

DB2 Replication & Publishing Overview

Introduction	1-2
Replication options	1-3
Replication uses	1-4
Data distribution	1-5
Data consolidation	1-6
Bi-directional replication	1-7
SQL replication components	1-9
Administration	1-10
Capture	1-11
Apply	1-15
Alert Monitor	1-18

Introduction

Replication involves the copying of data from one location to another.

Data is extracted by programs at a source location, transported to, and then loaded at the target location.

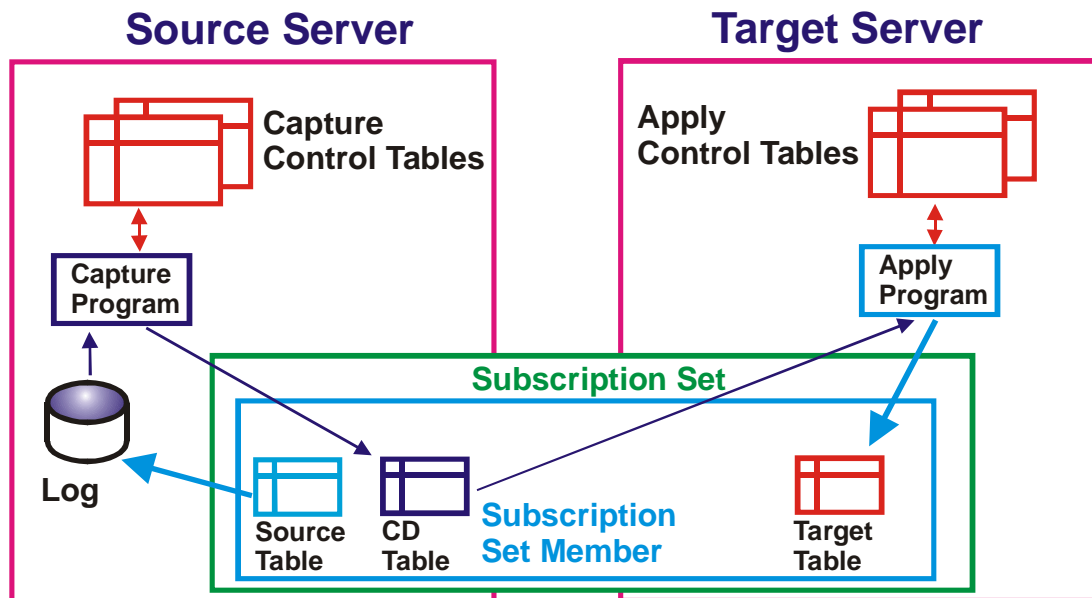
To improve efficiency, a common approach is to capture and transmit only the changes to the source data.

With DB2 Replication, the data can be filtered and transformed during the replication process. This means that different targets can receive data in different formats.

There are features which enable you to apply time constraints. This means that data can be delivered to different targets at different intervals.

DB2 Replication also provides management and monitoring facilities.

The replication process is designed to have a minimal impact on production systems by making use of the DB2 log, rather than accessing the source tables directly.



SQL Replication

DRA-10

Replication options

DB2 UDB Version 8.2 permits three types of replication:

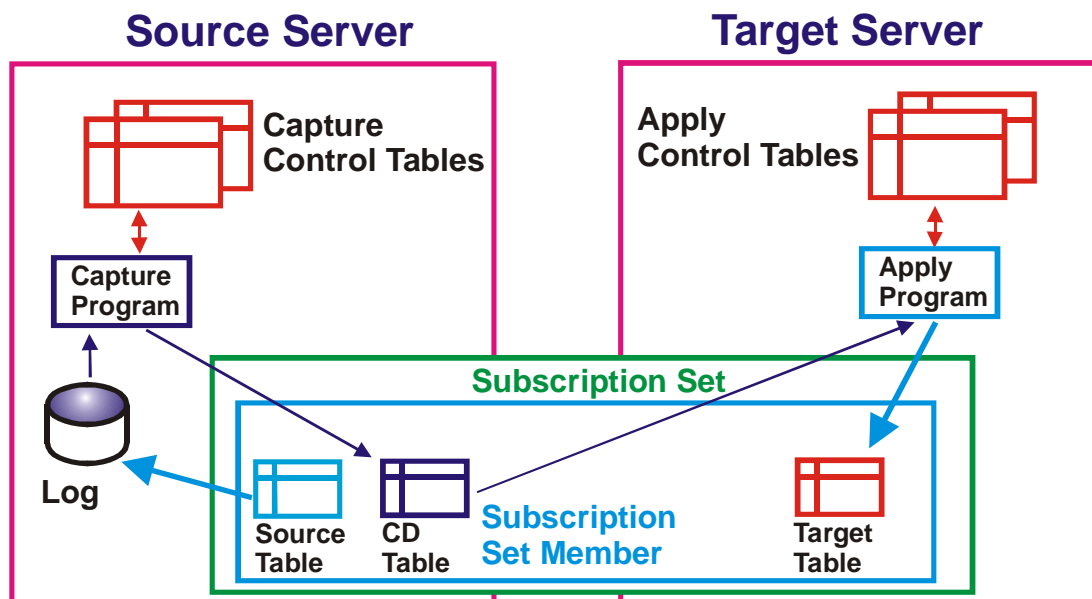
- SQL Replication
- Q Replication
- Event Publishing

