The run mode. To set the run mode, follow the Setting the Run Mode segment diagram to type **-f**, optionally type **d**, **p**, or **a**, and then optionally type **l** or **u**.
5. Follow the diagram to the terminator.

Keywords and punctuation

Keywords are words reserved for statements and all commands except system-level commands.

When a keyword appears in a syntax diagram, it is shown in uppercase letters. When you use a keyword in a command, you can write it in uppercase or lowercase letters, but you must spell the keyword exactly as it appears in the syntax diagram.

You must also use any punctuation in your statements and commands exactly as shown in the syntax diagrams.

Identifiers and names

Variables serve as placeholders for identifiers and names in the syntax diagrams and examples.

You can replace a variable with an arbitrary name, identifier, or literal, depending on the context. Variables are also used to represent complex syntax elements that are expanded in additional syntax diagrams. When a variable appears in a syntax diagram, an example, or text, it is shown in *lowercase italic*.

The following syntax diagram uses variables to illustrate the general form of a simple SELECT statement.

►—SELECT—*column_name*—FROM—*table_name*—◄

When you write a SELECT statement of this form, you replace the variables *column_name* and *table_name* with the name of a specific column and table.

How to provide documentation feedback

You are encouraged to send your comments about IBM Informix user documentation.

Use one of the following methods:

- Send email to docinf@us.ibm.com.
- In the Informix information center, which is available online at <http://www.ibm.com/software/data/sw-library/>, open the topic that you want to comment on. Click the feedback link at the bottom of the page, fill out the form, and submit your feedback.
- Add comments to topics directly in the information center and read comments that were added by other users. Share information about the product documentation, participate in discussions with other users, rate topics, and more!

Feedback from all methods is monitored by the team that maintains the user documentation. The feedback methods are reserved for reporting errors and omissions in the documentation. For immediate help with a technical problem, contact IBM Technical Support at <http://www.ibm.com/planetwide/>.

We appreciate your suggestions.

Chapter 1. Getting started with the Change Data Capture API

These topics describe the Change Data Capture API and how to use it.

The Change Data Capture API

The Change Data Capture API allows external client applications to capture transactional data from an IBM Informix database.

The Change Data Capture API provides functions to capture transactional data. You can use a variety of clients to run these functions, such as, JDBC, ODBC, ESQL/C, and DB-Access. The data is returned as CDC records by standard IBM Informix smart large object read functions. How the captured data is processed depends on your application. For example, you can write an application to replicate data from an IBM Informix database to another, heterogeneous, database.

The following types of operations are captured:

- INSERT
- DELETE
- UPDATE
- TRUNCATE

The Change Data Capture API starts capturing transactions from the current logical log and processes all transactions sequentially. The first time you start capturing data for a particular table, data capture starts at the current log position. If you later stop capture and the restart it, you can restart at the point in the logical logs where data capture was stopped. You cannot go backwards in time through the logical logs to capture the history of the table or perform random seeking in the logical logs.

At the beginning of data capture for a table, the Change Data Capture API provides the table schema information that you can use in your application to create a target table. However, any changes to the table schema after data capture begins are not captured by the Change Data Capture API.

The Change Data Capture API can only provide data as that data is changing; it does not provide an initial snapshot of the contents of the table. If you need a populated target table, you can externally load the existing data to the target table. Alternatively, you can create dummy updates to the table for each row so that the Change Data Capture API can capture those updates and populate the target table.

The Change Data Capture API does not capture changes to table schemas or any other database changes.

The Change Data Capture API can capture data only from databases that have logging enabled.

Change Data Capture API components

The Change Data Capture API consists of functions, a system database, CDC records, and error codes.

Functions

Change Data Capture functions are built-in SQL functions that you run by using the EXECUTE FUNCTION statement. You use these functions to control data capture. The `cdc_opensess()` function returns the CDC session ID, which is a smart large object file descriptor that you use to retrieve captured data. The `cdc_startcapture()` function specifies the table from which to capture data. Other functions specify to start or end data capture.

You must call Change Data Capture functions from a client application. You cannot call this function from a user-defined routine that runs within the database server.

System database

The `syscdc` system database contains the Change Data Capture functions and system tables. The system tables store information about Change Data Capture API error codes and record types.

Error codes

The Change Data Capture API functions return error codes. Most of the functions return an error code both if they succeed or fail. The Change Data Capture API error codes are listed in the `syscdcerrcodes` table of the `syscdc` database. You can query the `syscdcerrcodes` table to determine whether the function failed and if so, why it failed.

Smart large object read functions

You use smart large object read functions to read the captured data, by passing the smart large object file descriptor provided by the `cdc_opensess()` function. Smart large object read functions are not part of the Change Data Capture API; you can use smart large object read functions such as `mi_lo_read()` or `ifx_lo_read()`.

CDC Records

The Change Data Capture (CDC) records are returned by smart large object read functions and provide information about the transaction currently being captured as well as the actual captured data.

Smart large object read functions

You use smart large object read functions to transfer captured data to a buffer where your application can access it.

You can use any of the smart large object read functions listed in the following table, depending on your application language. You must use the same smart large object read function for all read calls during a particular session. Using different functions in the same session can result in incomplete delivery of captured data.

Table 1-1. Smart large object read functions

Read function	Arguments	Informix API	Application language
mi_lo_read()	A pointer to a connection descriptor A smart large object file descriptor A data buffer The maximum number of bytes to read	DataBlade API	Use in a C language application.
ifx_lo_read()	A smart large object file descriptor A data buffer	ODBC	Use in an ODBC application.
ifx_lo_read()	A smart large object file descriptor A data buffer The maximum number of bytes to read A pointer to an error code	ESQL/C	Use in a C language application.
IfxLoRead()	A smart large object file descriptor A data buffer The maximum number of bytes to read	JDBC	Use in a Java application.
IfxBlob.Read()	A data buffer	.NET	Use in a .NET application.

Read timeout

If no captured data is available to retrieve, the read call waits for data for the timeout period specified by the **cdc_opensess()** function. If the timeout period is exceeded, a CDC_REC_TIMEOUT record is returned to the read call. The read call passes the CDC_REC_TIMEOUT record into the data buffer and returns successfully.

Read buffer size

The size of the buffer specified in the read call should be at least 128 bytes. The maximum size of a read buffer is 2 GB. You can calculate the approximate minimum size of the buffer for your application by calculating the largest possible CDC record size, for example, a CDC_REC_INSERT record, and multiplying that value times the maximum number of records to return per read call that you specify in the **cdc_opensess()** function.

