

IBM DB2 Information Integrator



# Introduction to Replication and Event Publishing

*Version 8.2*



IBM DB2 Information Integrator



# Introduction to Replication and Event Publishing

*Version 8.2*

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 51.

This document contains proprietary information of IBM. It is provided under a license agreement and Copyright law protects it. 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.

You can order IBM publications online or through your local IBM representative:

- To order publications online, go to the IBM Publications Center at [www.ibm.com/shop/publications/order](http://www.ibm.com/shop/publications/order)
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at [www.ibm.com/planetwide](http://www.ibm.com/planetwide)

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 International Business Machines Corporation 2004. All rights reserved.

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

---

# Contents

<b>About this book</b> . . . . .	<b>v</b>
Who should read this book . . . . .	v
Text conventions used in this book . . . . .	v
Terminology used in this book . . . . .	v
Where to find more information . . . . .	vi
<b>Chapter 1. Introduction to DB2 replication and publishing</b> . . . . .	<b>1</b>
Introduction to DB2 replication and publishing—Overview . . . . .	1
Introduction to SQL replication . . . . .	1
Introduction to Q replication . . . . .	2
Introduction to event publishing. . . . .	3
Administrative interface for DB2 replication and publishing . . . . .	4
Utilities for DB2 replication and publishing . . . . .	4
<b>Chapter 2. SQL replication</b> . . . . .	<b>7</b>
SQL replication—Overview . . . . .	7
Infrastructure for an SQL replication environment . . . . .	8
Registration of sources in SQL replication. . . . .	9
Subscription sets in SQL replication . . . . .	11
Capture of data from DB2 UDB sources in SQL replication. . . . .	15
Application of data to DB2 UDB targets in SQL replication. . . . .	17
Capture of data from non-DB2 sources in SQL replication. . . . .	18
Application of data to non-DB2 targets in SQL replication. . . . .	19
<b>Chapter 3. Q replication</b> . . . . .	<b>21</b>
Q replication—Overview . . . . .	21
Infrastructure for a Q replication environment. . . . .	22
Sources and targets in Q replication . . . . .	23
Capture of data in Q replication . . . . .	25
Application of data to targets in Q replication. . . . .	26
Types of replication in Q replication . . . . .	27
<b>Chapter 4. Comparison of SQL replication and Q replication</b> . . . . .	<b>31</b>
Comparison of SQL replication and Q replication—Overview. . . . .	31
Comparison of the infrastructure of SQL replication and Q replication . . . . .	31
Comparison of sources and targets in SQL replication and Q replication . . . . .	32
Comparison of data capturing and applying in SQL replication and Q replication . . . . .	33
Which replication solutions should you use in common scenarios?. . . . .	33
<b>Chapter 5. Event publishing</b> . . . . .	<b>37</b>
Event publishing—Overview . . . . .	37
Infrastructure for an event publishing environment . . . . .	38
Sources in event publishing . . . . .	39
Capture of data in event publishing . . . . .	41
<b>Chapter 6. Comparison of SQL replication, Q replication, and event publishing</b> . . . . .	<b>43</b>
<b>Chapter 7. Comparison of Q replication to high availability disaster recovery (HADR)</b> . . . . .	<b>45</b>
<b>Chapter 8. Comparison of event publishing to DB2 UDB MQ user-defined functions</b> . . . . .	<b>47</b>
<b>Accessibility</b> . . . . .	<b>49</b>
Keyboard input and navigation. . . . .	49
Accessible display . . . . .	49
Font settings . . . . .	49
Nondependence on color . . . . .	49
Alternative alert cues . . . . .	49
Compatibility with assistive technologies . . . . .	49
Accessible documentation . . . . .	50
<b>Notices</b> . . . . .	<b>51</b>
Trademarks . . . . .	53
<b>Index</b> . . . . .	<b>55</b>
<b>Contacting IBM</b> . . . . .	<b>57</b>
Product information . . . . .	57
Comments on the documentation . . . . .	57



---

## About this book

This book introduces SQL replication, Q replication, and event publishing. You can use SQL replication, Q replication, and event publishing with other products in the IBM® replication solution to tailor a replication and publishing environment that suits your business needs.

---

## Who should read this book

This book is for database administrators, data replication specialists, and others who set up and maintain a SQL replication, Q replication, or event publishing environment. This book assumes that you are familiar with:

- Standard database terminology.
- Database design, database administration, database security, server connectivity, and networking.
- The operating systems that will be involved in your SQL replication, Q replication, or event publishing environment.
- The data that you want to replicate or publish.
- The applications that receive messages (event publishing only).

---

## Text conventions used in this book

This book uses these highlighting conventions:

- **Boldface type** indicates commands or user interface controls such as names of fields, folders, icons, or menu choices.
- Monospace type indicates examples of text that you enter exactly as shown.
- *Italic type* indicates variables that you should replace with a value. It is also used to indicate book titles and to emphasize words.

---

## Terminology used in this book

This book uses standard terminology for database, connectivity, copying, SQL, and LAN concepts.

Unless otherwise specified, the following meanings are assumed:

**Linux** Linux refers to SQL replication, Q replication, and event publishing for all Linux platforms (such as Linux Intel and Linux z/Series).

**UNIX** UNIX refers to SQL replication, Q replication, and event publishing for all UNIX platforms (such as HP UX, Solaris Operating Environment, and AIX).

### **Windows**

Windows refers to SQL replication, Q replication, and event publishing for Windows platforms (such as 2000 and XP).

**z/OS** z/OS refers to SQL replication, Q replication, and event publishing for z/OS and OS/390. z/OS is the next generation of the OS/390 operating system, and it also includes UNIX System Services (USS) on z/OS.

---

## Where to find more information

This section identifies other sources of information about SQL replication, Q replication, and event publishing that you might find useful.

---

If you want to ...	Refer to ...
Learn about the DB2 Information Integrator replication solution	<a href="http://www.ibm.com/software/data/dpropr/">www.ibm.com/software/data/dpropr/</a>
Learn how to set up and maintain SQL replication, Q replication, or event publishing	Either of the following places: <ul style="list-style-type: none"><li>• The DB2 Information Center at <a href="http://publib.boulder.ibm.com/infocenter/db2help/">http://publib.boulder.ibm.com/infocenter/db2help/</a></li><li>• <i>IBM DB2 Information Integrator SQL Replication Guide and Reference</i></li><li>• <i>IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference</i></li></ul>
View the information set for DB2 Universal Database and DB2 Information Integrator	See the DB2 Information Center <a href="http://publib.boulder.ibm.com/infocenter/db2help/">http://publib.boulder.ibm.com/infocenter/db2help/</a>
Find technical support resources and customer support options	<a href="http://www.ibm.com/software/data/integration/db2ii/support.html">http://www.ibm.com/software/data/integration/db2ii/support.html</a>
Learn about WebSphere MQ	See the <i>WebSphere MQ System Administration Guide</i> . Download the PDF or view the HTML version at <a href="http://www-306.ibm.com/software/integration/mqfamily/library/manualsa/manuals/crosslatest.html">http://www-306.ibm.com/software/integration/mqfamily/library/manualsa/manuals/crosslatest.html</a>
Learn about DB2 Information Integrator federated systems	See the <i>DB2 Information Integrator Federated Systems Guide</i> . The information in this manual is also included in the DB2 Information Center, which is at <a href="http://publib.boulder.ibm.com/infocenter/db2help/">http://publib.boulder.ibm.com/infocenter/db2help/</a> .

---



---

# Chapter 1. Introduction to DB2 replication and publishing

---

## Introduction to DB2 replication and publishing—Overview

DB2® Universal Database version 8.2 provides two different solutions that you can use to replicate data from and to relational databases: *SQL replication* and *Q replication*. In SQL replication, committed source changes are staged in relational tables before being replicated to target systems. In Q replication, committed source changes are written them in messages that are transported through WebSphere® MQ message queues to target systems.

DB2 UDB version 8.2 also provides a solution called *event publishing* for converting committed source changes into messages in an XML format and publishing those messages to applications such as message brokers.

The following topics briefly introduce the replication and publishing solutions. These topics also describe the administrative interface for these solutions and utilities to help you maintain replication and publishing environments.

- Introduction to SQL replication
- Introduction to Q replication
- Introduction to event publishing
- Administrative interface for DB2 replication and publishing
- Utilities for DB2 replication and publishing

**Related concepts:**

- “Administrative interface for DB2 replication and publishing” on page 4
- “Introduction to Q replication” on page 2
- “Introduction to SQL replication” on page 1
- “Utilities for DB2 replication and publishing” on page 4
- “Introduction to event publishing” on page 3

---

## Introduction to SQL replication

SQL replication captures changes to sources and uses staging tables to store committed transactional data. The changes are then read from the staging tables and replicated to corresponding target tables. With staging tables, data can be captured and staged once for delivery to multiple targets, in different formats, and at different delivery intervals.

You can use SQL replication for a variety of purposes that require replicated data, including capacity relief, feeding data warehouses and data marts, and auditing change history.

You can replicate continuously, at intervals, or for one time only. Replicating continuously can be useful if your applications need data in near-real-time, such as applications for making airline reservations. Replicating at intervals can be useful for replicating large batches of data during off-peak hours. You can also trigger replication through database events.

Sources and targets can be either in DB2<sup>®</sup> UDB databases or in non-DB2 relational databases. You can replicate from sources and to targets in the following relational database management systems: DB2 UDB on Linux, UNIX<sup>®</sup>, Windows<sup>®</sup>, z/OS<sup>™</sup>, and iSeries<sup>™</sup>; Microsoft<sup>®</sup> SQL Server; Oracle; Sybase; and Teradata (target only). If you plan to replicate to or from a non-DB2 relational database, you will need the federated server function of DB2 Information Integrator. See the *DB2 Information Integrator Federated Systems Guide* for more information about federated servers.

You have considerable flexibility with the data that you want to replicate. For example, you can choose to have all rows and columns replicated or you can choose just a subset of these. By subsetting rows and columns, you transport across your network only the data that you want.

You can also clean, aggregate, or otherwise manipulate data. If you manipulate data, you can do so centrally at the source and then distribute the manipulated data, or you can manipulate the data when you replicate it, so that some targets get manipulated data and others do not.

**Related concepts:**

- “Introduction to DB2 replication and publishing—Overview” on page 1
- “SQL replication—Overview” on page 7
- “Which replication solutions should you use in common scenarios?” on page 33
- Chapter 6, “Comparison of SQL replication, Q replication, and event publishing,” on page 43

---

## Introduction to Q replication

Q replication is a replication solution that can replicate large volumes of data at very low levels of latency. Q replication captures changes to source tables and converts committed transactional data to messages. This data is sent as soon as it is committed at the source and read by Q replication. The data is not staged in tables. The messages are sent to the target location through WebSphere<sup>®</sup> MQ message queues, where the messages are read from the queues and converted back into transactional data. The transactions are then applied to your target tables with a highly parallelized method that preserves the integrity of your data.

You can use Q replication for a variety of purposes that require replicated data, including failover, capacity relief, geographically distributed applications, and data availability during rolling upgrades or other planned outages.

Sources and targets can be relational tables in DB2<sup>®</sup> UDB servers on Linux, UNIX<sup>®</sup>, Windows<sup>®</sup>, or z/OS<sup>™</sup>. You can replicate from a DB2 UDB source on any of these platforms to a DB2 UDB target on the same platform or any of the other supported platforms. To use WebSphere MQ message queues for transporting data, you need WebSphere MQ on your source and target systems.

You can replicate a subset of columns and rows from the source tables. All subsetting occurs at the source location so that only the data you want is transported across the network. If you want to perform data transformations, you can pass replicated data to your own stored procedures.

**Related concepts:**

- “Introduction to DB2 replication and publishing—Overview” on page 1
- “Q replication—Overview” on page 21

- “Which replication solutions should you use in common scenarios?” on page 33
- Chapter 6, “Comparison of SQL replication, Q replication, and event publishing,” on page 43

---

## Introduction to event publishing

Event publishing captures changes to source tables and converts committed transactional data to messages in an Extensible Markup Language (XML) format. Each message can contain an entire transaction or only a row-level change. These messages are put on WebSphere® MQ message queues and read by a message broker or other applications. You can publish subsets of columns and rows from source tables so that you publish only the data that you need.

You can use event publishing for a variety of purposes that require published data, including feeding central information brokers and Web applications, and triggering actions based on updates, inserts, or deletes to source tables.

Source tables can be relational tables in DB2® UDB servers on Linux, UNIX®, Windows®, or z/OS™. To make use of WebSphere MQ message queues for transporting data, have WebSphere MQ available on your source and on the system where messages are received.

Event publishing can be very useful in a variety of different applications. Consider a scenario in which changing prices and inventory are published to potential buyers. For example, a food wholesaler procures perishable food products such as bananas from world markets in bulk and sells them to grocery food retailers and distributors.

The value of bananas decreases the longer that they are in the warehouse. The wholesaler wants to inform its potential buyers of the changing price and inventory data and can set up event publishing to do just that. Each time the price changes, an XML message can be sent to potential buyers, informing them of the “price change event.”

Each buyer (retailer or distributor) wants to maximize profit. These organizations can determine when to buy the bananas based upon price, age (or time to spoil), and historical knowledge regarding how quickly they can sell a certain quantity. If they purchase too early, they will pay a higher price for the bananas and will not achieve the maximum profit. Buying too late will likely result in spoilage and again profit will not be maximized. Profit maximization can be achieved by timing the purchase appropriately. Applications can receive the event notification messages and generate purchase orders automatically at the right time to maximize profit.

### Related concepts:

- “Introduction to DB2 replication and publishing—Overview” on page 1
- “Event publishing—Overview” on page 37
- Chapter 6, “Comparison of SQL replication, Q replication, and event publishing,” on page 43

---

## Administrative interface for DB2 replication and publishing

The Replication Center is a graphical user interface that you can use to define, operate, and monitor your replication and publishing environments. It comes with the DB2<sup>®</sup> Administration Client and runs on Linux, UNIX<sup>®</sup>, and Windows<sup>®</sup> systems. The Replication Center provides a single interface for administering your replication environments on different platforms across multiple systems.

The following list describes some of the Replication Center's most useful features:

- Launchpads that show you step by step how to configure basic replication and publishing environments.
- Wizards that help you set up simple to highly customized replication and publishing configurations.
- Customizable profiles that let you create replication objects with schemas, names, and other attributes that conform to your own conventions and storage requirements.
- Reports that show you how your replication and publishing environments are performing.

If you prefer, you can define and modify your replication and publishing environments through a command-line tool, the ASN command-line processor (ASNCLP). You can then operate the programs for replication and publishing through system commands.