SQL Replication captures changes in Change Data (CD) tables which are transmitted to the target locations.

Q Replication captures committed changes to transactional data in messages which are transmitted using WebSphere MQ queues.

Event Publishing extends Q Replication by placing the changes in an XML format. The messages can then be processed by products such as message brokers (WBIMB).

This course concentrates on SQL Replication.



SQL Replication

DRA-10

DB2 Versions

DB2 SQL Replication enables change-based replication for products in the IBM database family, including:

- DB2 Universal Database for z/OS and OS/390 Versions 6, 7, and 8
DB2 DataPropagator for z/OS and OS/390 V8
- DB2 Universal Database for iSeries on OS/400 V5R2
DataPropagator for iSeries V8
- DB2 Universal Database Version 8 for Windows, UNIX, and Linux
Included with the DB2 product, including 64-bit versions, Linux OS/390, and Linux OS/400
- Informix Dynamic Server Versions 7, 8, and 9
DB2 Universal Database Enterprise Server Edition Version 8 or DB2 Connect Version 8 is required.

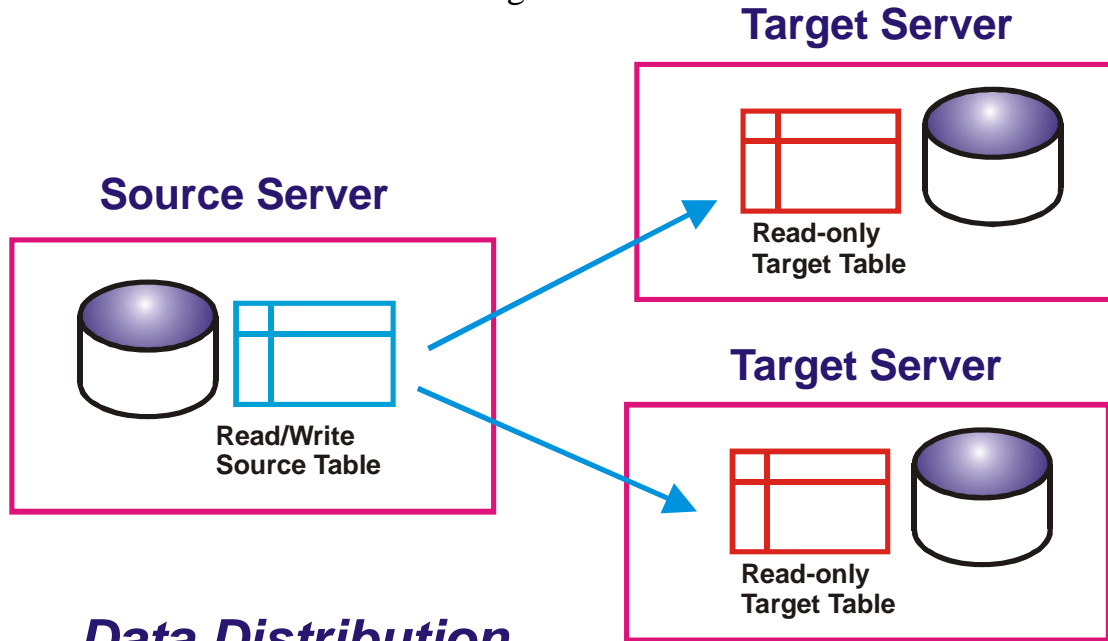
Replication Uses

There are many business reasons for using replication. These include:

- Distribution of data to other locations
- Consolidation of data from other locations
- Bidirectional exchange of data with other locations
- A combination of the above

Data Distribution

Data Distribution involves moving all or a subset of the data to one or more locations.