The amount of data returned by a read call is the lesser of the size of the buffer specified in the read call and the maximum number of records to return. No more than the maximum number of records is returned by one read call, even if the number of bytes contained in those records is less than the maximum number of bytes allowed by the read call. However, no more than the maximum number of bytes allowed by the read call will be returned, even if the number of records returned is less than the maximum number allowed. If a record does not fit into a buffer, as much of the record as can fit is returned, and subsequent read calls return the rest.

Smart large object file descriptor

The value for the smart large object file descriptor argument in the read functions is the CDC session ID returned by the `cdc_opensess()` function.

Smart large object read function for the IBM Informix .NET Provider

The smart large object read function for .NET works differently than for other client APIs. The following pseudo code illustrates the basic structure for reading smart large objects with .NET:

```
conn = new IfxConnection(..)// to SYSCDC database
execute function informix.cdc_opensess() // on the same connection
IfxBlob( IfxConnection connection )// construct it using the same connection
IfxBlob.Open(ReadOnly) // open it
IfxBlob.Read(char[] buff)
```

Related concepts:

[IfxBlob class \(.NET Provider Guide\)](#)

Related tasks:

[Accessing a smart large object \(DataBlade API Guide\)](#)

“Handling smart large objects” on page 1-7

Related reference:

[Smart large object examples \(JDBC Driver Guide\)](#)

[Smart large objects \(ESQL/C Guide\)](#)

CDC record sequence numbers

Most Change Data Capture (CDC) records returned to the client contain a sequence number.

The sequence number associated with a CDC record is a BIGINT data type.

The CDC record sequence number is not necessarily the same as the LSN of the IBM Informix logical log that is being captured.

You can compare sequence numbers for CDC records that are returned for the same transaction. Within a transaction, the sequence numbers of CDC records returned increase over time. A lower sequence number indicates that the CDC record was returned earlier than a CDC record with a higher sequence number.

You can compare the sequence numbers of CDC_REC_BEGINTX records or the sequence numbers of CDC_REC_COMMTX records for different transactions. Each committed transaction has one CDC_REC_BEGINTX record and one CDC_REC_COMMTX record. The sequence numbers for the CDC_REC_BEGINTX

and CDC_REC_COMMTX records are in monotonic order. A lower sequence number indicates that the associated transaction was begun or committed earlier than a transaction associated with a higher sequence number.

Data for capture

You can capture most IBM Informix data types. You can specify the data to capture at the column level.

The following data types are not supported for data capture:

- Simple large objects (TEXT and BYTE data types)
- User-defined data types
- Collection data types (SET, MULTISSET, LIST, and ROW data types)

Specifying what data to capture

You specify a table and which columns from that table to capture with the **cdc_startcapture()** function. You must run the **cdc_startcapture()** function once for each table that you want to capture. For information about which tables and columns are currently being captured, look in the **syscdctabs** table.

Ending capture of a table

To stop data capture of a specific table, run the **cdc_endcapture()** function. After you run **cdc_endcapture()** function, information about that table is removed from the **syscdctabs** table.

Preparing to use the Change Data Capture API

Before you can start using the Change Data Capture API, you must prepare the database and the database server.

Perform the following tasks to prepare for using the Change Data Capture API:

1. Turn on logging for all databases from which you intend to capture data changes. For information about logging, see the *IBM Informix Guide to SQL: Syntax*.
2. Run the following script as user **informix** from the \$INFORMIXDIR/etc directory: **syscdcv1.sql**
3. Verify that the **syscdcv1** database exists by creating a connection to it, as user **informix**. For example, you can use DB-Access to connect to the **syscdcv1** database.
4. Set the **DB_LOCALE** environment variable to be the same as the locale of the database from which you want to capture data.

Related tasks:

“Writing an application to capture data changes”

Writing an application to capture data changes

Use the Change Data Capture functions to control the data capture process. Process CDC records to extract the data. Query **syscdc** tables to retrieve the symbolic names and descriptions of CDC records and errors.

Complete the prerequisite tasks to prepare for using the Change Data Capture API.

Your application should contain the following structures and functions:

- A structure to store table schema information. You use the table schema to parse the column data.
- A function to interpret the table schema information and populate the table schema structure. You can obtain the table schema information from the CDC_REC_TABSCHEMA record.
- A function to retrieve and parse the column values from the data buffer.
- A function to handle errors. You can query the **syscdccerrcodes** table to determine the symbolic name and description of the error code.

Include the following tasks in your application to capture data changes:

1. As user **informix**, connect to the **syscdcv1** database on the database server to which the client is currently connected.
2. Open a capture session by running the **cdc_opensess()** function. The **cdc_opensess()** function returns a session ID.
3. Enable full-row logging for each table from which you want to capture data by running the **cdc_set_fullrowlogging()** function.
4. Specify which data to capture by running the **cdc_startcapture()** function. Run this function for each table from which you want to capture data.
5. Start the capture process by running the **cdc_activatesess()** function. CDC records, including those that contain captured data, are returned to the application.
6. Read the CDC records containing captured data with a smart large object read function such as **mi_lo_read()** by passing the session ID as the large object file descriptor. Use the same smart large object read function for all read calls.
7. Parse the data by column values. If you are writing your application in Java, you can use the **IfxToJavaType** class to convert a byte stream of an Informix representation of a data type to the appropriate Java data type and value.
8. Stop capturing data by running the **cdc_endcapture()** function for each table.
9. Disable full-row logging by running the **cdc_set_fullrowlogging()** function for each table. Make sure that no other applications or processes are dependent on full-row logging before you disable it.
10. Close the capture session by running the **cdc_closesess()** function.

Related concepts:

Chapter 2, “Change Data Capture functions,” on page 2-1

Related tasks:

“Preparing to use the Change Data Capture API” on page 1-5

Related reference:

Chapter 7, “Change Data Capture sample program,” on page 7-1

Handling errors

To process errors that are returned by Change Data Capture functions, reference error numbers by looking up their symbolic names in the **syscdccerrcodes** table.

Add code to your application to handle possible error conditions.

1. Declare error code variables for the types of errors that you intend to process separately.
2. Query the **syscdccerrcodes** table to find the error number corresponding to each of the symbolic names of the Change Data Capture error codes.
3. Set the error code variables to the Change Data Capture error numbers.

4. Add code to handle each error condition.

You can use the `cdc_errortext()` function to return the error text for a specified symbolic name.

Related reference:

Chapter 5, “Change Data Capture error codes,” on page 5-1

Handling smart large objects

The Change Data Capture API does not directly support the retrieval of smart large object column data from a captured BLOB or CLOB row. You must use the DataBlade API or client API smart large object read functions to retrieve smart large objects.