---

## Utilities for DB2 replication and publishing

Three utilities that come with the replication and publishing technologies help you monitor your replication and publishing environments and keep sources and targets synchronized:

### Replication Alert Monitor

The Replication Alert Monitor is a utility that checks the health of programs that are part of the replication and event publishing solutions. It checks for situations in which a program terminates, issues a warning or error message, reaches a threshold for a specified value, or performs a certain action. You tell the Replication Alert Monitor which situations to watch for. If any of them occur, the Replication Alert Monitor sends an e-mail message to the person or group of persons that are designated as the appropriate contacts for such a situation.

For example, you might have the Replication Alert Monitor notify you when a replication program is not running or has reached the maximum amount of memory that you expect it to use.

### Reconciliation utilities

Source and target tables can lose synchronization when, for example, a target table is changed by an application. The `tdiff` and `trepair` utilities can detect and repair differences between source and target tables in Q replication and SQL replication. By using them, you can avoid comparing tables manually and having to reload targets to resynchronize them with sources.

The `tdiff` utility generates a relational table of the differences between a source and target table. After you run the `tdiff` utility, you can run the

repair utility, which uses this table of differences to determine which rows to insert, update, or delete at the target table to synchronize it with the source table.

**Related concepts:**

- “Introduction to DB2 replication and publishing—Overview” on page 1



---

## Chapter 2. SQL replication

---

### SQL replication—Overview

The following topics describe the basic concepts involved in SQL replication. For more specific details, see the *IBM DB2 Information Integrator SQL Replication Guide and Reference*.

- Infrastructure for an SQL replication environment
- Registration of sources in SQL replication
- Subscription sets in SQL replication
- Capture of data from DB2 UDB sources in SQL replication
- Application of data to DB2 UDB targets in SQL replication
- Capture of data from non-DB2 sources in SQL replication
- Application of data to non-DB2 targets in SQL replication

The following figure shows a simple configuration in SQL replication. The topics listed above each highlight and describe different sections of this figure.

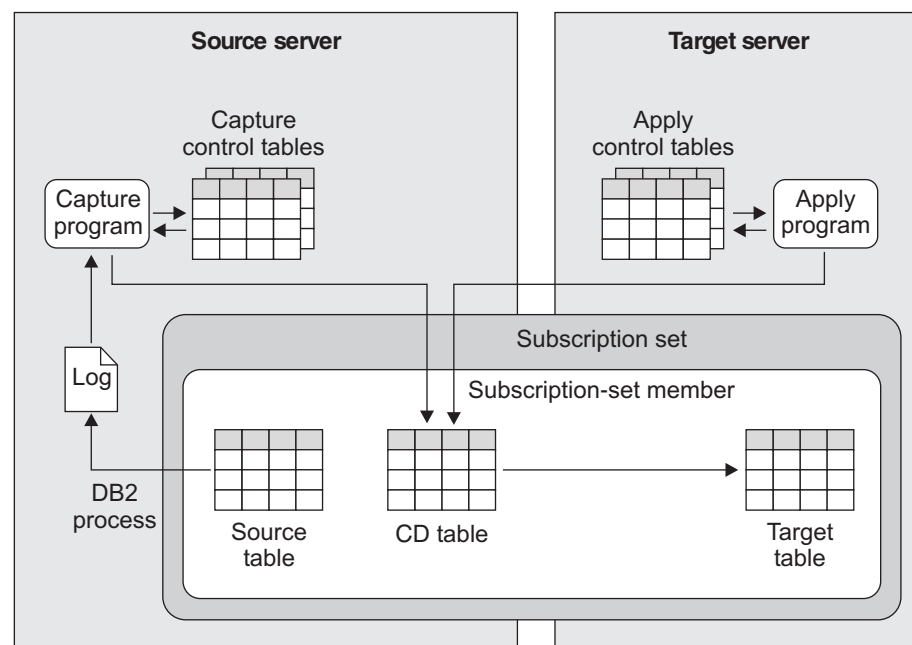


Figure 1. A simple configuration in SQL replication

#### Related concepts:

- “Application of data to non-DB2 targets in SQL replication” on page 19
- “Application of data to DB2 UDB targets in SQL replication” on page 17
- “Capture of data from non-DB2 sources in SQL replication” on page 18
- “Capture of data from DB2 UDB sources in SQL replication” on page 15
- “Infrastructure for an SQL replication environment” on page 8
- “Subscription sets in SQL replication” on page 11
- “Registration of sources in SQL replication” on page 9

---

## Infrastructure for an SQL replication environment

SQL replication allows you to replicate data from DB2<sup>®</sup> UDB sources to targets by using two programs: Capture and Apply. The Capture program runs on the source system. The Capture program reads DB2 recovery logs for changed source data and saves the committed changed data to staging tables. The Apply program typically runs on the target system. The Apply program retrieves captured data from staging tables and delivers the data to targets. Both programs use a set of DB2 tables to track the information that they require to do their tasks and to store information that they generate themselves, such as information that you can use to find out how well they are performing. You create these tables before you tell SQL replication what your replication sources and targets are.

The Capture program uses a set of DB2 UDB tables called *Capture control tables*. These tables contain information about replication sources and the current position of the Capture program in the DB2 recovery log. In most cases, the control tables for a Capture program need to be on the same DB2 UDB server as the sources associated with the program.

You can run multiple Capture programs on the same DB2 UDB server. Each Capture program uses its own set of Capture control tables. The schema associated with a set of Capture control tables identifies the Capture program that uses those control tables. This schema is called a *Capture schema*.

If your sources are located on a non-DB2 relational database, triggers are used to capture changes to them, though a set of control tables is still required.

The Apply program uses a set of DB2 UDB tables called *Apply control tables*. These tables contain information about your targets and where their corresponding sources are located. The control tables for the Apply program usually reside on the system where the Apply program runs. Unlike with the Capture program, you can create multiple Apply programs that use the same set of control tables. Each Apply program is identified in these control tables by a name called an *Apply qualifier*.

For each supported operating system or platform that is supported by DB2 UDB and for each non-DB2 relational database management system that is supported by SQL replication, you can set your own default specifications for the table spaces that will be used by control tables.

The following figure shows the infrastructure for a simple configuration in SQL replication.



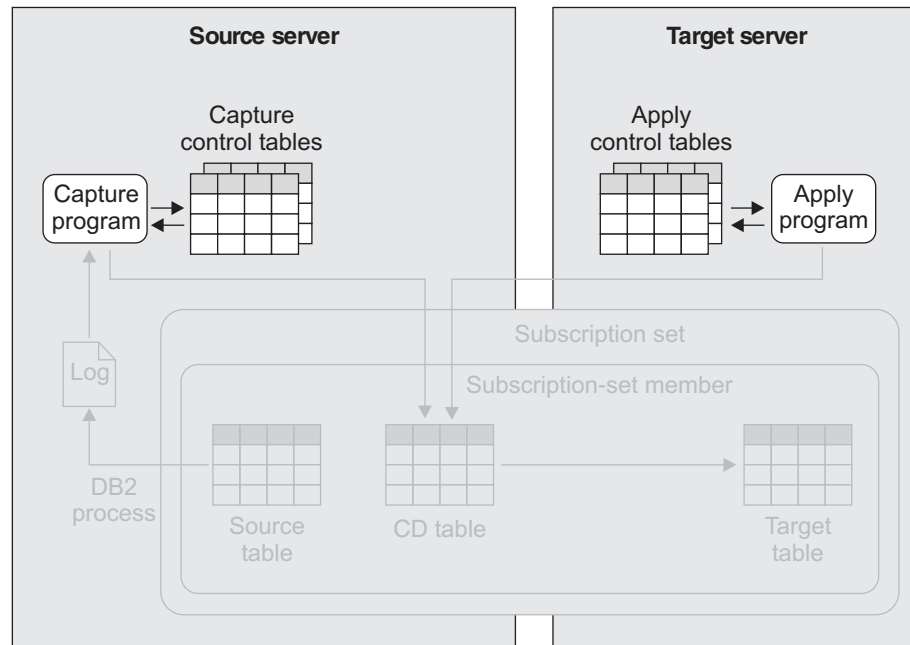


Figure 2. Infrastructure for a simple configuration in SQL replication. You create a set of DB2 relational tables on the source server called Capture control tables. Information about sources is written to these tables. The Capture program, which runs on the source server, uses this information to know what data it is supposed to capture. Information about targets will go in the Apply control tables, which are typically on your target server. The Apply program, which is also typically on your target server, uses this information to know which targets it is supposed to write data to.

**Related concepts:**

- “SQL replication—Overview” on page 7

## Registration of sources in SQL replication

When you know what you want your replication sources to be, you tell SQL replication about them by *registering* them. You can register sources that are DB2<sup>®</sup> UDB tables and views, or tables on non-DB2 relational databases.

**Note:** When you register source tables on non-DB2 relational databases, you use SQL replication together with DB2 Relational Connect. You map your source database to a federated database and you create nicknames for each source table.

You can specify one of three methods for replicating a table. The first is called *full-refresh replication*, and the second and third are types of *change-capture replication*:

**Replicate by refreshing the target table. Do not capture changes.**

At intervals that you can specify later, the Apply program invokes one or more DB2 UDB utilities to delete the content of the corresponding targets for the registered table and load the targets with content derived from the registered table. When you register the table, you need to provide only its schema and name. You can choose to replicate this way if a target is small or if you want to replicate infrequently. For this method of replication, no Capture program or Capture triggers are involved.

**Capture a row whenever the value changes in one or more columns that you want to replicate.**

This method, the most common way of using SQL replication, reduces the amount of data that is replicated when changes frequently affect only unregistered columns. For example, assume that you have 100 columns in your table and you register 50 of those columns for replication. Whenever a change is made to one of the 50 registered columns in your table, SQL replication captures the values of those 50 columns for the row in which the change occurred. If a row is changed, but the changes are limited to other 50 columns, no data is captured for replication.

**Capture a row whenever the value changes in any column of the table.**

This method is often chosen when changes in the table almost always occur in registered columns only, or when information is replicated for auditing purposes. For example, you want to know only which rows are changed in a table. Assume that you have 100 columns in your table and you register the primary key columns for replication. Any time a change is made to any of the columns in your table, SQL replication captures the values of the primary key columns for the row in which the change occurred.

When you register a table for either type of change-capture replication, you provide SQL replication with the same information and SQL replication creates the same objects:

- **If your source is a DB2 UDB table:**

You tell SQL replication the schema and name of the source table, as well as the columns that you want to replicate from it. The columns that you choose are called *registered* columns. They are also referred to as *after-image* columns because they contain values resulting from changes made to the source table. For example, if in one row of a registered table the value in registered column A is changed from 25 to 30, 30 is the after-image.

You can also choose to replicate the *before-images* of your registered columns, which are the values that existed in the columns before changes were made to the source table. In the example given in the preceding paragraph, 25 is the before-image. Before-images are useful for a number of purposes, including enabling the applying program to find rows in the target when key values have changed.

SQL replication creates a change-data (CD) table in which to record committed changes to the table. For every column that you register in the source table, there is an identical column in the CD table. If you choose to replicate before-images for a column, a column is created in the CD table to record those before-images.

- **If your source is a non-DB2 UDB table:**

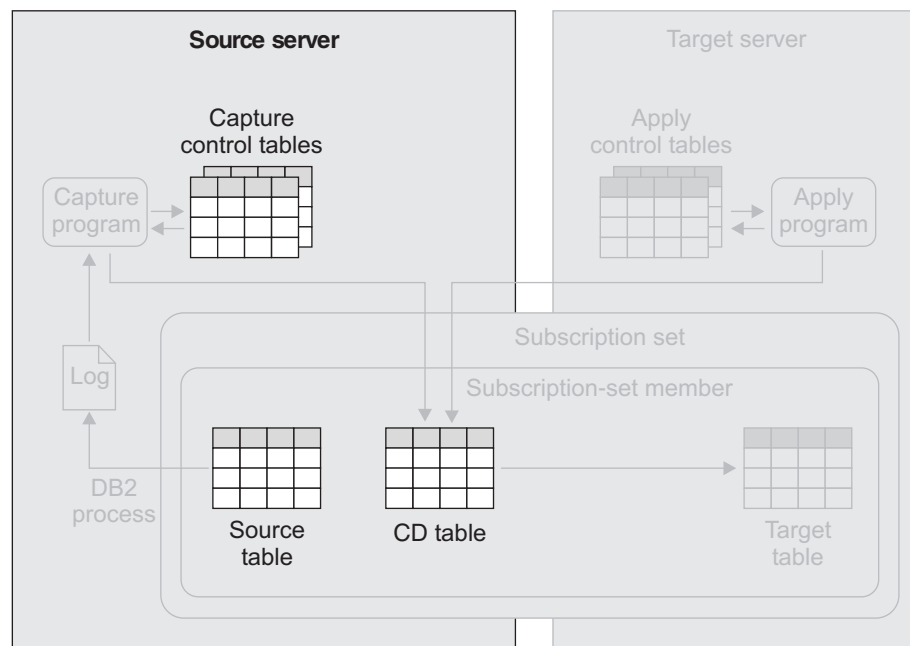
You tell SQL replication the schema and name of the nickname for the source table, as well as the columns that you want to replicate from the source table. As is the case when you register DB2 UDB tables, the columns that you choose are called registered columns, and you can replicate after-images and before-images. SQL replication creates three triggers, called *Capture triggers*, directly on the source table in the non-DB2 relational database: one trigger for inserts, another trigger for updates, and one more trigger for deletes. SQL replication also creates a staging table for the source table. The staging table is called a consistent change-data (CCD) table and is populated by the triggers. For example, when an update occurs in the source table, the update trigger on that table stores the

update as a record in the corresponding CCD table. Only committed transactional data is stored in the CCD table. A nickname for the CCD tables is created on the federated database.

For every source server, you can set your own default naming conventions for CD and CCD tables. You can also set your own defaults for the table spaces used for CD and CCD tables.

For more information about registering views, see the *IBM DB2 Information Integrator SQL Replication Guide and Reference*.

The following figure shows the objects that are involved when you register a source table.



*Figure 3. Registering a source table.* When you register a source table for replication, SQL replication stores the registration information in the Capture control tables. SQL replication also creates a CD table that the Capture program will use to record changes to the source table.

**Related concepts:**

- “SQL replication—Overview” on page 7

---

## Subscription sets in SQL replication

After you register your sources, you create *subscription sets*, in which you pair your sources with targets. Each source-target pair is referred to as a member of the subscription set in which it is created. You can use subscription sets to schedule the replication of data in one or more source-target pairs from one source server to one target server. The Apply program coordinates this replication so that it applies data to targets in a manner consistent with the original transactions on the source server.

For each subscription set that you create, you tell SQL replication the following details:

**Which Apply program to use for processing the subscription set**

When you create a subscription set, you associate it with an Apply program by selecting an Apply qualifier. Each Apply program can process more than one subscription set.