DRA-70

Data can be copied to a central data warehouse or decision support system. This often involves data transformation and denormalization.

You can copy subsets of the data to data marts to provide groups of users with local access. This allows you to use business intelligence tools with enterprise data, while maintaining the security and performance of production applications.

Data Distribution is also used to provide data to applications in the same or different environments. A simple example is maintaining a copy of the production data on another similar system. Other examples may need complex data transformation to fit new application requirements. The new application may be a Web application, a purchased package, or an application distributed on multiple laptop computers. Any data that is copied may have to be filtered and/or transformed for the target application.

Data Distribution can also be used to provide application co-existence if you are migrating from one environment to another. Legacy data can be copied to the new environment for reference by the new applications until such time as the legacy applications are migrated to the new environment.

Data Consolidation

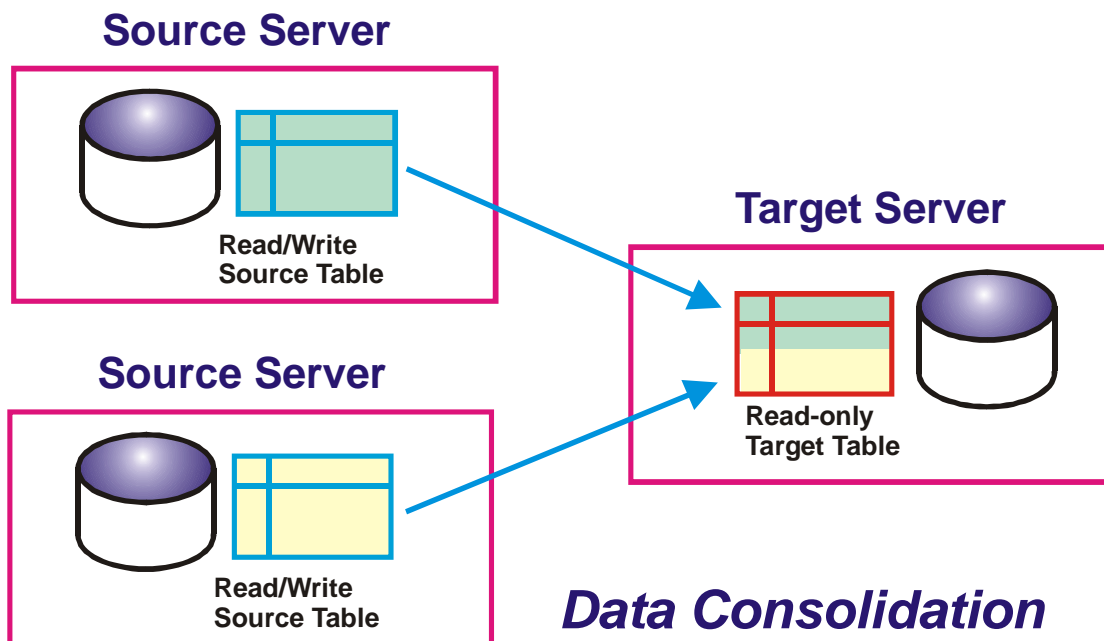
Your organization may have data on many different distributed systems.

For example, retail companies have data at each store. Manufacturing companies have data at each plant. Insurance companies have data at each branch office and perhaps on each laptop computer in the sales department.

Replication can be used to copy changes from each of the distributed sites to a central site. You would then be able to use the data for centralized analysis, reporting, and application processing.

Data Consolidation can be very useful for business intelligence applications such as OLAP or Data Mining.

The data at the source servers has the same structure. This means that the target is a UNION of the sources.