To retrieve the data in a smart large object column, follow these general steps:

1. Retrieve the data row that contains the smart large object with the Change Data Capture API.
2. Extract the values of columns that uniquely identify the data row, such as the primary key or a unique constraint.
3. Run an SQL SELECT statement with the identifying values to retrieve the data row.
4. Open the smart large object from the column in the data row.
5. Retrieve the smart large object data by using one of the following types of smart large object functions:
 - DataBlade API functions, such as `mi_lo_read()`, `mi_lo_to_buffer()`, or `mi_lo_to_file()`
 - SQL functions such as `LOTOFILE()`
 - ESQ/C functions such as `ifx_lo_read()`
6. Close the smart large object.

Related concepts:

 [IfxBlob class \(.NET Provider Guide\)](#)

“Smart large object read functions” on page 1-2

Related tasks:

 [Accessing a smart large object \(DataBlade API Guide\)](#)

Related reference:

 [Smart large object examples \(JDBC Driver Guide\)](#)

 [Smart large objects \(ESQ/C Guide\)](#)

Restarting data capture

You can restart data capture where the last data capture session ended.

The restart position is the sequence number of a CDC record that was returned in the previous data capture session. You can use the sequence number of the last CDC record processed in the previous data capture session. However, to preserve transactional integrity, you should determine last transaction for which a commit or rollback operation was not processed and restart capture at the beginning of that transaction. In this case, the restart position is the lowest sequence number of the `CDC_REC_BEGINTX` records for incomplete transactions. To avoid reprocessing already committed transactions, you should also determine the largest

sequence number of the CDC_REC_COMMTX record that you have already processed in a previous data capture session.

To restart data capture:

1. Determine the restart position. To preserve transactional integrity:
 - a. Find all captured transactions that did not return a CDC_REC_COMMTX or CDC_REC_RBTX record.
 - b. Compare the sequence numbers of the CDC_REC_BEGINTX records for the incomplete transactions. The lowest sequence number is the restart position.
2. Open a new capture session by running the **cdc_opensess()** function.
3. Run the **cdc_startcapture()** function for the table on which you want to restart capturing data.
4. Active the session by running the **cdc_activatesess()** function. Pass the appropriate sequence number as the *position* argument. Data capture restarts for the table at the last transaction that was processed.
5. Discard any transactions whose CDC_REC_COMMTX sequence number is less than that of the CDC_REC_COMMTX record with the largest sequence number that you processed in the previous data capture session.

Related reference:

“The **cdc_activatesess()** function” on page 2-1

Monitoring data capture

You can monitor the status of data capture by running the **onstat -g cdc** command.

To view the current status of a data capture session, run the **onstat -g cdc** command. For this command, and all other **onstat -g cdc** command options, you can specify a single session or view information about all current sessions.

To view the status of session buffers, run the **onstat -g cdc bufm** command.

To view information about session configuration, run the **onstat -g cdc config** command.

To view information about tables currently being captured, run the **onstat -g cdc table** command. You can provide a single table name or view information for all tables.

Related reference:

Chapter 6, “onstat -g cdc,” on page 6-1

Chapter 2. Change Data Capture functions

These topics describe the Change Data Capture functions.

Related tasks:

“Writing an application to capture data changes” on page 1-5

The `cdc_activatesess()` function

For an open capture session, starts capturing data from the specified log and log position.

The `syscdcsess` table is updated when the session is activated.

Syntax

```
►► cdc_activatesess—(—session_ID—, —position—)—►►
```

Function arguments

Table 2-1. The `cdc_activatesess()` arguments

Argument	Data Type	Description
<code>session_ID</code>	INTEGER	The session ID of the open capture session for which to start capturing data.
<code>position</code>	BIGINT	Must be 0 or the restart position.

Usage

After you open a session with the `cdc_opensess()` function, you use the `cdc_activatesess()` function to start capturing data at the specified log position. If you are starting data capture on a table for the first time, the `position` must be 0. If you have previously performed data capture, you can restart data capture where it left off by specifying a sequence number of a CDC record returned in the previous capture session.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns 0.

If unsuccessful, returns an integer corresponding to an error code and updates the `syscdcsess` table with the error information.

Related tasks:

“Restarting data capture” on page 1-7

The `cdc_closesess()` function

Closes a capture session that is associated with the specified session ID.

Any resources used by the capture session are released. The rows in the `syscdctabs` and `sysdcscs` tables containing the specified session ID are deleted.

Syntax

```
►► cdc_closesess(—session_ID—)◄◄
```

Function argument

Table 2-2. The `cdc_closesess()` argument

Argument	Data Type	Description
<i>session_ID</i>	INTEGER	The session ID of the capture session that you want to close.

Usage

Use the `cdc_closesess()` function to close a capture session that you no longer need. If the capture session was active, all data capture is immediately stopped when the session is closed.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns 0.

If unsuccessful, returns an integer corresponding to an error code and updates the `sysdcscs` table with the error information.

The `cdc_deactivatesess()` function

Stops capturing data for an active capture session.

The `sysdcscs` table is update to show that the capture session is not active.

Syntax

```
►► cdc_deactivatesess(—session_ID—)◄◄
```

Function argument

Table 2-3. The `cdc_deactivatesess()` argument

Argument	Data Type	Description
<code>session_ID</code>	INTEGER	The session ID of the capture session that you want to deactivate.

Usage

Use the `cdc_deactivatesess()` function to stop capturing data for a specific capture session.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns 0.

If unsuccessful, returns an integer corresponding to an error code and updates the `syscdcsess` table with the error information.

The `cdc_endcapture()` function

Ends capture for a specified table.

The row in the `syscdctabs` table associated with the specified session ID and table is deleted.

Purpose

►► `cdc_endcapture` (`—session_ID—`, `—MBZ—`, `—"database:owner.table_name"—`) ◀◀

Function arguments

Table 2-4. The `cdc_endcapture()` arguments

Argument	Data Type	Description
<code>session_ID</code>	INTEGER	The session ID of an open capture session.
<code>MBZ</code>	BIGINT	Must be 0. Reserved.

Table 2-4. The `cdc_endcapture()` arguments (continued)

Argument	Data Type	Description
<code>database:owner.table_name</code>	LVARCHAR	<p>The qualified name of the table from which to capture data. The qualified name includes the following elements:</p> <p><i>database</i> The name of the database in which the table exists.</p> <p><i>owner</i> The name of the owner of the table.</p> <p><i>table</i> The name of the table</p>

Usage

Use the `cdc_endcapture()` function to stop capturing data from a specific table. This function does not affect the session status; the session remains open and active.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns 0.

If unsuccessful, returns an integer corresponding to an error code and updates the `syscdcsess` table with the error information.

The `cdc_errortext()` function

Returns the error message text corresponding to the specified symbolic error name.

Symbolic error names are listed in the `syscdccerrcodes` table in the `syscdc` database.