**Where the source tables or views are located**

The sources in the members of a subscription set are all located on the same server.

**Where the target tables are located**

The targets in the members of a subscription set are all located on the same server.

**How often to replicate data from the sources to the targets**

You can choose to replicate data in a subscription set at regular intervals, continuously, or when an event occurs.

**Interval timing**

This method, also known as relative timing, is the simplest method of scheduling your subscription set. The interval is approximate, which depends on the workload or system resource available at the time. The interval pertains to all the subscription-set members within the subscription set. Therefore, all the tables will replicate within the subscription set at the each interval.

**Continuous replication**

This method lets the Apply program replicate data in a subscription set as frequently as it is able, depending on its workload and available resources.

**Event timing**

This method lets an application or a user determine when to start replicating. The application or user inserts a row into an Apply control table called the events table. When the Apply program sees that row, it begins replicating.

**Whether to use data blocking**

Data blocking allows you to specify how many minutes worth of change data the Apply program can replicate at a time. The number of minutes that you specify determines the size of the data block. If there is a backlog of change data that is greater than the size of the data block, the Apply program processes the subscription set by replicating a block of data at a time until all of the current change data is replicated. One benefit of data blocking is that by retrieving smaller sets of data, the Apply program can lessen both the network load and the temporary space required for the retrieved data. Another benefit is that if there is an error, the Apply program has to rollback only the current block of data rather than the entire set of change data.

For example, suppose that you use SQL replication to populate the tables of a data warehouse every 24 hours. Rather than replicate in one block the entire set of changes made to your source tables every 24 hours, the Apply program can replicate 20 minutes worth of changes at a time until the entire 24-hour block of data is replicated. You can adjust the size of the block to match your own requirements and restrictions.

**Whether to issue one COMMIT for all applied data or to issue interim commits**

The Apply program can apply data in table mode or in transaction mode. In table mode, the Apply program processes the fetched changes for each table separately, and then issues a single commit after all data is applied.

In transaction mode, the Apply program applies the fetched changes to all of the target tables, in the order that the changes occurred in their corresponding transactions on the source server. The Apply program commits these transactions at transaction boundaries. You specify how many transactions to apply before each commit.

### **Whether to transform data in the subscription set with SQL scripts or stored procedures**

When data in the source tables of the subscription set is replicated to corresponding targets, you can use SQL scripts or stored procedures to transform that data.

For each source-target pair that you create within a subscription set, you specify whether to create a new target table or use an existing one. If you create a new target table, you can specify one of the following types:

#### **User copy**

This read-only target type is the most common for basic data replication. The structure of this type of target table can be the same as the source table, or only a subset of the source columns. Before-images or after-images and calculated columns can also be defined. This target type requires replicated columns that uniquely describe each row of the target table. These columns can be either a unique index or primary key.

#### **Point-in-time**

The structure of this read-only type of target table is similar to that of a user copy target table. A timestamp column is added so that the Capture program can indicate when data was committed in the source table. You can select this target type if you want to keep track of the time when the changes were applied to the target table.

#### **Aggregate**

This is a read-only target table that summarizes the entire contents of a source table or its changed data. The target columns are defined from the SQL column functions such as SUM, COUNT, MIN, MAX, and AVG. These columns contain the computed value of the SQL function instead of the actual source data. A timestamp is included to indicate when the Apply program performed the aggregation. The two types of aggregation target tables are:

- **Base aggregate.** Summarizes the entire contents of the source table every time the source table is replicated. For example, you could use this target type if you want to keep track of year-to-date summaries or average sales by salesperson or region.
- **Change aggregate.** Summarizes the changes that occurred at the source table since the source table was last replicated. For example, you could use this target type to track monthly sales totals by salesperson, region, or customer.

#### **CCD**

This read-only target type contains committed changes that occurred at the source, with replication control fields to determine the type operation from the source table (insert, update, and delete). There are different types of CCD tables and the type that you should use depends on your goals and your replication requirement. With a CCD table, you can track the history of source changes in various ways, depending on how you define the CCD table. For example, you can track before and after comparisons of the data, when changes occurred, and which user ID made the update to the source table.

You can also set up a multi-tier configuration, where your target CCD table acts as a source to other target tables. Your source table is the first tier, your CCD the second tier, and the targets mapped to your CCD the third tier. One reason to set up a multi-tier replication environment is to provide stable sources for the third-tier targets. Because you can collect changes from tier 1 in CCD tables at tier 2, you can reduce the number of changes replicated to the targets at tier 3. You can also avoid many of the database connections to your source system, thus moving the connection cost to the second tier.

## Replica

This target type is used in update-anywhere replication. The other target types are read-only, but this target type allows writes which can be replicated back to the source table, which acts as a master table. If any conflicts occur with the data in the master table, the change replicated from the replica will be rejected. Replica target tables cannot be defined at non-DB2 relational databases.

For each source-target pair, you also specify:

- **Which source columns to replicate to the target.** Although you can subset columns when you register your source, you can subset further when you create a subscription-set member.
- **How to map columns from the source to columns in the target.** If your target table already exists in the target location, you tell SQL replication how the registered columns in your source table or view map to the columns in the target table. If you want SQL replication to create the target table for you, the registered source columns are automatically mapped to the columns in the target table.

In both cases, you can transform data by mapping a source column to a target column with a corresponding data type or by mapping a source column to a calculated column.

- **A predicate for replicating a subset of rows to the target.** You can subset the rows that you replicate from the source.
- **A method for loading the target table.** In most cases, the data from the source table is loaded into the target table so that the target table is identical to the source table before replication begins. SQL replication allows for two methods of loading a target table: automatic loads and manual loads.

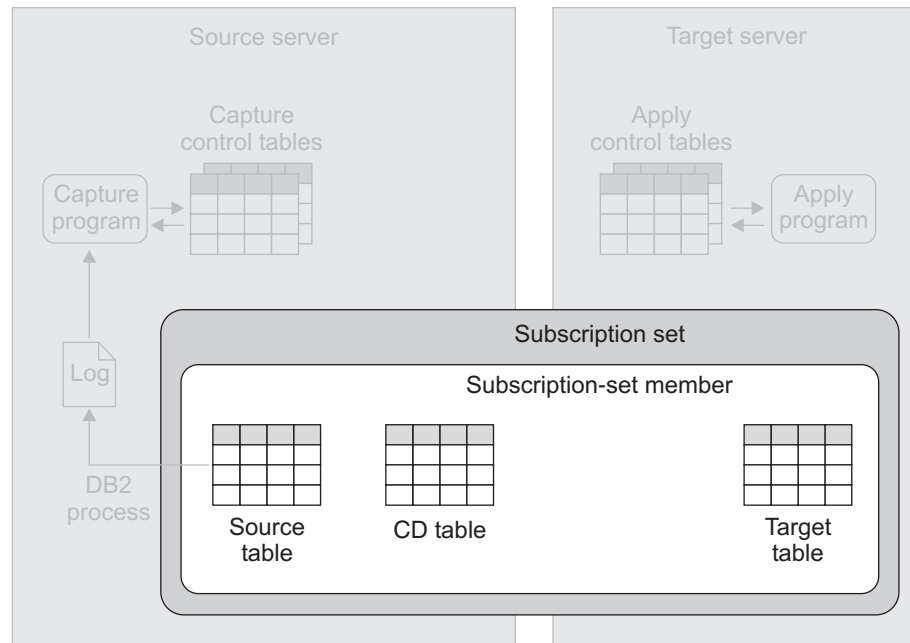
For automatic loads, you can tell the Apply program to call one utility or a pair of utilities. The Apply program can call the LOAD FROM CURSOR option of the LOAD utility, the EXPORT and LOAD utilities, or the EXPORT and IMPORT utilities, depending on the platform on which the Apply program is running. You can also tell the Apply program to choose the option that is most appropriate for the target table.

If you decide to load the target table manually, you can choose any method that you prefer. After performing the load, you notify the Apply program that the target table is ready and replication begins.

When you create a subscription-set member, you can also decide whether to replicate only a subset of the registered source columns and whether to replicate only a subset of the rows from the source.

For every target server, you can set your own default naming conventions for target tables. You can also set your own defaults for the table spaces that are used for your target tables.

The following figure shows a subscription set in a simple configuration in SQL replication.



*Figure 4. Creating a subscription set.* You map sources to targets as part of creating subscription sets. A subscription set groups together one or more source-target pairs, also called subscription-set members. In the figure, a source table is mapped to a target table. The data replicated to the target is first staged in the CD table for the source.

**Related concepts:**

- “SQL replication—Overview” on page 7

---

## Capture of data from DB2 UDB sources in SQL replication

After you register your DB2<sup>®</sup> UDB replication sources and create subscription sets to pair your sources with targets, you are ready to begin capturing the changes made to your sources. You use a program called the Capture program to do this.

After you start the Capture program and that program receives signals from the Apply program to indicate that the sources and targets are synchronized, the Capture program reads the DB2 log sequentially for changes to the source tables in which you are interested. If it reads a change to one of your source tables, the Capture program adds the change to the corresponding database transaction that it is retaining in memory. Transactions in memory are potentially subsets of the corresponding transactions in the log; they contain only changes to your source tables. The Capture program collects changes in memory until it reads either a ROLLBACK or a COMMIT statement for the transaction in which those changes are made. Upon reading a ROLLBACK statement, the Capture program erases from memory the changes associated with the rolled-back transaction. Upon reading a COMMIT, the Capture program stores the changes associated with the committed transaction.

For example, suppose that you registered Table A and Table B as replication sources. For each of these tables, SQL replication creates a CD table as part of the registration process. After you start the Capture program and after that program receives signals from the Apply program to indicate that the target tables are

synchronized with the source tables, the Capture program reads the DB2 logs for changes to those source tables. Application 1 makes a series of changes to Table A. Each change is recorded in the DB2 logs. The Capture program collects these changes in memory. Application 1 issues a ROLLBACK statement. When the Capture program reads this statement, it erases from memory the changes that are associated with that transaction.

Application 2 makes a series of changes to Table B. As before, the Capture program collects these changes in memory. Application 2 then issues a COMMIT statement. When the Capture program reads this statement, it appends a copy of each change to the CD table for Table B.

At least one Capture program must run on each server where replication sources are located. However, you might want to have more than one Capture program on a single server for the following reasons:

**You want to increase parallelism**

You can use multiple Capture programs to parallelize traffic. Multiple Capture programs, which could also be helpful in large sysplexes, can improve performance and achieve higher throughput. The trade-off is additional CPU overhead associated with multiple log readers. Using multiple Capture programs also requires more DB2 UDB connections.

**You need to meet different replication requirements**

You can create multiple Capture programs to direct the flow of source changes toward different uses. For example, if you need to replicate a large source table with very low latency and your current Capture program is already capturing data for a large number of source tables, you could use another Capture program to capture changes made only to that table.

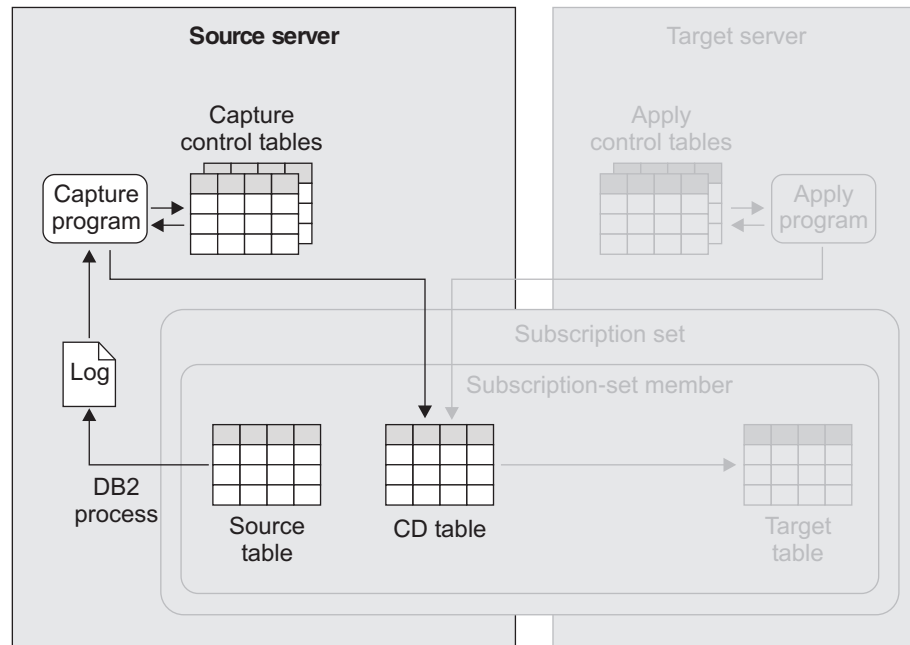
For example, if you need to replicate one set of source tables with very low latency, and you need to replicate another set of source tables to feed a data warehouse on a daily basis, the performance and tuning requirements might be very different. You could use a Capture program for each of these separate sets of source tables, and tune each Capture program to the needs of the respective applications.

**You need to use separate encoding schemes**

On z/OS™ systems, you can use multiple Capture programs to work with UNICODE or EBCDIC encoding schemes separately.

The following figure shows the Capture program capturing data in a simple configuration in SQL replication.





*Figure 5. Capturing data from a DB2 UDB source table.* The Capture program reads the DB2 recovery log for changes made to your sources and then stages those changes in relational tables. In the figure, the Capture program reads the DB2 log for changes to the source table and stages committed transactional data in the source’s change-data (CD) table. The Capture program continuously updates its control tables with information that you can use to monitor its progress.

**Related concepts:**

- “SQL replication—Overview” on page 7

---

## Application of data to DB2 UDB targets in SQL replication

To start replicating data from sources to targets, you start the Apply program. After you do so, the Apply program begins processing the subscription sets that you assigned to it. The Apply program processes any active subscription sets one at a time, according to the scheduling or event criteria that you specified when you created the subscription sets. For example, if you chose to replicate subscription set Set\_One every 60 minutes, the Apply program will replicate that subscription set as close to every 60 minutes as possible. If the Apply program has other work to finish when the 60-minute interval elapses, Set\_One will be replicated as soon as the Apply program finishes its other work.

The following list describes how the Apply program processes each member of a subscription set.

**For sources that you registered for change-capture replication**

The first time that the Apply program processes the corresponding subscription set, you can have the Apply program populate the targets with the content of the sources. You can tell the Apply program to call one or more utilities, such as the EXPORT and LOAD utilities, to populate the targets. The utilities that the Apply program chooses depends on the platform on which the Apply program is running.