DRA-80

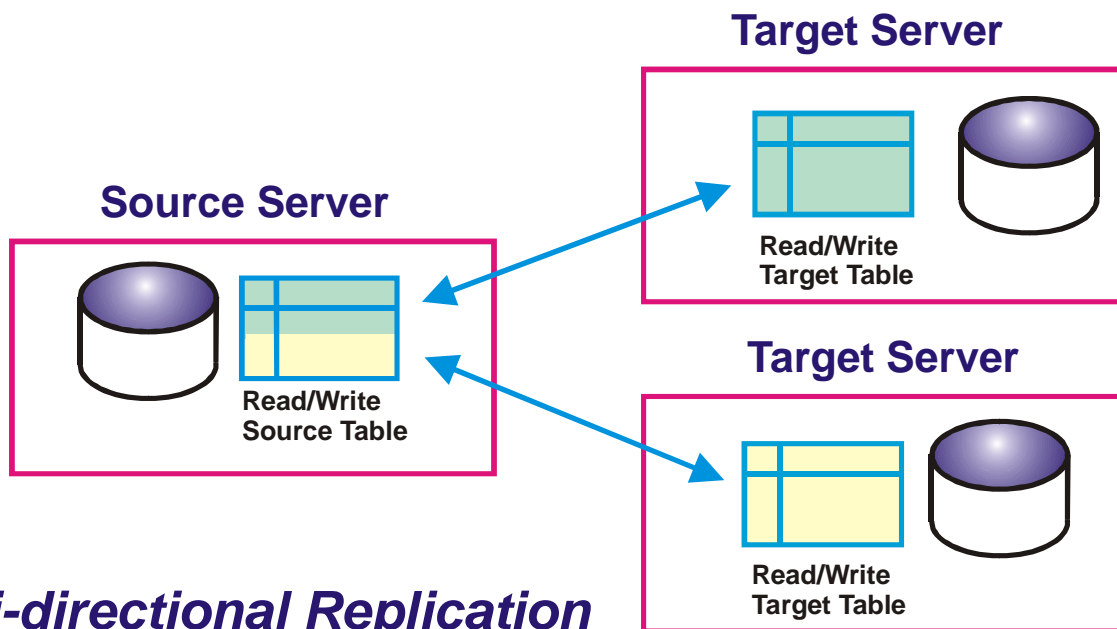
Bidirectional Replication

This method is used if the data can be updated at multiple locations.

The replication process must coordinate the changes made at any or all of the sites.

One location acts as the master location and distributes changes to the target locations.

Changes made at the targets flow to other target sites through the master.



Bi-directional Replication

DRA-90

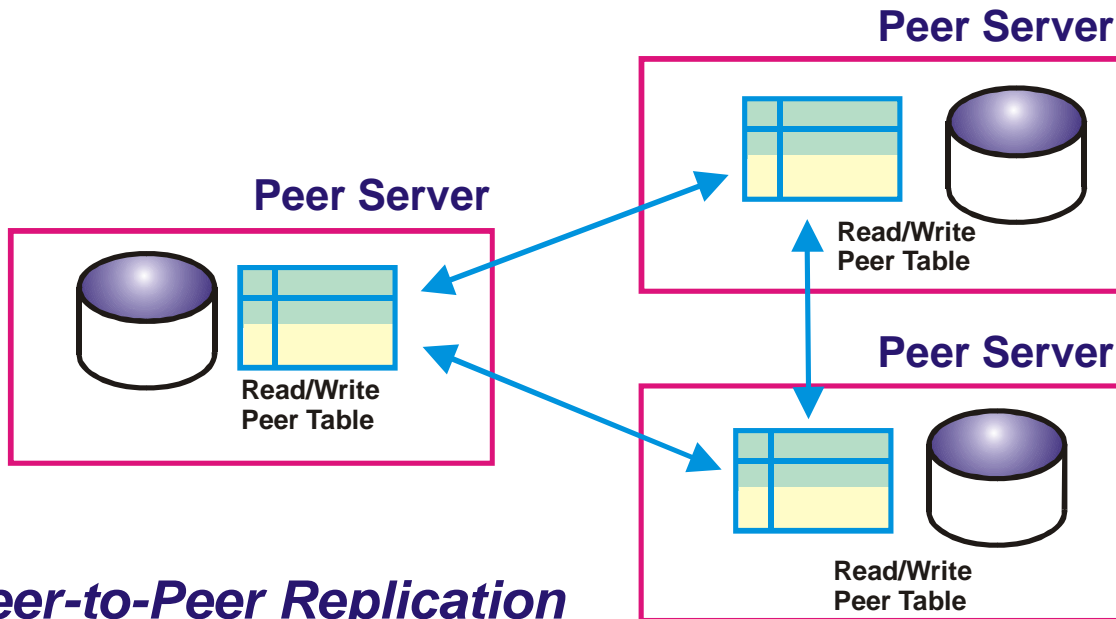
Bidirectional replication can be used for mobile applications where the target may be a computer in a branch office or in a delivery truck.

There may be many targets which are only occasionally connected to the source system. Often, the connection may be via phone lines, so efficiency is important.

This technique is sometimes known as 'master-slave' or 'update anywhere' replication.

Peer-to-Peer Replication

Peer-to-Peer Replication is a type of bidirectional replication which does not have a designated master location.