Syntax

►► `cdc_errortext`—(`—'`*error_name*`—'`, `—'`*locale_name*`—'`)—►►

Function arguments

Table 2-5. The `cdc_errortext()` arguments

Argument	Data Type	Description
<code>error_name</code>	LVARCHAR	The symbolic name of the error.
<code>locale_name</code>	LVARCHAR	The name of the locale in which to display the error text. If locale name parameter is SQL NULL or a string of 0 length ("") the default locale is used.

Usage

Use the `cdc_errortext()` function to return the error text for an error that you received from another CDC function. Not all error texts are available in all locales. If the `cdc_errortext()` function does not return the text in the locale you specified, try to run the function again with a different locale, such as 'en_us.819' or 'en_us.033'.

Return values

If successful, returns SQLCODE 0 and the error message text.

If unsuccessful, returns with a nonzero SQLCODE:

- 23109: Invalid locale specification.
The locale name is not correct or the specified locale was not found.
- 1824: Message cannot be found.
The locale is valid but the message was not found in the message file for that locale. Specify a different locale, such as `en_us.033`.
- Other SQLCODES represent internal errors.

Example

The following example returns the error text for the error `CDC_E_TABCAPTURED` in the `en_us.033` locale:

```
> select cdc_errortext('CDC_E_TABCAPTURED', 'en_us.033') from syscdcvrs;
```

(expression) The specified table is already being captured by the CDC session.

1 row(s) retrieved.

Related reference:

Chapter 5, “Change Data Capture error codes,” on page 5-1

The `cdc_opensess()` function

Opens a capture session and creates a session ID.

A row is inserted into the `syscdcsess` table for the session.

Syntax

```
►► cdc_opensess(—"server_name"—,—"session_ID"—,—"timeout"—,—————►  
►—"max_recs"—,—"major_version"—,—"minor_version"—)—————►◄
```

Function arguments

Table 2-6. The `cdc_opensess()` arguments

Argument	Data Type	Description
<i>server_name</i>	LVARCHAR	The name of the server. Must be the server to which the client application that is calling the <code>cdc_opensess()</code> function is connected.
<i>session_ID</i>	INTEGER	Must be 0.
<i>timeout</i>	INTEGER	Specifies the timeout behavior of a read call on the captured data: <0 Do not timeout. 0 Return immediately if no data is available. 1 or more The number of seconds to wait for data before timing out.
<i>max_recs</i>	INTEGER	The maximum number of CDC records to return per read function call. This value takes precedence over the maximum number of bytes to return that is specified in the smart large object read function.
<i>major_version</i>	INTEGER	The major version number of the Change Data Capture API. Must be 1.
<i>minor_version</i>	INTEGER	The minor version number of the Change Data Capture API. Must be 1 for new applications. Can be 0 for existing applications.

Usage

Use the `cdc_opensess()` function to open a communication session between the client application and the database server. The session ID returned by the `cdc_opensess()` is the smart large object file descriptor that you supply to the smart large object read function. To start capturing data, you must then use the `cdc_activatesess()` function and the `cdc_startcapture()` function.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Important: If you have multiple applications that use the Change Data Capture API and connect to the same Informix server, all applications must use the same values for the *major_version* and *minor_version* arguments.

Return values

If successful, returns an integer that is the session ID.

If unsuccessful, returns an integer corresponding to an error code.

The `cdc_recboundary()` function

Restarts data capture from the beginning of the CDC record currently being returned.

Syntax

► `cdc_recboundary`—(*—session_ID—*)—►

Function argument

Table 2-7. The `cdc_recboundary()` argument

Argument	Data Type	Description
<i>session_ID</i>	INTEGER	The session ID of the open capture session.

Usage

Use the `cdc_recboundary()` function if you need to restart capture from the beginning of the current log record.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns a positive integer representing the number of complete or partial log records that were captured but skipped during the current session.

If unsuccessful, returns an integer corresponding to an error code and updates the `syscdcsess` table with the error information.

The `cdc_set_fullrowlogging()` function

Enables or disables full-row logging for a table.

Purpose

You must run this function to enable full-row logging on a table before you can start capturing data from it.

The `DB_LOCALE` environment variable must be set to the same locale as the database locale when you run this function.

►►—cdc_set_fullrowlogging—(—"database:owner.table_name"—,—logging—)————►►

Function arguments

Table 2-8. The `cdc_set_fullrowlogging()` arguments

Argument	Data Type	Description
<i>database:owner.table_name</i>	LVARCHAR	The qualified name of the table. The qualified name includes the following elements: <i>database</i> The name of the database in which the table exists. <i>owner</i> The name of the owner of the table. <i>table</i> The name of the table.
<i>logging</i>	INTEGER	<ul style="list-style-type: none"> • 0 Disable full-row logging • 1 Enable full-row logging

Usage

Use the `cdc_set_fullrowlogging()` function to enable full-row logging on a table from which you intend to perform data capture. This function must be run as user **informix**. After you stop capturing data from a table, you can disable full-row logging.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns 0.

If unsuccessful, returns an integer corresponding to an error code and updates the `syscdcsess` table with the error information.

The `cdc_startcapture()` function

Specifies the data to start capturing from a table.

If the capture session is both open and active (you have run the `cdc_activatesess()` function), data capture starts immediately on the specified columns in the specified table. Otherwise, data capture starts when you activate the open capture session.

The `DB_LOCALE` environment variable must be set to the same locale as the database locale when you run this function.

A row is added in the `syscdctabs` table associated with the specified session ID and table.

Syntax

```

▶ cdc_startcapture(—session_ID—, —MBZ—,
▶ "database:owner.table_name", "column_name", —user_data—)

```

Function arguments

Table 2-9. `cdc_startcapture()` arguments

Argument	Data Type	Description
<i>session_ID</i>	INTEGER	The session ID of an open capture session.
<i>MBZ</i>	BIGNIT	Must be 0. Reserved.
<i>database:owner.table_name</i>	LVARCHAR	The qualified name of the table from which to capture data. The qualified name includes the following elements: <i>database</i> The name of the database in which the table exists. <i>owner</i> The name of the owner of the table. <i>table</i> The name of the table.
<i>column_name</i>	LVARCHAR	A comma-separated list of column names in the specified table, from which to capture data.
<i>user_data</i>	INTEGER	The table identifier.

Usage

Use the `cdc_startcapture()` function to specify a table and columns within that table from which to start capturing data. You cannot include columns with simple large objects, user-defined data types, or collection data types.

The table identifier is a number you use in your application to uniquely identify each table that will participate in data capture.

You must call this function from a client application. You cannot call this function from a user-defined routine that runs within the database server.

Return values

If successful, returns 0.