Then, at the time intervals that you specified when you created your subscription sets or whenever an event occurs, the Apply program reads the CD tables of your sources for rows that were inserted into the CD table

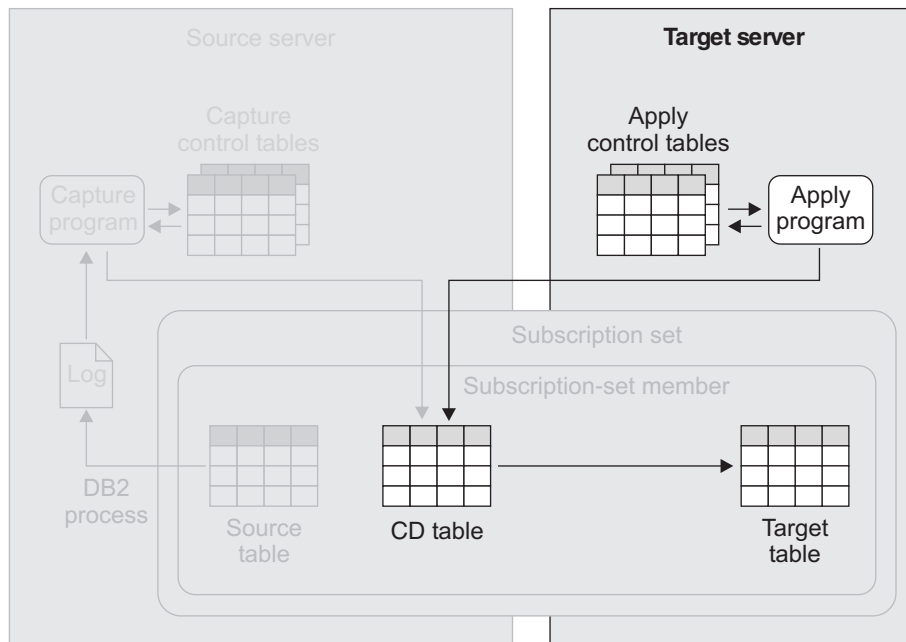
since the Apply program last looked. The rows in the CD tables indicate whether they are records of deletes, updates, or inserts to the corresponding sources.

The Apply program uses the data from the CD table to insert, update, and delete rows in the target. Predicates are used to identify the rows to be updated or deleted.

**For sources that you registered for full-refresh replication**

At intervals that you specify the Apply program populates the targets with the content of the sources. You can tell the Apply program to call one or more utilities, such as the EXPORT and LOAD utilities, to refresh your target tables. The utilities that the Apply program chooses depends on the platform on which the Apply program is running.

The following figure shows the Apply program applying data in a simple configuration in SQL replication.



*Figure 6. Applying data to DB2® UDB targets.* The Apply program reads data from staging tables and makes the appropriate changes to targets. In the figure, the Apply program reads data from the source table's CD table and then makes the appropriate changes to the target table. The Apply program continuously updates its control tables with information that you can use to monitor its progress.

**Related concepts:**

- “SQL replication—Overview” on page 7

---

## Capture of data from non-DB2 sources in SQL replication

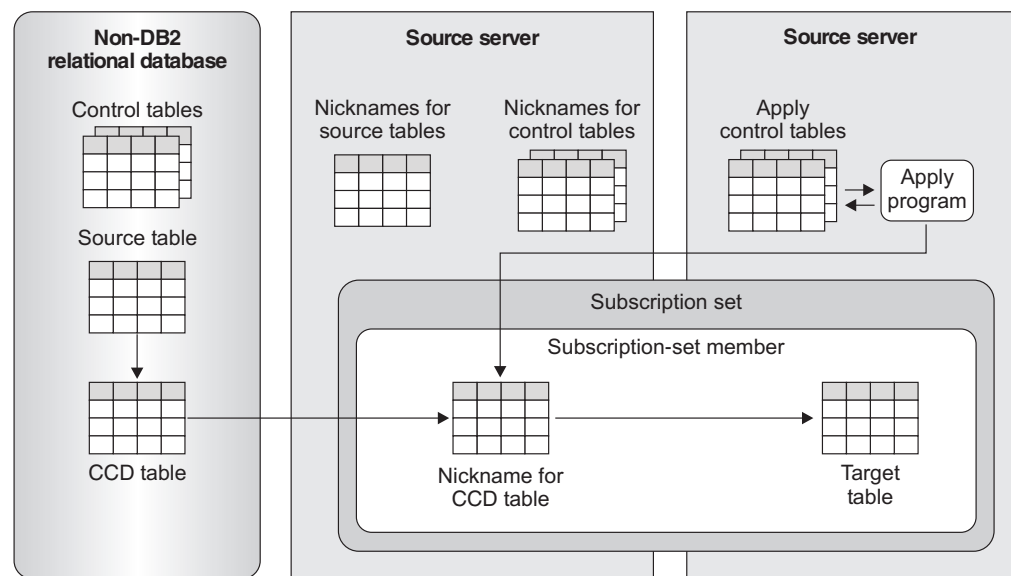
The Capture program is not involved in capturing changes made to tables on non-DB2 relational databases. Instead, when you register such a table, SQL replication creates three after triggers (called *Capture triggers*) on the table (an INSERT trigger, a DELETE trigger, and an UPDATE trigger) and also creates a consistent change-data (CCD) table, which is similar to a CD table because it stores records of changes made to the source table. SQL replication also creates corresponding nicknames for the CCD tables on the same DB2® Information

Integrator (II) federated server where the nicknames for the source tables are located. The Apply program accesses the CCD tables through these nicknames.

When one of these triggers is activated by a change to the source table, the trigger records the change in the CCD table. If the application transaction is rolled back, the new records in the staging table are rolled back, too. Only changes that are part of committed transactions are replicated.

SQL replication also creates a fourth trigger on a non-DB2 relational database that is used as a source database. This is the pruning trigger that deletes replicated records from the CCD tables after the Apply program successfully replicates changes to target tables.

The following figure shows how changes to a source table on a non-DB2 relational database are replicated to a DB2 UDB target table.



*Figure 7. Replicating data from non-DB2 sources to DB2 UDB targets.. Triggers populate a staging table called a consistent-change-data (CCD) table with records of changes made to a source table. The Apply program reads the CCD table through that table's nickname on the DB2 II federated database. The Apply program then makes the appropriate changes to the target table.*

**Related concepts:**

- "SQL replication—Overview" on page 7

## Application of data to non-DB2 targets in SQL replication

As when replicating to DB2<sup>®</sup> UDB targets, the Apply program can replicate to non-DB2 targets in two ways: through change-capture replication and through full-refresh replication.

**Change-capture replication for targets on a non-DB2 relational database**

If the source is on a DB2 server, the Apply program reads the change-data (CD) table for the source. The Apply program then applies the changes to the target by accessing the target table's nickname on a federated database.

If the source is on a non-DB2 relational database, the Apply program reads the CCD table for the source by accessing the CCD table's nickname. It

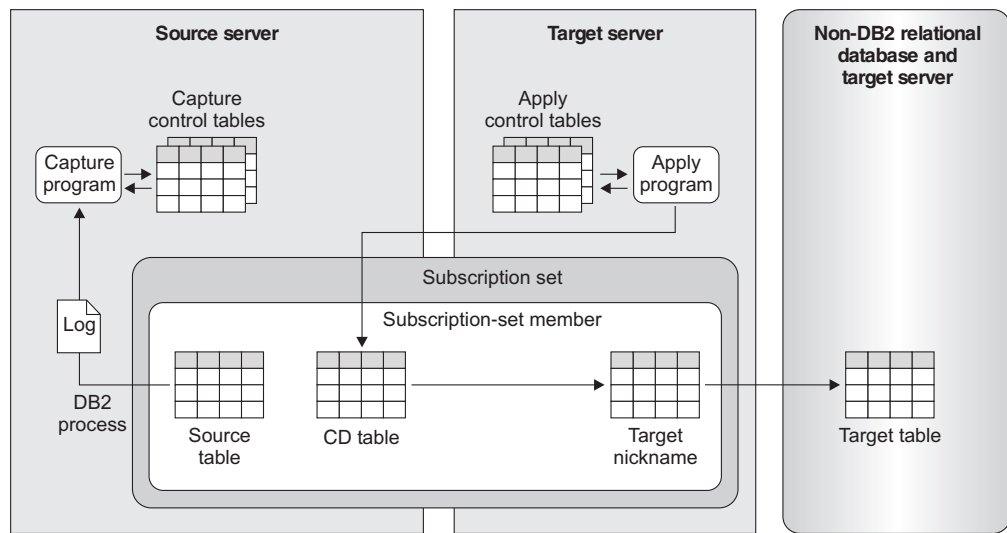
then applies the changes to the target table by accessing the target table's nickname. The nickname for the CCD table and the nickname for the target table can be in the same DB2 federated database or in different ones.

**Full-refresh replication for targets on a non-DB2 relational database**

If the source is on a DB2 server, the Apply program reads the source directly, deletes the rows of the target table by accessing the target table's nickname, and inserts the source rows into the target table by accessing the target table's nickname.

If the source is on a non-DB2 relational database, the Apply program reads the rows of the source by accessing the source's nickname. It then deletes the rows of the target by accessing the nickname, and inserts the source rows into the target table by accessing the target table's nickname. The nickname for the source table and the nickname for the target table can be in the same DB2 federated database or in different ones.

The following figure shows how changes to a DB2 UDB source table are replicated to a table on a non-DB2 relational database.



*Figure 8. Replicating from DB2 UDB sources to non-DB2 targets. The Capture program populates a CD table with the changes made to the corresponding source table. The Apply program reads the data in the CD table and applies the changes to the target table through the target's nickname on the DB2 Information Integrator federated database.*

**Related concepts:**

- "SQL replication—Overview" on page 7

---

## Chapter 3. Q replication

---

### Q replication—Overview

The following topics describe the concepts involved in Q replication. For more specific details, see the *IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference*.

- Infrastructure for a Q replication environment
- Sources and targets in Q replication
- Capture of data in Q replication
- Application of data to targets in Q replication
- Types of replication in Q replication

The following figure shows a simple configuration in Q replication. The topics listed above each highlight and describe different sections of this figure.

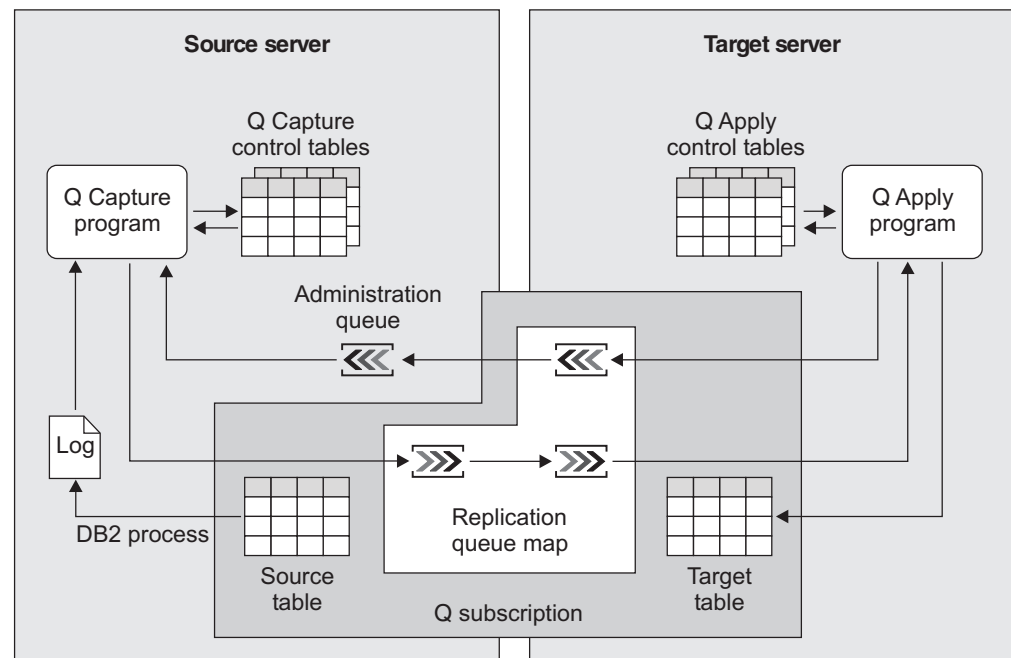


Figure 9. A simple configuration in Q replication

#### Related concepts:

- “Application of data to targets in Q replication” on page 26
- “Capture of data in Q replication” on page 25
- “Infrastructure for a Q replication environment” on page 22
- “Sources and targets in Q replication” on page 23
- “Types of replication in Q replication” on page 27

---

## Infrastructure for a Q replication environment

Q replication allows you to replicate committed transactional data from DB2<sup>®</sup> UDB sources to targets by using two programs: Q Capture and Q Apply. The Q Capture program runs on the source system. The Q Capture program reads DB2 recovery logs for changed source data and writes the changes to WebSphere<sup>®</sup> MQ queues. The Q Apply program runs on the target system. The Q Apply program retrieves captured changes from queues and writes the changes to targets. Both programs use a set of DB2 tables to track the information that they require to do their tasks and to store information that they generate themselves, such as information that you can use to find out how well they are performing. You create these tables before you tell Q replication what your replication sources and targets are.

The Q Capture program uses a set of control tables called *Q Capture control tables*. These tables contain information about your replication sources, the targets that correspond to them, and which WebSphere MQ queue manager and queues are being used by the Q Capture program. These tables also contain data that you can use to check and monitor the Q Capture program's performance, such as data about the Q Capture program's current position in the DB2 recovery log.

You can run multiple Q Capture programs on the same DB2 UDB server. Each Q Capture program uses its own set of control tables. The schema associated with a set of Q Capture control tables identifies the Q Capture program that uses those control tables. This schema is called a *Q Capture schema*.

The Q Apply program uses a set of control tables called *Q Apply control tables*. These tables contain information about targets and where their corresponding sources are located, as well as information about which WebSphere MQ queue manager and queues are being used by the Q Apply program. Like the Q Capture control tables, these tables also contain data that you can use to check and monitor the Q Apply program's performance.

As with Q Capture programs, you can run multiple Q Apply programs on the same DB2 UDB server. Each Q Apply program uses its own set of control tables. The schema associated with a set of Q Apply control tables identifies the Q Apply program that uses those control tables. This schema is called a *Q Apply schema*.

The following figure shows the infrastructure for a simple configuration in Q replication.