Peer-to-Peer Replication

DRA-100

Each Peer Server copies changes from all other locations directly.

This is also known as 'multi-master' replication.

Peer-to-peer replication uses include:

- maintaining disaster recovery sites
- providing fail-over systems for high availability
- balancing query workload across multiple locations

SQL Replication Components

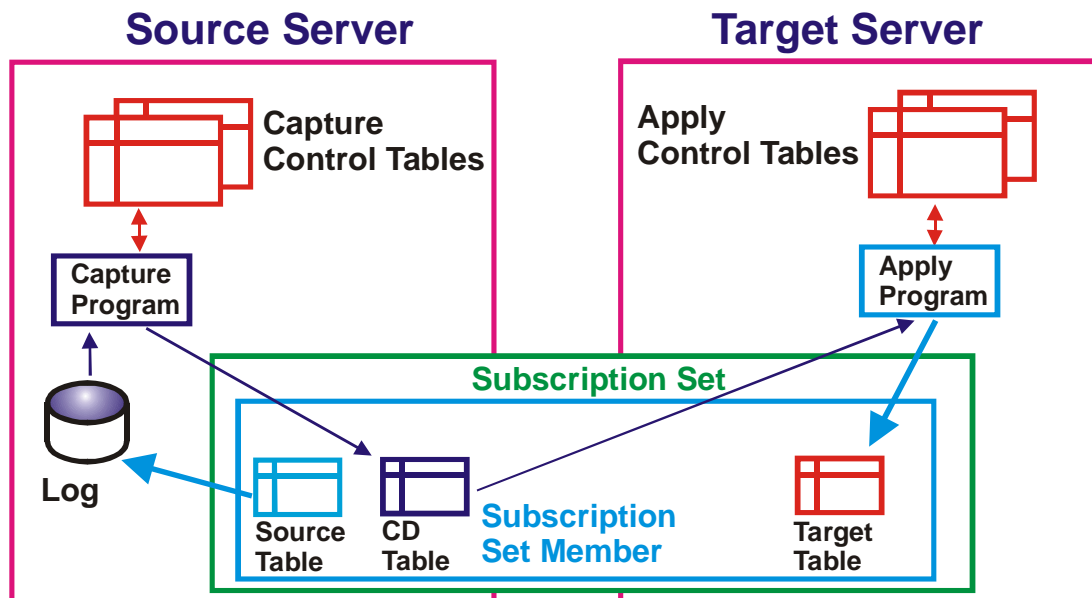
SQL replication uses four components:

- Administration
- Capture
- Apply
- Alert Monitor

The components communicate via relational tables, called control tables.

The control tables are created and populated using the Replication Center.

The Capture, Apply, and Alert Monitor programs read and update information in the control tables.