If unsuccessful, returns an integer corresponding to an error code and updates the `syscdcsess` table with the error information.

Chapter 3. Change Data Capture records

These topics describe the CDC records returned from calls to read functions from an open capture session.

Format of CDC records

The Change Data Capture (CDC) records contain a header that is common to all records, followed by a specific header for the type of CDC record.

The CDC_REC_INSERT, CDC_REC_DELETE, CDC_REC_UPDBEF, and CDC_REC_UPDAFT records also contain column data.

The header common to all CDC records describes the size and type of the CDC record.

Table 3-1. The header common to all CDC records

Section	Size	Description
Header size	4 bytes	The number of bytes in the common and CDC record-specific headers.
Payload size	4 bytes	The number of bytes of data in the record after the common and CDC record-specific headers.
Packet scheme	4 bytes	The packetization scheme number of one of the packetization schemes contained in the syscdcpacketschemes table. The only packetization scheme is 66, CDC_PKTScheme_LRECBINARY.
Record number	4 bytes	The record number of one of the CDC records contained in the syscdcrectypes table.

Related reference:

- “The CDC_REC_BEGINTX record”
- “The CDC_REC_COMMTX record”
- “The CDC_REC_DELETE record” on page 3-3
- “The CDC_REC_DISCARD record” on page 3-4
- “The CDC_REC_ERROR record” on page 3-4
- “The CDC_REC_INSERT record” on page 3-5
- “The CDC_REC_RBTX record” on page 3-6
- “The CDC_REC_TABSCHEMA record” on page 3-6
- “The CDC_REC_TIMEOUT record” on page 3-7
- “The CDC_REC_TRUNCATE record” on page 3-8
- “The CDC_REC_UPDAFT record” on page 3-8
- “The CDC_REC_UPDBEF record” on page 3-9

The CDC_REC_BEGINTX record

Indicates the beginning of a transaction.

The header for the CDC_REC_BEGINTX record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-2. Format of the CDC_REC_BEGINTX record

Section	Size	Description
Sequence number	8 bytes	The sequence number of the record.
Transaction ID	4 bytes	The transaction ID.
Start time	8 bytes	The UTC time at which the transaction began, in time_t format.
User ID	4 bytes	The operating system user ID of the user who started the transaction.

Related reference:

- “Format of CDC records” on page 3-1
- “The syscdcrectypes table” on page 4-1

The CDC_REC_COMMTX record

Indicates that a transaction has been committed.

The header for the CDC_REC_COMMTX record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-3. Format of the CDC_REC_COMMTX record

Section	Size	Description
Sequence number	8 bytes	The sequence number of the record.
Transaction ID	4 bytes	The transaction ID.

Table 3-3. Format of the CDC_REC_COMMTX record (continued)

Section	Size	Description
Commit time	8 bytes	The UTC time at which the transaction was committed, in time_t format.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_DELETE record

Provides the row that was removed as a result of a DELETE operation.

The CDC_REC_DELETE record consists of these fields:

- The common header.
- The record-specific header.
- Fields listing the size of each variable-length column in the row, if any.
- Column data for each fixed-length column, if any.
- Column data for each variable-length column, if any.

The value in the header size field in the common header represents the number of bytes occupied by the combination of the common header, the record-specific header, and the fields listing the size of variable-length columns.

The value in the payload size field in the common header represents the number of bytes of the column data for both fixed-length and variable length columns.

The record-specific header

The header specific to the CDC_REC_DELETE record follows the common header.

Table 3-4. The CDC_REC_DELETE record header

Section	Size	Description
Sequence number	8 bytes	The sequence number associated with the DELETE operation.
Transaction ID	4 bytes	The transaction ID.
User data	4 bytes	The table identifier passed to the <code>cdc_startcapture()</code> function and stored in the <code>syscdtabs</code> table.
Flags	4 bytes	Reserved.

Variable-length column size fields

If there are variable-length columns in the row being deleted, a 4-byte field for each of those columns appears containing the column size. The order of the column size fields is the same as the order of the columns in the CDC_REC_TABSCHEMA record.

Fixed-length column data

The data from the fixed-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Variable-length column data

The data from the variable-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_DISCARD record

Indicates that some operations of the transaction should be discarded.

CDC records for the same transaction that follow this record should be discarded.

The header specific to the CDC_REC_DISCARD record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-5. Format of the CDC_REC_DISCARD record

Section	Size	Description
Sequence number	8 bytes	The sequence number of the record. Any CDC records that have the same transaction ID value and that have a sequence number greater than or equal to this sequence number should be discarded.
Transaction ID	4 bytes	The transaction ID.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_ERROR record

Indicates that an error occurred and the session is no longer valid.

The header specific to the CDC_REC_ERROR record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-6. Format of the CDC_REC_ERROR record

Section	Size	Description
Flags	4 bytes	Hexadecimal flag: <ul style="list-style-type: none"> • 0x1 indicates that the capture session is no longer valid and the only valid operation is to run the <code>cdc_closesess()</code> function to close the session. • any other value indicates that the session is still valid.
Error code	4 bytes	The error code.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_INSERT record

Provides the row that resulted from an INSERT operation.

The CDC_REC_INSERT record consists of these fields:

- The common header.
- The record-specific header.
- Fields listing the size of each variable-length column in the row, if any.
- Column data for each fixed-length column, if any.
- Column data for each variable-length column, if any.

The value in the header size field in the common header represents the number of bytes occupied by the combination of the common header, the record-specific header, and the fields listing the size of variable-length columns.

The value in the payload size field in the common header represents the number of bytes of the column data for both fixed-length and variable length columns.

The record-specific header

The header specific to the CDC_REC_INSERT record follows the common header.

Table 3-7. The CDC_REC_INSERT record header

Section	Size	Description
Sequence number	8 bytes	The sequence number associated with the INSERT operation.
Transaction ID	4 bytes	The transaction ID.

Table 3-7. The CDC_REC_INSERT record header (continued)

Section	Size	Description
User data	4 bytes	The table identifier passed to the <code>cdc_startcapture()</code> function and stored in the <code>syscdtabs</code> table.
Flags	4 bytes	Reserved.

Variable-length column size fields

If there are variable-length columns in the row being inserted, a 4-byte field for each of those columns appears containing the column size. The order of the column size fields is the same as the order of the columns in the CDC_REC_TABSCHEMA record.

Fixed-length column data

The data from the fixed-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Variable-length column data

The data from the variable-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_RBTX record

Indicates that the transaction has been rolled back.

The header specific to the CDC_REC_RBTX record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-8. Format of the CDC_REC_RBTX record

Section	Size	Description
Sequence number	8 bytes	The sequence number associated with the ROLLBACK operation.
Transaction ID	4 bytes	The transaction ID.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_TABSCHEMA record

Describes the table from which data is being captured.