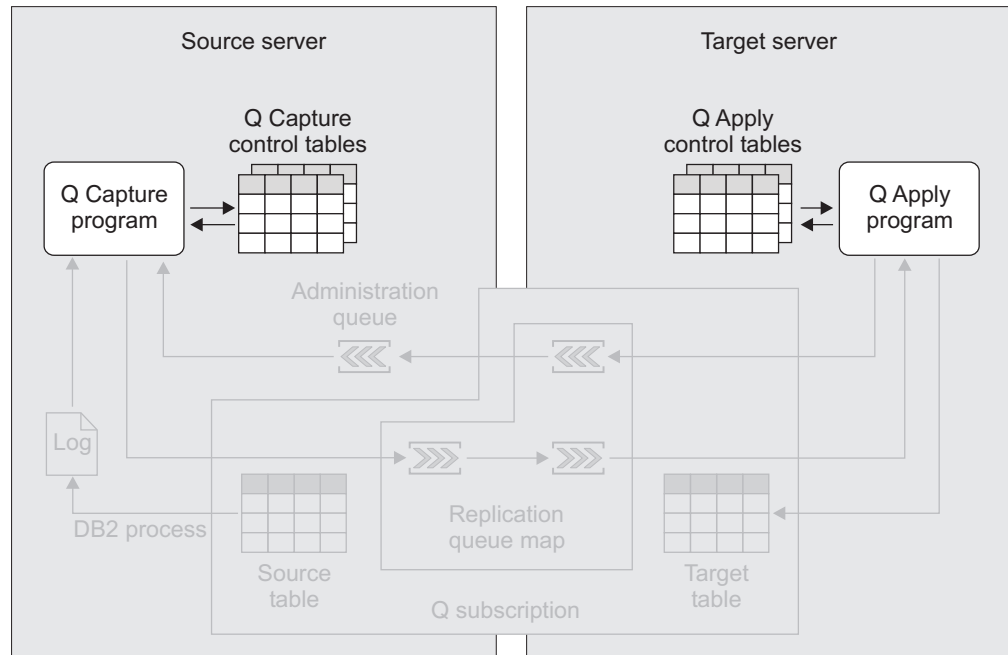


Figure 10. Infrastructure for a simple configuration in Q replication. You create a set of DB2 relational tables on the source server called Q Capture control tables. Information about your sources and targets is written to these tables. The Q Capture program, which runs on the source server, uses this information to know what data it is supposed to capture and send to the Q Apply program. Information about your sources and targets goes in the Q Apply control tables, which are on your target server. The Q Apply program, which runs on the target server, uses this information to know which target tables it is supposed to write data to.

**Related concepts:**

- “Q replication—Overview” on page 21

---

## Sources and targets in Q replication

You pair source tables with targets by means of *Q subscriptions*. When you create a Q subscription, you specify the following primary attributes:

**The source server and source table**

A source server is a DB2® UDB server that contains at least one set of Q Capture control tables. You can use a DB2 UDB relational table located on a source server as your source table.

**The target server and target table or stored procedure**

A target server is a DB2 UDB server that contains at least one set of Q Apply control tables. You can use a DB2 UDB relational table located on a target server as your target table, or you can tell Q replication to define the table for you when you create the Q subscription.

You can use a stored procedure as a target if you want to transform replicated data.

**Which source columns to replicate to the target table or stored procedure**

You can subset the columns that you replicate from the source.

**How to map the source columns to the target columns or parameters in a stored procedure**

If your target table already exists in the target location, you tell Q replication how the selected columns in your source table correspond to

the columns in the target table. If you want Q replication to create the target table for you, the selected source columns are automatically mapped to the columns in the target table.

If you are replicating to a stored procedure, you tell Q replication how the selected columns in your source table correspond to the parameters in your stored procedure.

#### **A predicate for replicating a subset of rows to the target table**

You can also subset the rows that you replicate from the source.

#### **The method of loading the target table**

In most cases, the data from the source table is loaded into the target table so that the target table is identical to the source table before replication begins. Q replication allows for two methods of loading a target table after you start a Q subscription: automatic loads and manual loads.

For automatic loads, the Q Apply program manages the loading of target tables. The Q Apply program calls one utility or a pair of utilities to perform the load. You can tell the Q Apply program to call the LOAD FROM CURSOR option of the LOAD utility, the EXPORT and LOAD utilities, or the EXPORT and IMPORT utilities, depending on the platform on which the Q Apply program is running. You can also tell the Q Apply program to determine itself which of these options is most appropriate for the Q subscription.

For manual loads, you handle the loading of target tables, and then signal the replication programs when loading is done.

You can also specify that when you start a Q subscription, you want replication to begin immediately without the need for coordination of the loading process with the Q Capture and Q Apply programs. You might want to do this when you have added a large number of new tables to your replication environment. In this case, you might select a down period, such as the weekend, quiesce the source tables, load the target tables yourself, and then start the Q Capture and Q Apply programs.

#### **The WebSphere® MQ message queues to use for transporting messages**

The Q Capture program puts messages on a queue called a *send queue*. The Q Apply program receives the messages on a queue called a *receive queue*. (These names are used only in Q replication, not in WebSphere MQ.) You choose a *replication queue map* to tell Q replication which send queue and receive queue you want to use for the messages for your Q subscription.

The following figure shows the parts of a Q subscription in a simple configuration in Q replication.



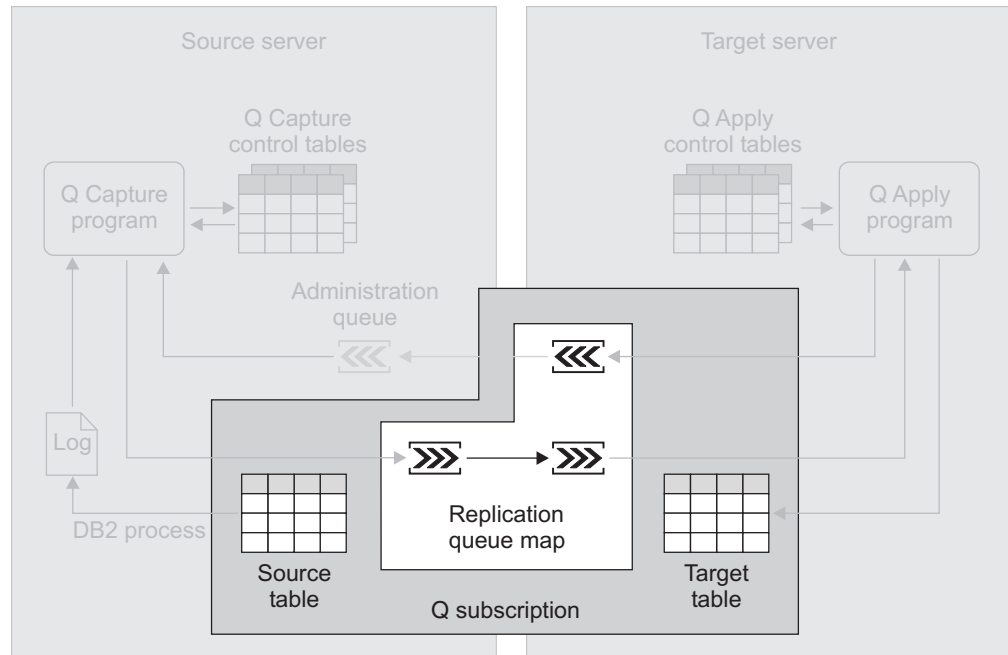


Figure 11. Creating a Q subscription. You map sources to targets by creating Q subscriptions. One Q subscription maps one source table to one target table. Data for the Q subscription is replicated across a send queue and a receive queue, both of which are part of a replication queue map.

**Related concepts:**

- “Q replication—Overview” on page 21

## Capture of data in Q replication

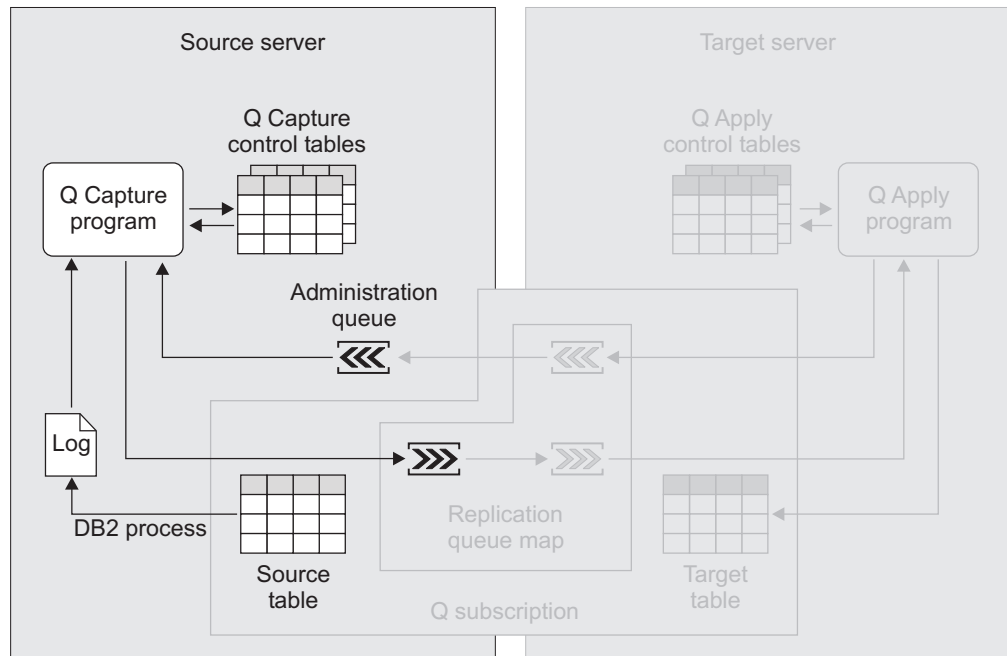
The Q Capture program reads the DB2® log sequentially for changes to source tables. If it reads a change to one of your source tables, the Q Capture program adds the change to the corresponding database transaction that it is retaining in memory. Transactions in memory are therefore potentially subsets of the corresponding transactions in the log because they contain only changes to the source tables in your Q subscriptions. When the Q Capture program reads the COMMIT statement for a transaction, it converts the transaction into a message and puts the messages on a send queue.

For example, you create two Q subscriptions for two different source tables on the same source server: QSUB1 and QSUB2. Both Q subscriptions use the same replication queue map. The Q Capture program reads a COMMIT for a database transaction that involves changes to the source tables in both Q subscriptions. The Q Capture program converts the changes into a message and writes the message to the send queue that is part of the replication queue map.

You can run more than one Q Capture program on the same source server. Although one Q Capture program can capture changes made to many sources and send those changes to many target servers, in some situations you might benefit from having more than one Q Capture program. For example, You can use multiple Q Capture programs to parallelize traffic. Multiple Q Capture programs, which could also be helpful in large sysplexes, can improve performance and

achieve higher throughput. The trade-off is additional CPU overhead associated with multiple log readers. Using multiple Q Capture programs also requires more DB2 UDB connections.

The following figure shows the Q Capture program capturing data in a simple configuration in Q replication.



*Figure 12. Capturing data in Q replication.* The Q Capture program reads the DB2 recovery log for changes made to your sources and converts committed transactional data to messages, which it puts on a send queue. The Q Capture program continuously updates its control tables with information that you can use to monitor its progress.

**Related concepts:**

- “Q replication—Overview” on page 21

---

## Application of data to targets in Q replication

The Q Apply program reads the messages that contain committed transactional data for Q subscriptions. These messages arrive at the target server on receive queues. The Q Apply program converts the messages to SQL and applies the transactions to the relevant target tables. The Q Apply program is multi-threaded and can apply several transactions concurrently, whenever those transactions are not dependent on one another. If transactions are dependent on one another, the Q Apply program applies them in the order in which they were committed at the source server.

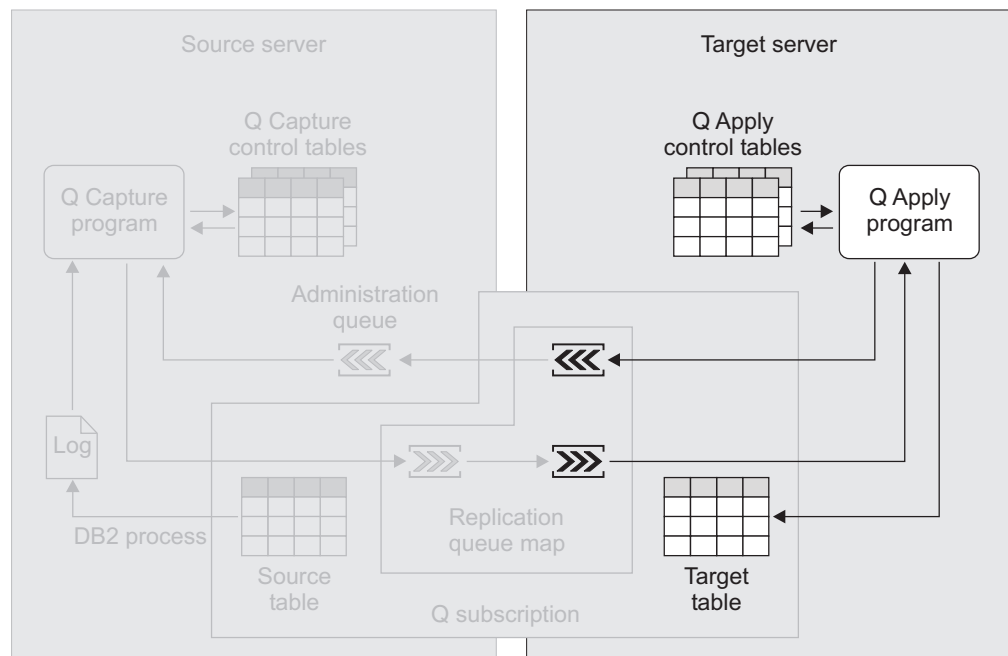
The Q Apply program can receive messages for a large number of Q subscriptions on a single receive queue and apply the transactions at very high rates of speed. In most cases, you can use a single replication queue map between one source server and one target server and not experience noticeable latency.

A single pair of Q Capture and Q Apply programs can be configured with multiple replication queue maps. If multiple applications are running at the source server that update independent sets of tables, then you should consider defining multiple

replication queue maps to allow for parallel delivery and application of data for each independent set of tables. A Q Apply program will create a multi-threaded process for each receive queue.

In most cases, one Q Apply program per server is sufficient, even for replication configurations with a high volume of transactions being replicated to a large number of tables. A single Q Apply program can be configured to process one or more receive queues. You can also run multiple Q Apply programs on a server if, for example, you require the data on one or more receive queues to be processed differently than the data on other receive queues.

The following figure shows the Q Apply program applying data in a simple configuration in Q replication.



*Figure 13. Applying data in Q replication.* The Q Apply program reads messages from a receive queue, converts them to SQL, and applies data to target tables. The Q Apply program continuously updates its control tables with information that you can use to monitor its progress.

**Related concepts:**

- “Q replication—Overview” on page 21

---

## Types of replication in Q replication

With Q replication, you can configure three different types of replication.