The value in the payload size field in the common header represents the number of bytes occupied by the column name and data type list.

The header specific to the CDC_REC_TABSCHEMA record follows the common header.

Table 3-9. Format of the CDC_REC_TABSCHEMA record

Section	Size	Description
User data	4 bytes	The table identifier that was specified in the <code>cdc_startcapture()</code> function for the table being captured.
Flags	4 bytes	Must be 0.
Fixed-length size	4 bytes	The number of bytes of data in fixed-length columns in the table.
Fixed-length columns	4 bytes	The number of fixed-length columns in the table being captured. A 0 indicates that there are no fixed-length columns.
Variable-length columns	4 bytes	The number of variable-length columns in the table being captured. A 0 indicates that there are no variable-length columns.
Column names and data types	variable byte length	A comma-separated list of column names and data types in UTF-8 format. The column list conforms to the syntax of the column list in a CREATE TABLE statement. Names of any fixed-length columns appear before names of any variable-length columns. The number of columns equals the number of fixed-length columns plus the number of variable-length columns.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_TIMEOUT record

Indicates that the read call did not return data before the timeout period specified in the `cdc_opensess()` function.

The header specific to the CDC_REC_TIMEOUT record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-10. Format of the CDC_REC_TIMEOUT record

Section	Size	Description
Sequence number	8 bytes	The sequence number of the last data retrieved from the source database.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_TRUNCATE record

Indicates that a TRUNCATE operation was performed on a table.

The header specific to the CDC_REC_TRUNCATE record follows the common header. No data follows the headers; the payload size in the common header is 0.

Table 3-11. Format of the CDC_REC_TRUNCATE record

Section	Size	Description
Sequence number	8 bytes	The sequence number associated with the TRUNCATE operation.
Transaction ID	4 bytes	The transaction ID.
User data	4 bytes	The table identifier passed to the <code>cdc_startcapture()</code> function and stored in the <code>syscdtabs</code> table.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_UPDAFT record

Provides the image of a row after an UPDATE operation.

The CDC_REC_UPDAFT record consists of these fields:

- The common header.
- The record-specific header.
- Fields listing the size of each variable-length column in the row, if any.
- Column data for each fixed-length column, if any.
- Column data for each variable-length column, if any.

The value in the header size field in the common header represents the number of bytes occupied by the combination of the common header, the record-specific header, and the fields listing the size of variable-length columns.

The value in the payload size field in the common header represents the number of bytes of the column data for both fixed-length and variable length columns.

The record-specific header

The header specific to the CDC_REC_UPDAFT record follows the common header.

Table 3-12. The CDC_REC_UPDAFT record header

Section	Size	Description
Sequence number	8 bytes	The sequence number associated with the UPDATE operation.
Transaction ID	4 bytes	The transaction ID.
User data	4 bytes	The table identifier passed to the <code>cdc_startcapture()</code> function and stored in the <code>syscdtabs</code> table.
Flags	4 bytes	Reserved.

Variable-length column size fields

If there are variable-length columns in the row being updated, a 4-byte field for each of those columns appears containing the column size. The order of the column size fields is the same as the order of the columns in the CDC_REC_TABSCHEMA record.

Fixed-length column data

The data from the fixed-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Variable-length column data

The data from the variable-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

The CDC_REC_UPDBEF record

Provides the image of a row before an UPDATE operation.

The CDC_REC_UPDBEF record consists of these fields:

- The common header.
- The record-specific header.
- Fields listing the size of each variable-length column in the row, if any.
- Column data for each fixed-length column, if any.
- Column data for each variable-length column, if any.

The value in the header size field in the common header represents the number of bytes occupied by the combination of the common header, the record-specific header, and the fields listing the size of variable-length columns.

The value in the payload size field in the common header represents the number of bytes of the column data for both fixed-length and variable length columns.

The record-specific header

The header specific to the CDC_REC_UPDBEF record follows the common header.

Table 3-13. The CDC_REC_UPDBEF record header

Section	Size	Description
Sequence number	8 bytes	The sequence number associated with the UPDATE operation.
Transaction ID	4 bytes	The transaction ID.
User data	4 bytes	The table identifier passed to the <code>cdc_startcapture()</code> function and stored in the <code>syscdtabs</code> table.
Flags	4 bytes	Reserved.

Variable-length column size fields

If there are variable-length columns in the row being updated, a 4-byte field for each of those columns appears containing the column size. The order of the column size fields is the same as the order of the columns in the CDC_REC_TABSCHEMA record.

Fixed-length column data

The data from the fixed-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Variable-length column data

The data from the variable-length columns, if any, appears in the order that the corresponding columns are listed in the CDC_REC_TABSCHEMA record.

Related reference:

“Format of CDC records” on page 3-1

“The syscdcrectypes table” on page 4-1

Chapter 4. The syscdc system database

The **syscdc** system database contains tables that store information about the Change Data Capture API.

The **syscdc** database can only be accessed or connected to by the user **informix**. It uses the UTF-8 locale. You cannot alter the tables in the **syscdc** database; you can only query them.

The syscdcerrcodes table

Contains the error codes used by the Change Data Capture API.

Use this table to look up the symbolic name and description that correspond to an error code.

*Table 4-1. The **syscdcerrcodes** table*

Column	Data Type	Description
errcode	INTEGER	Numeric value of the error.
errname	VARCHAR(16)	Symbolic name of the error.
errdesc	VARCHAR(127)	Error description.

Related reference:

Chapter 5, “Change Data Capture error codes,” on page 5-1

The syscdcrectypes table

Contains the record types used by the Change Data Capture API.

Use this table to look up the symbolic name and description that correspond to a record code.

*Table 4-2. The **syscdcrectypes** table*

Column	Data Type	Description
recnum	INTEGER	Numeric value of the record type.
recname	VARCHAR(16)	Symbolic name of the record type.
recdesc	VARCHAR(127)	Record type description.

Related reference:

- "The CDC_REC_BEGINTX record" on page 3-2
- "The CDC_REC_COMMTX record" on page 3-2
- "The CDC_REC_DELETE record" on page 3-3
- "The CDC_REC_DISCARD record" on page 3-4
- "The CDC_REC_ERROR record" on page 3-4
- "The CDC_REC_INSERT record" on page 3-5
- "The CDC_REC_RBTX record" on page 3-6
- "The CDC_REC_TABSCHEMA record" on page 3-6
- "The CDC_REC_TIMEOUT record" on page 3-7
- "The CDC_REC_TRUNCATE record" on page 3-8
- "The CDC_REC_UPDAFT record" on page 3-8
- "The CDC_REC_UPDBEF record" on page 3-9

Chapter 5. Change Data Capture error codes

If a Change Data Capture function encounters a problem, it returns an error code. Most functions return 0 if they succeed.

Error numbers are not guaranteed to remain the same in subsequent releases. Always use the symbolic names in your application code. You can view the error message text corresponding to a symbolic error name by using the `cdc_errortext()` function.

Table 5-1. Change Data Capture error codes

Symbolic Name	Description
CDC_E_OK	Operation succeeded.
CDC_E_NOCDCDB	The <code>syscdc</code> database does not exist.
CDC_E_APIVERS	The requested CDC API behavior version is not valid or is unsupported.
CDC_E_NODB	The specified database does not exist.
CDC_E_DBNOTLOGGED	The specified database is not logged.
CDC_E_NOTAB	The specified table does not exist.
CDC_E_TABPROPERTIES	The table properties do not support capture: it is a temporary table, a view, or otherwise not logged.
CDC_E_NOCOL	The specified column does not exist.
CDC_E_NOSES	The specified CDC session does not exist.
CDC_E_NOREOPEN	The CDC session cannot be reopened.
CDC_E_TABCAPTURED	The specified table is already being captured by the CDC session.
CDC_E_TABNOTCAPTURED	The specified table is not being captured by the CDC session.
CDC_E_ARGNULL	An argument to the function has the SQL NULL value, which is not allowed.
CDC_E_LSN	Data at the requested log sequence number is unavailable for capture.
CDC_E_DUPLSESS	A CDC session is already active.
CDC_E_ARG	A parameter passed to the function is not valid.
CDC_E_ARG1	The first parameter passed to the function is not valid.
CDC_E_ARG2	The second parameter passed to the function is not valid.
CDC_E_ARG3	The third parameter passed to the function is not valid.
CDC_E_ARG4	The fourth parameter passed to the function is not valid.
CDC_E_ARG5	The fifth parameter passed to the function is not valid.

Table 5-1. Change Data Capture error codes (continued)

Symbolic Name	Description
CDC_E_ARG6	The sixth parameter passed to the function is not valid.
CDC_E_INTERNAL	Internal error. Contact IBM Support.
CDC_E_NOMEM	Memory allocation failed.
CDC_E_MUSTCLOSE	The CDC capture session cannot continue and must be closed.
CDC_E_BADSTATE	The resource state does not allow the attempted operation.
CDC_E_BADCHAR	A byte sequence that is not a valid character in the character code set was encountered.
CDC_E_INTERRUPT	The CDC session was interrupted.
CDC_E_UNIMPL	Unimplemented feature.
CDC_E_LOCALEMISMATCH	The locale setting in the environment does not match the locale of the database.

Related tasks:

“Handling errors” on page 1-6

Related reference:

“The sysdcerrcodes table” on page 4-1

“The cdc_errortext() function” on page 2-4

Table 6-1. The **onstat -g cdc** syntax elements (continued)

Element	Purpose
table	<p>Displays information about the tables being captured, including:</p> <ul style="list-style-type: none"> • The number of tables being captured in a session. • The full name of each table being captured. • The time when data capture on each table started. • With the long option, information about the captured columns for each table. <p>If you specify a fully-qualified table name, only the information for that table is displayed. If you do not specify a table name, information for all tables is displayed.</p>

Examples

The following examples display sample output of the **onstat -g cdc** command with some of its options.

Example 1: Detailed session information

The following command generates output that shows detailed information about the session 159383591:

```
onstat -g cdc 159383591 long
```

```
CDC subsystem structure at 0x44252318
  CDC session structure at 0x4d8e0d00
    CDC session id: 159383591 (0x9800027)
    Associated SQL session id: 304
    Number of tables captured: 1
    State: ACTIVATED (0x50534555)
    Create time: 1238530254 (Tue Mar 31 15:10:54 2009)
    Open time: 1238530254 (Tue Mar 31 15:10:54 2009)
    Activate time: 1238530256 (Tue Mar 31 15:10:56 2009)
    Activate Sequence Number: 0x0
    Total client read calls: 9
    Last client read time: 1238530321 (Tue Mar 31 15:12:01 2009)
    Last Sequence Number returned to client: 0x150004b774
    Total number records examined: 4385
    Total number records kept (approximate): 1937
    Total number I/U/D records examined: 1046
    Total number I/U/D records kept (approximate): 582
    Client required to close: NO
    Read exit error code: 0
```

Example 2: Configuration information

The following command generates output that shows information about the configuration of open sessions:

```
onstat -g cdc config
```

```
CDC subsystem structure at 0x44252318
  CDC session structure at 0x4dba3d00
    CDC session id: 160432167 (0x9900027)
    Read Timeout (seconds): 3
    Maximum buffers per read call: 4
    Survive DATALOST errors: NO

  CDC session structure at 0x4d8e0d00
    CDC session id: 159383591 (0x9800027)
    Read Timeout (seconds): 3
```

```
Maximum buffers per read call: 4
Survive DATALOST errors: NO
```

```
CDC session structure at 0x4c022d00
CDC session id: 158335015 (0x9700027)
Read Timeout (seconds): 3
Maximum buffers per read call: 4
Survive DATALOST errors: NO
```

Example 3: Buffer information

The following command generates output that shows information about the buffers being used by currently open sessions:

```
onstat -g cdc 0 bufm
```

```
CDC subsystem structure at 0x44252318
CDC session structure at 0x4dba3d00
  CDC session id: 160432167 (0x9900027)

  Buffer Manager at 0x4dba5028
    Number of allocated buffers high watermark: 268
    Number of currently allocated buffers: 267
    Minimum prepend for allocated buffers: 172

CDC session structure at 0x4d8e0d00
CDC session id: 159383591 (0x9800027)

  Buffer Manager at 0x4d8e2028
    Number of allocated buffers high watermark: 271
    Number of currently allocated buffers: 270
    Minimum prepend for allocated buffers: 172

CDC session structure at 0x4c022d00
CDC session id: 158335015 (0x9700027)

  Buffer Manager at 0x4c6e5028
    Number of allocated buffers high watermark: 269
    Number of currently allocated buffers: 267
    Minimum prepend for allocated buffers: 172
```

Example 4: Table information

The following command generates output that shows information about the session 158335015 for the table named **account**:

```
onstat -g cdc 158335015 table bank:pinch.account
```

```
CDC subsystem structure at 0x44252318
CDC session structure at 0x4c022d00
  CDC session id: 158335015 (0x9700027)
  Captured Table Manager found at 0x4c048b20
  Number of tables captured: 1

    Captured Table structure at 0x4c6e5160
      Full Table Name: bank:pinch.account
      Version Sequence Number: 0xe00238388
      Time capture started: 1238530249 (Tue Mar 31 15:10:49 2009)
```

Related tasks:

“Monitoring data capture” on page 1-8

Chapter 7. Change Data Capture sample program

The Change Data Capture sample program provides an example of using the Change Data Capture API to capture and process data.