### Unidirectional replication

In unidirectional replication, changes that occur to a source are replicated to a target. Target tables are typically used exclusively by read-only applications. You can replicate all of the rows and columns from a source, or you can choose to replicate column and row subsets. No replication takes place from the target back to the source.

You can have one source replicating to one target, one source replicating to multiple targets, multiple sources replicating to one target, or multiple

sources replicating to multiple targets. For each source and its corresponding target, there is one Q subscription. Q subscriptions should use the same replication queue map if their targets are logically related; by grouping Q subscriptions this way, you ensure that the data applied to the targets is consistent with the original transactions on the source server.

At each source server, you create at least one set of Q Capture control tables. At each target server, you create at least one set of Q Apply control tables.

For Q subscriptions in unidirectional replication, you can have the Q Apply program call a stored procedure and pass the source data as input parameters to the stored procedure. Instead of the source columns mapping to target columns, the source columns map to parameters in the stored procedure. By mapping source columns directly to parameters in a stored procedure, you avoid the need to parse the incoming data and have a clean, simple programming model. The Q Apply program calls a stored procedure for each row operation instead of inserting the rows into a table. The stored procedure is responsible for getting the source data to its final destination.

### **Bidirectional replication**

In bidirectional replication, replication occurs between tables on two servers. Tables on one server that are involved in replication are identical in structure to the corresponding tables on the other server. Two corresponding tables have the same number of columns, the same column names, and compatible data types, though they can have different schemas and names. You cannot replicate subsets of rows. Changes made to a table on either server are replicated to the corresponding table on the other server.

Applications on one server can make changes to a table at the same time that applications on the other server make changes to the corresponding table. If conflicts occur in the data that is replicated between corresponding tables, you can choose which of the two tables takes precedence. Conflicts are found by comparing old values to current values. This method of detection might not find all conflicts in your data, but requires less overhead than other methods.

For every two corresponding tables, there are two Q subscriptions. For example, if you are replicating between Table\_One on server Red and Table\_Two on server Blue, there are these two Q subscriptions:

- One Q subscription replicating from Table\_One to Table\_Two
- One Q subscription replicating from Table\_Two to Table\_One

When you create the Q subscriptions, if you want an initial load performed automatically, you can choose which table contains the data that you want to start with. In this example, if you specify that Table\_One contains the data that you want to start with, Table\_Two will be loaded with the data from Table\_One.

A Q Capture program and a Q Apply program run at both servers. The Q Capture program and Q Apply program on each server have matching schemas. For example:

- If on one server the schema of the Q Capture program is Green, the schema of the Q Apply program on that server is also Green.

- If on the other server the schema of the Q Capture program is Yellow, the schema of the Q Apply program on that server is also Yellow.

### Peer-to-peer replication

In peer-to-peer replication, replication occurs between tables on two or more servers. Tables on one server that are involved in replication are identical in structure to the corresponding tables on the other servers. All corresponding tables have the same number of columns, the same column names, and compatible data types, though they can have different schemas and names. You cannot replicate subsets of rows. Changes made to a table on any server are replicated to the corresponding tables on the other servers.

Applications on one server can make changes to a table at the same time that applications on the other servers make changes to the corresponding tables. Conflicts are detected and resolved with the help of version columns and triggers that are added to the tables when you create the Q subscriptions. Convergence is possible with peer-to-peer replication, which means that if changes to replicated tables are stopped and all changes are propagated, corresponding tables will be identical.

For each pair of corresponding tables, there are two Q subscriptions. For example, if you are replicating between Table\_One on server Red, Table\_Two on server Blue, and Table\_Three on server Green, there are three pairs of corresponding tables:

- Table\_One and Table\_Two
- Table\_Two and Table\_Three
- Table\_One and Table\_Three

For each of these pairs, there are two Q subscriptions. For example, between Table\_One and Table\_Two, there are these two Q subscriptions:

- One Q subscription replicating from Table\_One to Table\_Two
- One Q subscription replicating from Table\_Two to Table\_One

A Q Capture program and a Q Apply program run at each server. The Q Capture program and Q Apply program on each server have matching schemas. For example, if you are replicating between two servers:

- If on one server the schema of the Q Capture program is Blue, the schema of the Q Apply program on the same server is also Blue.
- If on the other server the schema of the Q Capture program is Red, the schema of the Q Apply program on the same server is also Red.

### Related concepts:

- “WebSphere MQ objects required for unidirectional replication (remote)” in the *IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference*
- “WebSphere MQ objects required for bidirectional or peer-to-peer replication (two remote servers)” in the *IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference*
- “WebSphere MQ objects required for unidirectional replication on the same system” in the *IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference*
- “WebSphere MQ objects required for peer-to-peer replication (three or more remote servers)” in the *IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference*



---

## Chapter 4. Comparison of SQL replication and Q replication

---

### Comparison of SQL replication and Q replication—Overview

The following topics describe the similarities and differences between SQL replication and Q replication.

- Comparison of the infrastructure of SQL replication and Q replication
- Comparison of sources and targets in SQL replication and Q replication
- Comparison of data capturing and applying in SQL replication and Q replication
- Which replication solutions should you use in common scenarios?

**Related concepts:**

- “Comparison of the infrastructure of SQL replication and Q replication” on page 31
- “Comparison of data capturing and applying in SQL replication and Q replication” on page 33
- “Comparison of sources and targets in SQL replication and Q replication” on page 32
- “Which replication solutions should you use in common scenarios?” on page 33

---

### Comparison of the infrastructure of SQL replication and Q replication

Table 1 compares the infrastructure of SQL replication and Q replication.

*Table 1. Comparison of the infrastructure in SQL replication and Q replication*

Points of comparison	SQL replication	Q replication
<b>Location of the control tables for the capturing program</b>	You must create the control tables for a Capture program on the DB2® UDB server where the Capture program runs. In most cases, this server is the same DB2 UDB server where the sources associated with the program are located.	You must create the control tables for a Q Capture program on the DB2 UDB server where the Q Capture program runs. This server is the same DB2 UDB server where the source tables associated with that program are located.
<b>Number of capturing programs possible per DB2 UDB server</b>	Multiple, each with its own set of control tables and identified by the schema of those control tables.	Multiple, each with its own set of control tables and identified by the schema of those control tables.
<b>Location of the control tables for applying program</b>	You can create the control tables for the Apply program on any server that can connect to the source server and the target server. In general, the control tables should be located on the server where the Apply program runs.	You must create the control tables for a Q Apply program on the DB2 UDB server where the target tables associated with that program are located.
<b>Number of applying programs possible per DB2 UDB server</b>	Multiple, each sharing one set of control tables and identified by a string called an Apply qualifier.	Multiple, each with its own set of control tables and identified by the schema of those control tables.

**Related concepts:**

- “Comparison of SQL replication and Q replication—Overview” on page 31

## Comparison of sources and targets in SQL replication and Q replication

Table 2 compares how sources and targets are treated in SQL replication and Q replication.

Table 2. Comparison of sources and targets in SQL replication and Q replication

Points of comparison	SQL replication	Q replication
<b>Source and target platforms</b>	Sources and targets can be on the following relational database management systems: <ul style="list-style-type: none"> <li>• DB2<sup>®</sup> UDB for Linux, UNIX<sup>®</sup>, and Windows<sup>®</sup></li> <li>• DB2 UDB for z/OS<sup>™</sup></li> <li>• DB2 UDB for iSeries<sup>™</sup></li> <li>• Informix<sup>®</sup></li> <li>• Microsoft<sup>®</sup> SQL Server</li> <li>• Oracle</li> <li>• Sybase</li> <li>• Teradata (for targets only)</li> </ul>	Sources and targets can be on the following relational database management systems: <ul style="list-style-type: none"> <li>• DB2 UDB for Linux, UNIX, and Windows</li> <li>• DB2 UDB for z/OS</li> </ul>
<b>Database objects that can be targets</b>	DB2 UDB tables and views, and tables on non-DB2 relational databases.	DB2 UDB tables and stored procedures.
<b>Pairing of sources and targets</b>	You register a source. This information is stored in the Capture control tables. Then, you create one or more subscription-set members to map this registered source to targets. Information about these subscription-set members is stored in Apply control tables.	You create a Q subscription to map a source to a single target. No registration of the source is required. A source can be replicated to multiple targets by creating one Q subscription for each target. Information about the Q subscription is stored in Q Capture control tables and Q Apply control tables.
<b>Grouping of sources-target pairs</b>	You group source–target pairs into subscription sets. Each source–target pair is referred to as a subscription-set member.	You can group Q subscriptions by replication queue map. No subscription set object exists in Q replication.
<b>Subsetting of source columns and rows allowed?</b>	Yes.	Yes.
<b>Transforming data</b>	You can transform data by using stored procedures or SQL statements run by the Apply program. On the source server, the stored procedures or SQL statements can transform staged data. On the target server, they can transform data in targets.  The stored procedure interface limits transformations. The Apply program passes no parameters to a stored procedure and runs it once for an entire subscription set rather than for each subscription–set member.  You can also use triggers to transform data that the Capture program writes to staging tables.	You can transform data by using stored procedures as targets. The Q Apply program calls the stored procedures and passes the changed data to them as parameters. This method of using stored procedures allows considerable versatility for data transformation.

### Related concepts:



- “Comparison of SQL replication and Q replication—Overview” on page 31

## Comparison of data capturing and applying in SQL replication and Q replication

Table 3 compares data capturing and applying in SQL replication and Q replication.

*Table 3. Comparison of data capturing and applying in SQL replication and Q replication*

Points of comparison	SQL replication	Q replication
<b>Location of the capturing program</b>	The Capture program runs on the DB2 <sup>®</sup> UDB server where its control tables are located. In most cases, the source tables are also on the same server.	The Q Capture program runs on the DB2 UDB server where its control tables are located. The source tables are also on the same server.
<b>How DB2 changed data is captured</b>	The Capture program reads the DB2 recovery log and stores committed transactional data in staging tables.	The Q Capture program reads the DB2 recovery log and converts committed transactional data to messages.
<b>How DB2 UDB changed data is transported</b>	The Apply program fetches data from staging tables and applies the data to targets with DB2 SQL.	The Q Capture program puts data as messages in queues. WebSphere <sup>®</sup> MQ moves the messages to the target system. The Q Apply program gets these messages from queues and applies data to the targets using DB2 SQL.
<b>Location of the applying program</b>	Apply programs can run on any DB2 UDB server in your network, provided that they can connect to the servers that contain the source, target, and Apply control tables.	Q Apply programs run on target servers.
<b>How data is applied to DB2 UDB targets</b>	<p>The Apply program can process a subscription set in table mode or transaction mode.</p> <p>In table mode, the Apply program processes the fetched changes for each table separately, and then issues a single commit after all data is applied.</p> <p>In transaction mode, the Apply program applies the fetched changes to all of the target tables, in the order that the changes occurred in their corresponding transactions on the source server. The Apply program commits these transactions at transaction boundaries. You specify how many transactions to apply before each commit.</p>	The Q Apply program can apply transactions concurrently whenever they do not involve dependent changes to the target tables. The Q Apply program uses various methods of conflict handling and resolution to make bidirectional and multidirectional replication possible.

### Related concepts:

- “Comparison of SQL replication and Q replication—Overview” on page 31

## Which replication solutions should you use in common scenarios?

Table 4 on page 34 describes the most common scenarios that are possible with SQL replication and Q replication and recommends which replication technology to use for each scenario. The recommendations are general and might not apply for your particular needs.

Table 4. Recommended replication solutions for common replication scenarios

Configuration	Description of configuration	Recommended replication solution
<b>Data consolidation</b>	You can replicate data from many sources to a central repository. Data consolidation configurations are most often used to build data warehouses or operational data stores.	You can use either SQL replication or Q replication successfully in this scenario.
<b>Data distribution</b>	You can replicate data from a source to one or more targets that reside anywhere in a distributed network. Applications use the local target tables, so that they do not overload the network or central server. Each target server might only be interested in or allowed to access a subset of the data. You can choose to have a subset of columns and rows replicated so that each site sees only the data that it needs to see.	If you have a large number of targets, SQL replication might be the better choice because it allows you to stage data once and replicate that data to many targets.  Otherwise, you can use either SQL replication or Q replication successfully in this scenario.
<b>Update-anywhere replication</b>	You can replicate data between a master data source and one or more replicas of the source. Changes made to the master source are propagated to the replicas, and changes made to the replicas are also propagated to the master source. Whenever a conflict occurs between the data sent from the master source and data sent from a replica, the data from the master source takes precedence.	In a typical hub-and-spoke configuration with more than one replica, SQL replication is required. Such a configuration replicates from one master data source to multiple replicas.  For a scenario with a single master and a single replica, Q replication is recommended.
<b>Hot standby</b>	You can replicate data from a production server to a backup server. If the production server goes down, applications can instantly switch to using data on the backup server. When the production server is back online, committed changes made at the backup server can be replicated to the production server and applications can then use the data on the production server again.	Q replication is recommended. Either bidirectional or peer-to-peer replication can be used in this configuration, depending on the needs of your applications.
<b>Peer-to-peer replication</b>	You can replicate data between table copies on two or more servers. Committed changes made to one table copy are replicated to all other corresponding table copies. This type of replication can support geographically-distributed applications.  For example, an online business might have three servers on three continents. Each server contains the same product and ordering information. Customers typically access the product data and provide ordering information at the server on their home continents. However, if the server on one continent goes down, those customers can be rerouted to a server on another continent. Later, their ordering information can be replicated back to the server for their continent so that local financial and shipping offices have access to the information.	Q replication is recommended because of its strengths in conflict detection and resolution. Peer-to-peer replication is easy to set up in Q replication through the Replication Center.

Table 4. Recommended replication solutions for common replication scenarios (continued)

Configuration	Description of configuration	Recommended replication solution
<b>Query off-loading</b>	<p>When an OLTP server is overloaded, you can improve the performance of the OLTP application on that server by off-loading the query workload to another server.</p> <p>For example, a cable television company might have a large number of technicians handling customer tickets during each working day. The tables used to keep track of these tickets receive a large number of updates, inserts, and deletes per day. These tables are on a production system that handles other aspects of the business at the same time. The company could replicate relevant tables to another system and run a reporting tool on that system.</p>	You can use either SQL replication or Q replication successfully in this scenario.

**Related concepts:**

- “Comparison of SQL replication and Q replication—Overview” on page 31



---

## Chapter 5. Event publishing

---

### Event publishing—Overview

The following topics describe the concepts involved in event publishing. For more specific details, see the *Replication and Transaction Publishing Guide and Reference*.

- Infrastructure for an event publishing environment
- Sources in event publishing
- Capture of data in event publishing

The following figure shows a simple configuration in event publishing. The topics listed above each highlight and describe different sections of this figure.