The IBM Informix sample program, `cdcapi.ec`, is located in the `INFORMIXDIR/demo/cdc` directory. The program creates an application that captures data from multiple tables. The program runs Change Data Capture functions, reads CDC records, and displays the column values of the captured data to `stdout`. The program also queries the `syscdc` system tables to display information about CDC records and error messages. The program terminates when it encounters an error or a `CDC_REC_TIMEOUT` record.

The program has a command-line interface that you use to enter the database name, the table name, column names, and the timeout value.

This program requires that the `getopt` parser function is implemented on your computer.

Related tasks:

“Writing an application to capture data changes” on page 1-5

Appendix. Accessibility

IBM strives to provide products with usable access for everyone, regardless of age or ability.

Accessibility features for IBM Informix products

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Informix products. These features support:

- Keyboard-only operation.
- Interfaces that are commonly used by screen readers.
- The attachment of alternative input and output devices.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

Related accessibility information

IBM is committed to making our documentation accessible to persons with disabilities. Our publications are available in HTML format so that they can be accessed with assistive technology such as screen reader software.

IBM and accessibility

See the *IBM Accessibility Center* at <http://www.ibm.com/able> for more information about the IBM commitment to accessibility.

Dotted decimal syntax diagrams

The syntax diagrams in our publications are available in dotted decimal format, which is an accessible format that is available only if you are using a screen reader.

In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), the elements can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read punctuation. All syntax elements that have the same dotted decimal number (for example, all syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, the word or symbol is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is read as 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol that provides information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, that element is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 refers to a separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? Specifies an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element (for example, 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! Specifies a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.
- * Specifies a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be

repeated. For example, if you hear the line 5.1* data-area, you know that you can include more than one data area or you can include none. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
 2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you can write HOST STATE, but you cannot write HOST HOST.
 3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.
- + Specifies a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times. For example, if you hear the line 6.1+ data-area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. As for the * symbol, you can repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy,

modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

A

- Accessibility A-1
 - dotted decimal format of syntax diagrams A-1
 - keyboard A-1
 - shortcut keys A-1
 - syntax diagrams, reading in a screen reader A-1
- application development with Change Data Capture API 1-5

C

- capture session
 - activating 2-1
 - closing 2-2
 - deactivating 2-2
 - opening 2-5
 - specifying data 2-8
- capture session for Change Data Capture 1-5
- cdc_activatesess() function 2-1
- cdc_closesess() function 2-2
- cdc_deactivatesess() function 2-2
- cdc_endcapture() function 2-3
- cdc_errortext() function 2-4
- cdc_opensess() function 2-5
- CDC_REC_BEGINTX record 3-2
- CDC_REC_COMMTX record 3-2
- CDC_REC_DELETE record 3-3
- CDC_REC_DISCARD record 3-4
- CDC_REC_ERROR record 3-5
- CDC_REC_INSERT record 3-5
- CDC_REC_RBTX record 3-6
- CDC_REC_TABSCHEMA record 3-6
- CDC_REC_TIMEOUT record 3-7
- CDC_REC_TRUNCATE record 3-8
- CDC_REC_UPDAFT record 3-8
- CDC_REC_UPDBEF record 3-9
- cdc_reboundary() function 2-7
- cdc_set_fullrowlogging() function 2-7
- cdc_startcapture() function 2-8
- Change Data Capture
 - application development 1-5
 - data types supported 1-5
 - error handling 1-6
 - logging 2-7
 - monitoring 1-8
 - onstat -g cdc 1-8
 - overview 1-1
 - restarting capture 1-7
 - sample program 7-1
 - smart large objects 1-2, 1-7
 - syscdc database 4-1
 - transactions captured 1-1
- Change Data Capture API
 - monitoring 6-1
 - prerequisites 1-5
- Change Data Capture API components 1-2
- Change Data Capture error codes 5-1
- Change Data Capture functions
 - cdc_activatesess() 2-1
 - cdc_closesess() 2-2
 - cdc_deactivatesess() 2-2

Change Data Capture functions (continued)

- cdc_endcapture() 2-3
- cdc_errortext() 2-4
- cdc_opensess() 2-5
- cdc_reboundary() 2-7
- cdc_set_fullrowlogging() 2-7
- cdc_startcapture() 2-8
- Change Data Capture records
 - CDC_REC_BEGINTX 3-2
 - CDC_REC_COMMTX 3-2
 - CDC_REC_DELETE 3-3
 - CDC_REC_DISCARD 3-4
 - CDC_REC_ERROR 3-5
 - CDC_REC_INSERT 3-5
 - CDC_REC_RBTX 3-6
 - CDC_REC_TABSCHEMA 3-6
 - CDC_REC_TIMEOUT 3-7
 - CDC_REC_TRUNCATE 3-8
 - CDC_REC_UPDAFT 3-8
 - CDC_REC_UPDBEF 3-9
 - format of 3-1
 - sequence number 1-4
- Change Data Capture system tables
 - syscdccerrorcodes 4-1
 - syscdcrectypes 4-1
- compliance with standards vi

D

- data buffer for Change Data Capture 1-5
- data capture
 - data selection 2-8
 - logging 2-7
 - restarting 1-7
 - starting 1-7, 2-1, 2-5
 - stopping 2-2
- data types for Change Data Capture 1-5
- Disabilities, visual
 - reading syntax diagrams A-1
- Disability A-1
- Dotted decimal format of syntax diagrams A-1

E

- error codes for Change Data Capture 5-1
- error handling for Change Data Capture 1-6
- error text
 - returning 2-4

I

- industry standards vi

L

- logging for Change Data Capture 1-5, 2-7

O

onstat -g cdc 6-1

P

prerequisites for Change Data Capture API 1-5

R

rewinding data capture 2-7

S

sample program for Change Data Capture 7-1

Screen reader

 reading syntax diagrams A-1

sequence number for CDC records 1-4

session ID for Change Data Capture 1-5, 2-5

Shortcut keys

 keyboard A-1

smart large object read functions 1-5

smart large objects 1-2, 1-7

standards vi

status of data capture session 1-8

status of table capture 1-8

Syntax diagrams

 reading in a screen reader A-1

syscdc database 4-1

syscdccerrcodes table 4-1

syscdcrectypes table 4-1

T

table schema for Change Data Capture 1-5

target destination for Change Data Capture 1-5

transactions

 Change Data Capture 1-1

V

Visual disabilities

 reading syntax diagrams A-1



Printed in USA

SC27-4509-00