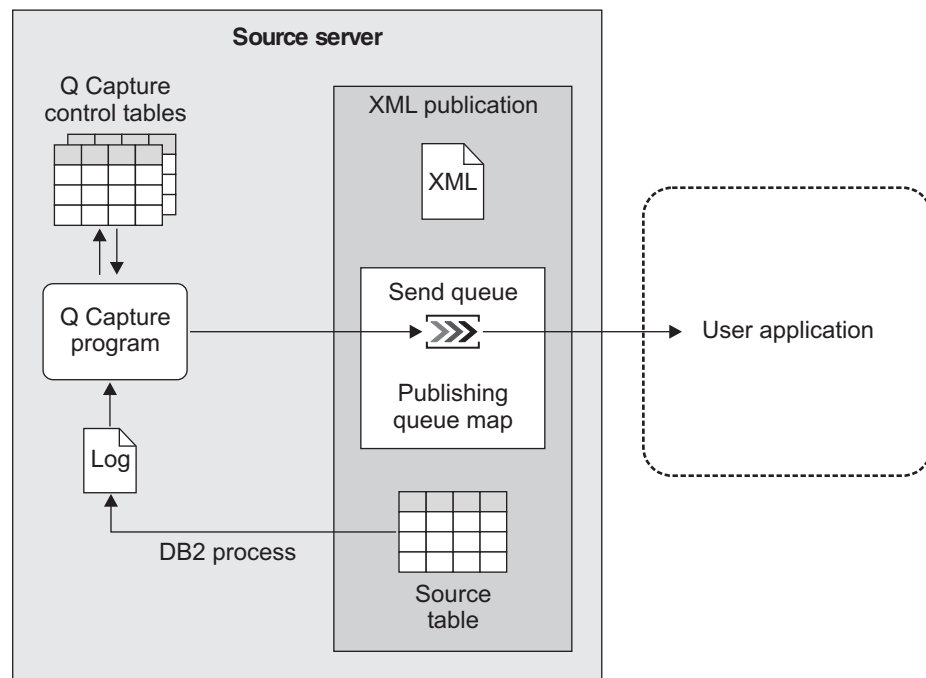


Figure 14. A simple configuration in event publishing

#### Related concepts:

- "Capture of data in event publishing" on page 41
- "Infrastructure for an event publishing environment" on page 38
- "Sources in event publishing" on page 39
- "WebSphere MQ objects required for event publishing" in the *IBM DB2 Information Integrator Replication and Event Publishing Guide and Reference*

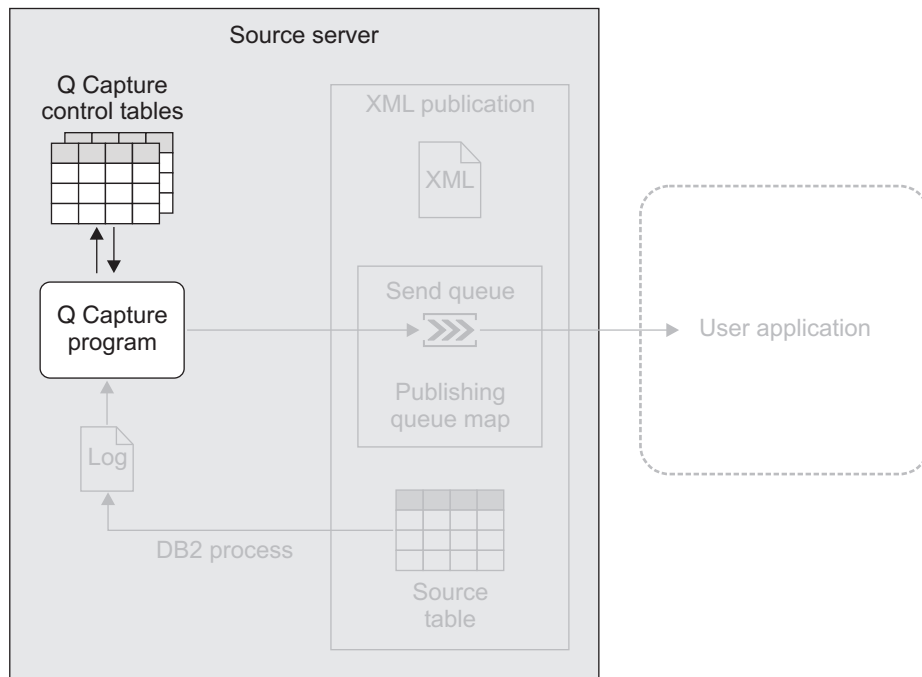
## Infrastructure for an event publishing environment

Event publishing allows you to publish committed transactional or row-level data from DB2® UDB tables as messages in WebSphere® MQ message queues. The messages can be read and interpreted directly by user applications, or they can first be interpreted by a message broker such as WebSphere Business Integration Message Broker or a DB2 MQ listener daemon.

Event publishing captures data by using a program called the *Q Capture program*. This is the same Q Capture program that is used in Q replication. In fact, it is possible for one Q Capture program to be involved in Q replication and event publishing at the same time. The Q Capture program uses a set of control tables, called *Q Capture control tables*, to store the information that the Q Capture program requires to perform its tasks (such as information about what its sources are and what to publish from its sources) and to store information that it generates itself (such as information about how well it is performing).

You can run multiple Q Capture programs on the same DB2 UDB server. Each Q Capture program uses its own set of control tables. The schema associated with a set of Q Capture control tables identifies the Q Capture program that uses those control tables. This schema is called a *Q Capture schema*.

The following figure shows the infrastructure in a simple configuration in event publishing.



*Figure 15. Infrastructure for a simple configuration in event publishing.* You create a set of DB2 relational tables on the source server called Q Capture control tables. Information about your sources is written to these tables. The Q Capture program, which runs on the source server, uses this information to know what data to capture and publish.

### Related concepts:

- “Event publishing—Overview” on page 37

---

## Sources in event publishing

With event publishing, you create objects called *XML publications* to define how changes from a single source table are published in XML format to a WebSphere® MQ message queue. You can then have an application retrieve and use those XML messages.

When you create an XML publication, you specify the following attributes:

### **The source server and source table to publish data from**

A source server is a DB2® UDB server that contains at least one set of Q Capture control tables. You can use a DB2 UDB relational table located on a source server as your source table.

### **The columns that contain data that you want to publish**

You can subset the columns that you publish from the source.

### **A predicate for publishing a subset of rows**

You can subset the rows that you publish from the source.

### **Which publishing queue map to use**

The Q Capture program writes data to a WebSphere MQ object that is referred to in event publishing as a *send queue*. The send queue passes the data to another queue that a user application can read the data from. Event publishing conveniently lets you group a send queue and other options for sending data messages into an object called a *publishing queue map*. For every XML publication, you select a publishing queue map to use. Multiple XML publications can use the same publishing queue map.

The publishing queue map that you select determines whether each message that the Q Capture program sends contains committed data for all source rows changed in a transaction or for only a single source row changed in a transaction.

For example, you might want each message to contain committed data for all source rows involved in a transaction if your application handles purchase orders. In this case, you would want a collection of data published together, such as product type, price, shipping address, and billing information, so that with one message you could start processes in systems for billing, shipping, and inventory tracking.

You might want each message to contain committed data for one source row involved in a transaction if you are providing streaming information to a Web application that provides stock quotes. In this case, each change to a stock price is unassociated with any other data.

### **Whether values in unchanged non-key columns should be published together with updates to other non-key columns**

By default, if there are updates to one or more of the columns that you selected, the updates are published but the values in the unchanged columns are not. For example, from table T1 you want to publish columns A1 (the primary key), A2, and A3. If a transaction updates a value in column A2 and then commits, the Q Capture program will publish a message that contains only the new value in column A2, as well as the value in the key column A1.

You can choose to have the messages contain updated and unchanged values in non-key columns. If column A2 is updated, the Q Capture

program will publish a message that contains the new value in column A2 and the unchanged value in column A3, in addition to the value in the key column A1.

You might want to choose this option if it is easier to write your application to always expect a value for each column. You also might want to choose this option if you are writing an application to audit changes to your table and that application needs a complete snapshot of each row.

**Whether to include new and old values for the data in updated non-key columns, or only the new values**

When there are updates to one or more of the columns that you selected, the message containing those updates will provide only the new values by default. For example, from table T2 you want to publish columns B1 (the primary key) and B2. If a transaction updates column B2 from 5 to 6, the message that contains that update will provide only the value 6.

You can choose to have the messages provide the old values together with the new ones. In this case, the message that contains the update to column B2 would provide both the value 5 and the value 6.

Choose to send old values if the application that receives the published changes requires both old and new values. For example, your application might be aggregating information and finding the difference between the old and new values. Also, if you are publishing significant price changes to a Web application and that application requires the old prices, sending the old prices along with the new prices saves the application from having to look up the old prices elsewhere.

**Important:** XML publications are not part of Q subscriptions. Q subscriptions are used to replicate information only to a Q Apply program. XML publications are used only to publish data in XML format to user applications.

The following figure shows an XML publication in a simple configuration in event publishing.



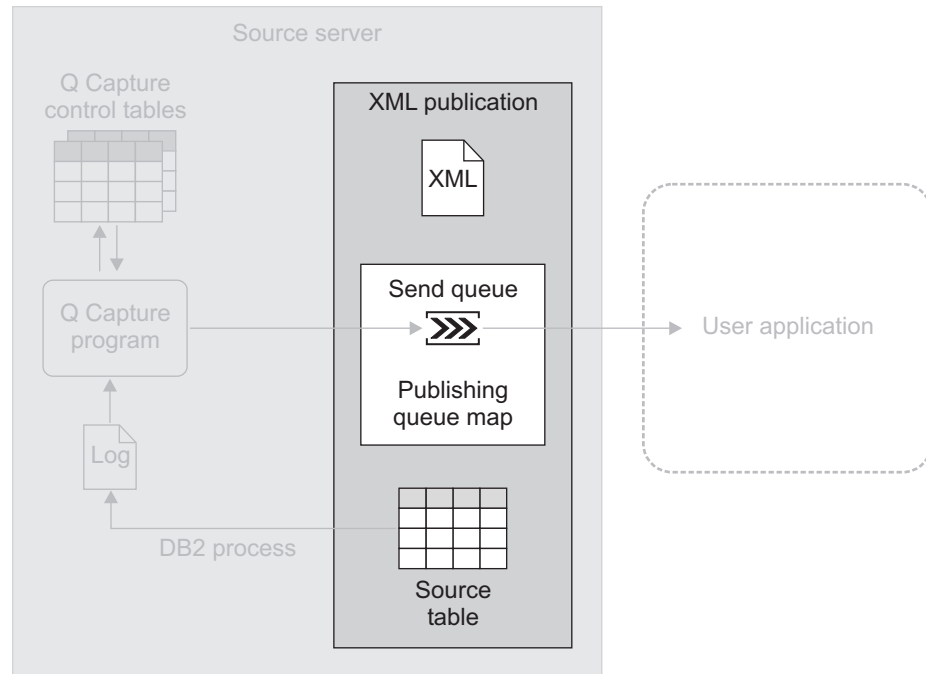


Figure 16. An XML publication. An XML publication tells the Q Capture program which changes to capture for a source table. It also tells the Q Capture program which send queue to use for transporting XML messages.

**Related concepts:**

- “Event publishing—Overview” on page 37

---

## Capture of data in event publishing

The Q Capture program reads the DB2® log sequentially for changes to the source tables. If it reads a change to one of your source tables, the Q Capture program adds the change to the corresponding database transaction that it is retaining in memory. Transactions in memory are therefore potentially subsets of the corresponding transactions in the log; they contain only changes to your source tables. When the Q Capture program reads the COMMIT statement for a transaction, it converts the transaction into one or more messages, depending on whether you want to send row-level messages or transaction-level messages and depending on how many send queues are being used by the source tables that are involved in the transaction. The Q Capture program then puts the message on the corresponding send queues.

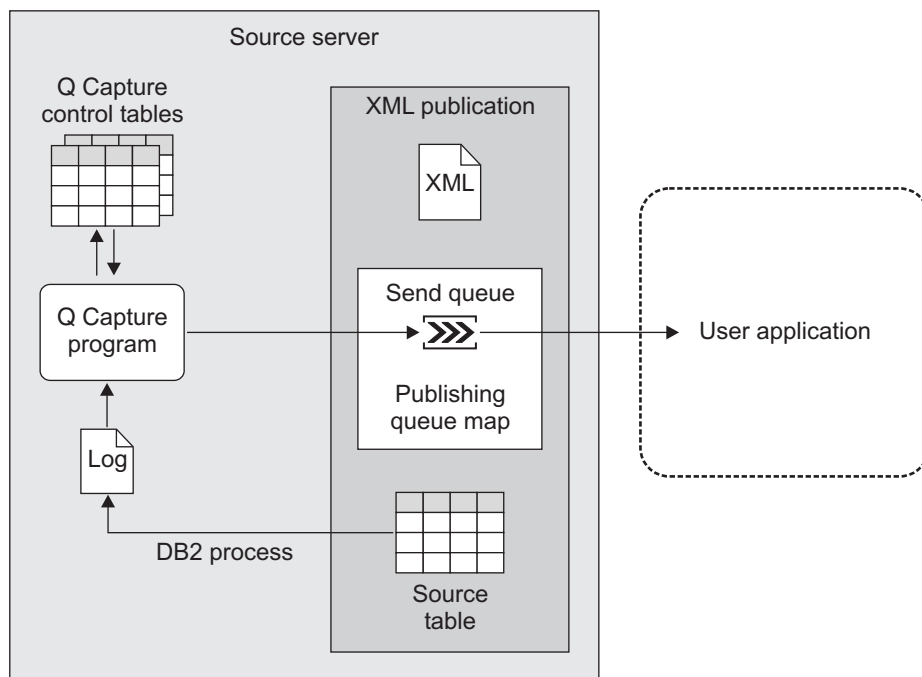
For example, you create three XML publications: PUB1, PUB2, and PUB3. PUB1 and PUB2 use the same publishing queue map (which uses SENDQ1), and you want to use these XML publications to publish messages that contain committed row-level changes. For PUB3, you choose a different publishing queue map (which uses SENDQ2), and you want to use this XML publication to publish messages that contain committed transaction-level changes. A publishing queue map can send either row-level messages or transaction-level messages, but it cannot send both types of messages.

The Q Capture program reads a COMMIT for a database transaction that involves changes to the source tables in all three XML publications. The Q Capture program converts the parts of the transaction that involve the source tables in PUB1 and

PUB2 to row-level messages that it writes to SENDQ1. The Q Capture program also converts the parts of the transaction that involve the source table in PUB3 to a transaction-level message that it writes to SENDQ2.

You can run more than one Q Capture program on the same source server. Although one Q Capture program can capture changes made to many sources and send those changes to many user applications, in some situations you might benefit from running more than one Q Capture program. For example, multiple Q Capture programs, which could also be helpful in large sysplexes, can improve performance and achieve higher throughput. The trade-off is additional CPU overhead associated with multiple log readers. Using multiple Q Capture programs also requires more DB2 UDB connections.

The following figure shows the Q Capture program capturing data in a simple configuration in event publishing.



*Figure 17. Capturing data in event publishing.* The Q Capture program reads the DB2 recovery log for changes made to sources and converts committed transactional data to XML messages, which the Q Capture program puts on a send queue. The send queue transports data to any number of possible applications that can read XML data. The Q Capture program continuously updates its control tables with information that you can use to monitor its progress.

**Related concepts:**

- “Event publishing—Overview” on page 37

## Chapter 6. Comparison of SQL replication, Q replication, and event publishing

Table 5 compares the main features of SQL replication, Q replication, and event publishing.

Table 5. Comparison of the main features in SQL replication, Q replication, and event publishing

Points of comparison	SQL replication	Q replication	Event publishing
<b>Uses</b>	Multiple, including capacity relief, feeding data warehouses and data marts, auditing change history.	Multiple, including failover, capacity relief, supporting geographically distributed applications, data availability for planned or rolling upgrades or outages.	Multiple, including feeding central information brokers and Web applications, and triggering actions based on updates, inserts, or deletes to source tables.
<b>How data is replicated or published</b>	Committed transactional data is captured and stored in staging tables by a capturing program. An applying program reads the information from the staging tables and applies the transactions to target tables.	Committed transactional data is captured and put on WebSphere® MQ message queues by a capturing program. An applying program reads the information from message queues and applies the transactions to target tables.	Committed transactional data is captured and put on WebSphere MQ message queues by a capturing program.
<b>Sources</b>	Tables (either in DB2® UDB or other supported relational database management systems), views	Tables in DB2 UDB	Tables in DB2 UDB
<b>Targets</b>	Tables (either in DB2 UDB or other supported relational database management systems), views	Tables in DB2 UDB	User application
<b>Filtering possible?</b>	Yes. Column and row filtering.	Yes. Column and row filtering.	Yes. Column and row filtering.
<b>Where the filtering is done?</b>	Source, target, or both.	Source	Source
<b>Data transformation possible?</b>	Yes. Data cleansing, data aggregation, and calculated columns in target tables. Manipulate data once and replicate to many targets. Manipulate data and replicate to selected targets.	Yes, by using user-developed stored procedures.	Yes. Any required data transformation is performed by the user application.
<b>Utilities supported for automatic loads of target tables</b>	EXPORT/IMPORT, EXPORT/LOAD, LOAD FROM CURSOR option of the LOAD utility, depending on the target platform.	EXPORT/IMPORT, EXPORT/LOAD, LOAD FROM CURSOR option of the LOAD utility, depending on the target platform.	No automatic loads because there are no target tables.
<b>Supported operating systems</b>	Linux, UNIX®, Windows®, z/OS™, iSeries™	Linux, UNIX, Windows, z/OS	Linux, UNIX, Windows, z/OS

Table 5. Comparison of the main features in SQL replication, Q replication, and event publishing (continued)

Points of comparison	SQL replication	Q replication	Event publishing
Supported relational database management systems	DB2 UDB, Oracle, Sybase, Teradata, Microsoft® SQL Server	DB2 UDB	DB2 UDB
Administrative interfaces	Replication Center or command line	Replication Center or command line	Replication Center or command line

**Related concepts:**

- “SQL replication—Overview” on page 7
- “Q replication—Overview” on page 21
- “Event publishing—Overview” on page 37

## Chapter 7. Comparison of Q replication to high availability disaster recovery (HADR)

DB2® high availability disaster recovery (HADR) provides a high-availability technology to help you recover from complete site failures, as well as to support applications that demand ultrafast failover for partial site failures. HADR is also useful for rolling upgrades and other kinds of planned outages. HADR ships database log records for an entire database from a source copy of the database (called the primary) to a target copy of the database (called the standby). The standby database cannot be accessed by applications. The standby is initialized using a restore or a split mirror, which is an identical and independent copy of disk volumes that can be attached to a different system and can be used in various ways.

When you start HADR, it retrieves log records and replays those records on the standby until the standby catches up to the in-memory log set of the primary. The primary then writes log data to local disk and sends them to the standby. Updates to the standby database occur by rolling forward log data that is generated on the primary database and shipped to the standby database.

You can also use Q replication to protect your data from complete site failures and to support application failover for partial site failures, as well as for rolling upgrades and other planned outages. Table 6 summarizes the differences between HADR and Q replication when used for the same purposes as HADR.

Table 6. Comparison of HADR with Q replication

Points of comparison	High availability disaster recovery	Q replication
What is the scope of the setup procedure?	Entire DB2 UDB database.	Tables within a DB2 UDB database or subsystem.
How is data propagated to the standby?	Log operations are shipped to the standby database and replayed continuously through forward recovery.	Committed transactional data are captured from the DB2 recovery log and applied to target tables.
Is synchronous replication possible?	Yes.	No.
Is asynchronous replication possible?	Yes.	Yes.
Are client applications automatically rerouted to the standby?	Yes.	Yes.
Which operating systems are supported?	Linux, UNIX®, and Windows®.	Linux, UNIX, Windows, z/OS™.
Can applications read data from the standby?	No.	Yes.
Can applications write data to the standby?	No.	Yes.
Is SQL Data Definition Language (DDL) automatically replicated in addition to SQL Data Manipulation Language (DML)?	Yes.	No.

Table 6. Comparison of HADR with Q replication (continued)

Points of comparison	High availability disaster recovery	Q replication
On what hardware, operating system, and version of DB2 UDB can the standby run?	Hardware, operating system, and version of DB2 UDB must be identical to those where the source database is located.	Hardware, operating system, and version of DB2 UDB can be different from those where the source database is located.
Are there tools for monitoring performance?	Yes.	Yes.
Is network compression and encryption built in?	No.	Yes.
Can the source and standby databases be partitioned?	No.	Yes.

For more information about HADR, see the *Data Recovery and High Availability Guide and Reference*.

---

## Chapter 8. Comparison of event publishing to DB2 UDB MQ user-defined functions

You can use WebSphere® MQ functions that are available in DB2® Universal Database as user-defined functions in scenarios similar to those scenarios that you can set up with event publishing. These functions allow users to use SQL to access WebSphere MQ message queues from DB2 UDB. These functions can be used by applications written in any of the programming languages supported by DB2 UDB in a wide variety of scenarios, including:

### **Data collection**

The application receives information in the form of messages from one or more sources. An information source can be any application. The application receives the data from queues and stores the data in database tables for additional processing.

### **Workload distribution**

The application posts work requests to a queue that is shared by multiple instances of the same application. When an application instance is ready to perform some work, it receives a message that contains a work request from the head of the queue. Multiple instances of the application can share the workload that is represented by a single queue of pooled requests.

### **Application signaling**

In a situation where several processes collaborate, you can use messages to coordinate their efforts. These messages might contain commands or requests to perform work.

Event publishing differs from the DB2 WebSphere MQ functions in that you can publish events as messages without having to modify the applications that generate these events. If you want to use the DB2 WebSphere MQ functions, you have to code your applications to work with them.

For more information about the DB2 WebSphere MQ functions, see the *IBM DB2 Information Integrator Application Developer's Guide*.

### **Related concepts:**

- “How to use WebSphere MQ functions within DB2” in the *IBM DB2 Information Integrator Application Developer's Guide*





---

## Accessibility

Users with physical disabilities, such as restricted mobility or limited vision, can use software products successfully by using accessibility features. These are the major accessibility features in DB2 Information Integrator Version 8:

- You can operate all features by using the keyboard instead of the mouse.
- You can customize the size and color of your fonts.
- You can receive either visual or audio alert cues.
- DB2 supports accessibility applications that use the Java™ Accessibility API.
- DB2 documentation is provided in an accessible format.

---

## Keyboard input and navigation

You can operate the DB2 database tools, such as Control Center, Data Warehouse Center, and Replication Center, by using only the keyboard. You can use keys or key combinations instead of a mouse to perform most operations.

In UNIX-based systems, the position of the keyboard focus is highlighted. This highlighting indicates which area of the window is active and where your keystrokes will have an effect.

---

## Accessible display

The DB2 database tools have features that enhance the user interface and improve accessibility for users with low vision. These accessibility enhancements include support for customizable font properties.

### Font settings

For the DB2 database tools, you can use the Tools Settings notebook to select the color, size, and font for the text in menus and windows.

### Nondependence on color

You do not need to distinguish between colors to use any of the functions in this product.

---

## Alternative alert cues

You can specify whether you want to receive alerts through audio or visual cues, using the Tools Settings notebook.

---

## Compatibility with assistive technologies

The DB2 Information Integrator graphical interface supports the Java Accessibility API, enabling the use of screen readers and other assistive technologies that are used by people with disabilities.

---

## Accessible documentation

Documentation for the DB2 family of products is available in HTML format. You can view documentation according to the display preferences set in your browser. You can use screen readers and other assistive technologies.

---

## 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 all 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 give 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/region or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

**The following paragraph does not apply to the United Kingdom or any other country/region 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites 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 that 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.

This information may contain 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 may contain 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.

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.

---

## Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM  
DB2  
Informix  
iSeries  
UNIX  
WebSphere  
Windows  
z/OS

The following terms are trademarks or registered trademarks of other companies:

Microsoft and Windows 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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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



---

# Index

## A

- Apply control tables
  - definition 8
  - location 31
- Apply program
  - definition 8
  - location 33
- Apply qualifier
  - definition 8
- applying data
  - Q replication 26
  - SQL replication
    - DB2 UDB targets 17
    - non-DB2 targets 19
- ASNCLP program 4

## B

- bidirectional replication 27

## C

- Capture control tables
  - definition 8
  - location 31
- Capture program
  - definition 8
  - location 33
- Capture schema
  - definition 8
- Capture triggers 18
- capturing data
  - event publishing 41
  - Q replication 25
  - SQL replication
    - DB2 UDB sources 15
    - non-DB2 sources 18
- CCD tables 18, 19
- CCD targets 11
- CD tables 9, 17, 18
- change-capture replication 9
- comparison of replication and publishing solutions 43

## D

- DB2 UDB MQ user-defined functions 47

## E

- event publishing
  - capturing data 41
  - comparison to DB2 UDB MQ user-defined functions 47
  - infrastructure 38
  - introduction 3
  - sources 39
- event timing 11

## F

- full-refresh replication 9

## H

- high availability disaster recovery (HADR)
  - comparison to Q replication 45

## I

- infrastructure
  - event publishing 38
  - Q replication 22
  - SQL replication 8
- interval timing 11

## L

- loading target tables
  - Q replication 23
  - SQL replication 17

## M

- mapping
  - sources to targets
    - Q replication 23
    - SQL replication 11

## P

- peer-to-peer replication 27
- point-in-time targets 11
- publishing queue map 39

## Q

- Q Apply control tables
  - definition 22
  - location 31
- Q Apply program
  - definition 22
  - location 33
- Q Apply schema
  - definition 22
- Q Capture control tables
  - definition
    - event publishing 38
    - Q replication 22
  - location
    - event publishing 38
    - Q replication 31
- Q Capture program
  - definition
    - event publishing 38
    - Q replication 22

- Q Capture program (*continued*)
  - location
    - event publishing 38
    - Q replication 33

- Q Capture schema
  - definition
    - event publishing 38
    - Q replication 22

- Q replication
  - applying data 26
  - capturing data 25
  - infrastructure 22
  - introduction 2
  - source-target pairs 23

- Q subscriptions
  - definition 23

## R

- registering sources in SQL replication 9
- replica targets 11
- replication
  - bidirectional 27
  - change-capture 9
  - full-refresh 9
  - peer-to-peer 27
  - unidirectional 27
- Replication Alert Monitor 4
- replication and publishing solutions
  - comparison 43
- Replication Center 4
- replication queue map 23

## S

- source tables
  - event publishing 39
  - Q replication 25
  - SQL replication 9
- source-target pairs
  - Q replication 23
  - SQL replication 11
- SQL replication
  - applying data
    - DB2 UDB targets 17
    - non-DB2 targets 19
  - capturing data
    - DB2 UDB sources 15
    - non-DB2 sources 18
  - change-capture replication 9
  - full-refresh replication 9
  - infrastructure 8
  - introduction 1
  - registering sources 9
  - source-target pairs 11
- SQL replication and Q replication
  - comparison
    - applying data 33
    - capturing data 33
    - infrastructure 31

SQL replication and Q replication  
(*continued*)  
    comparison (*continued*)  
        sources and targets 32  
subscription sets  
    definition 11

## **T**

target tables  
    Q replication 23  
    SQL replication 11  
tdiff 4  
trepair 4  
triggers  
    Capture 18

## **U**

unidirectional replication 27  
user copy targets 11

## **X**

XML publications 39



---

## Contacting IBM

To contact IBM in the United States or Canada, call one of the following numbers:

- For customer service: 1-800-IBM-SERV (1-800-426-7378)
- For DB2 marketing and sales: 1-800-IBM-4YOU (1-800-426-4968)

To learn about available service options, call one of the following numbers:

- In the United States: 1-888-426-4343
- In Canada: 1-800-465-9600

To locate an IBM office in your country or region, see the IBM Directory of Worldwide Contacts on the Web at [www.ibm.com/planetwide](http://www.ibm.com/planetwide).

---

## Product information

Information about DB2 Information Integrator is available by telephone or on the Web.

If you live in the United States, you can call one of the following numbers:

- To order products or to obtain general information: 1-800-IBM-CALL (1-800-426-2255)
- To order publications: 1-800-879-2755

On the Web, go to

<http://www.ibm.com/software/data/integration/db2ii/support.html>. This site contains the latest information on the technical library, ordering books, client downloads, newsgroups, fix packs, news, and links to Web resources.

To locate an IBM office in your country or region, see the IBM Directory of Worldwide Contacts on the Web at [www.ibm.com/planetwide](http://www.ibm.com/planetwide).

---

## Comments on the documentation

Your feedback helps IBM to provide quality information. Please send any comments that you have about this book or other DB2 Information Integrator documentation. You can use any of the following methods to provide comments:

- Send your comments using the online readers' comment form at [www.ibm.com/software/data/rcf](http://www.ibm.com/software/data/rcf).
- Send your comments by electronic mail (e-mail) to [comments@us.ibm.com](mailto:comments@us.ibm.com). Be sure to include the name of the product, the version number of the product, and the name and part number of the book (if applicable). If you are commenting on specific text, please include the location of the text (for example, a title, a table number, or a page number).







Printed in USA

GC18-7567-00