SQL Replication

DRA-10

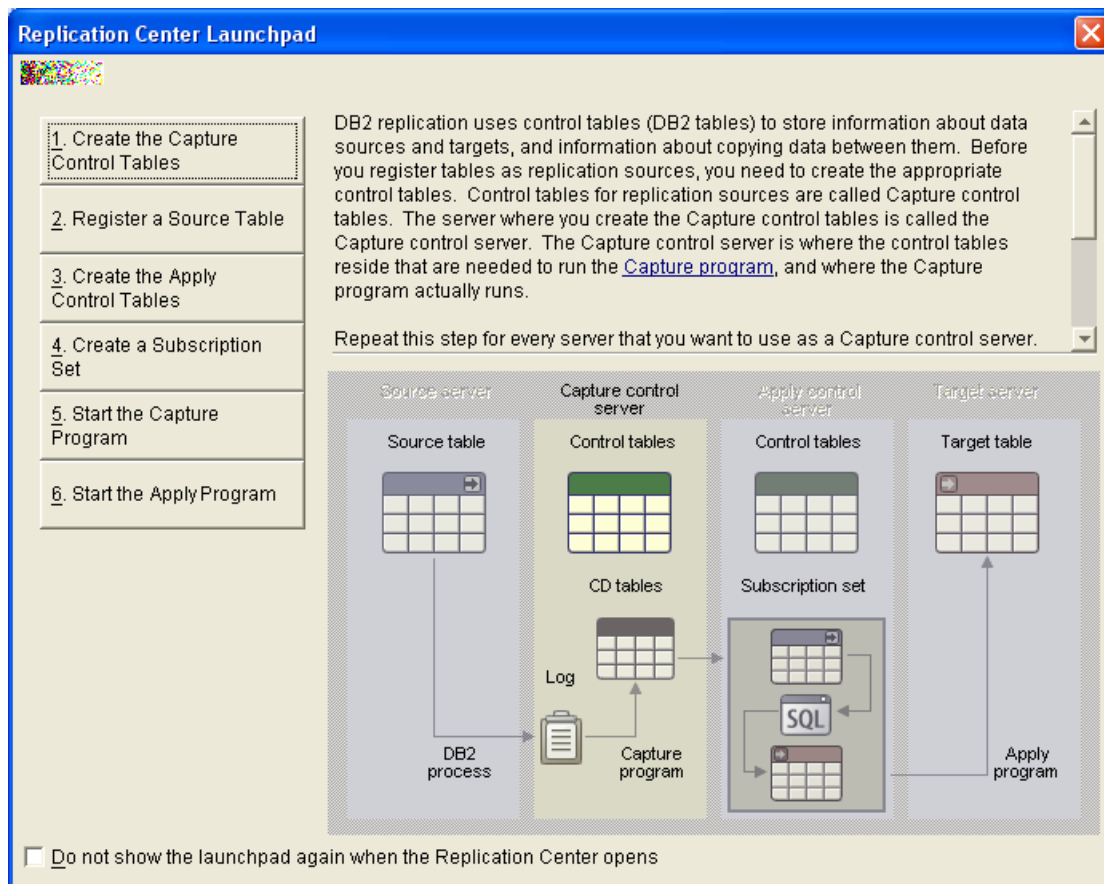
Administration

Administration is performed using the Replication Center. This is a graphical user interface used to:

- define replication sources
- map sources to targets
- manage and monitor the Capture and Apply processes on local and remote systems

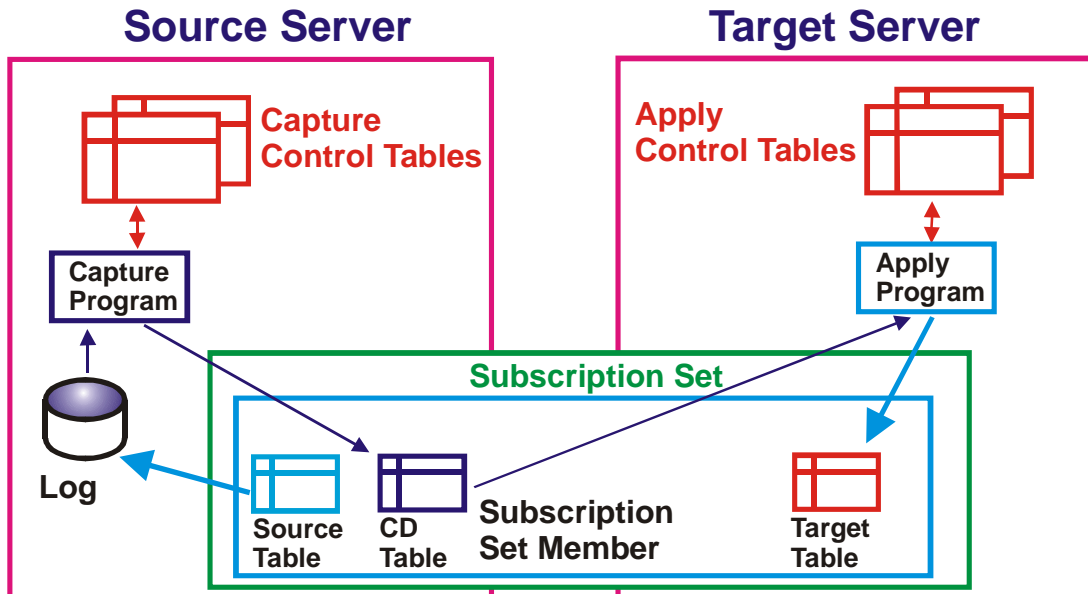
The Replication Center runs on Windows and UNIX/Linux systems and needs to have connectivity to both the source and target servers.

The DB2 V8 Administration Client for Windows and UNIX includes the Replication Center.



Capture

Changes to DB2 source tables are captured by a Capture program running at the source server. The Capture program uses Control Tables to determine which source tables are used in replication.



Control Tables for Capture and Apply

DRA-20

The DB2 source server can be:

- DB2 for z/OS and OS/390 Versions 6,7 and 8
- DB2 for iSeries on OS/400 V5R2
- DB2 for Windows and UNIX Version 8

Changes to Informix source tables are captured by triggers which are created automatically when you define the replication source.

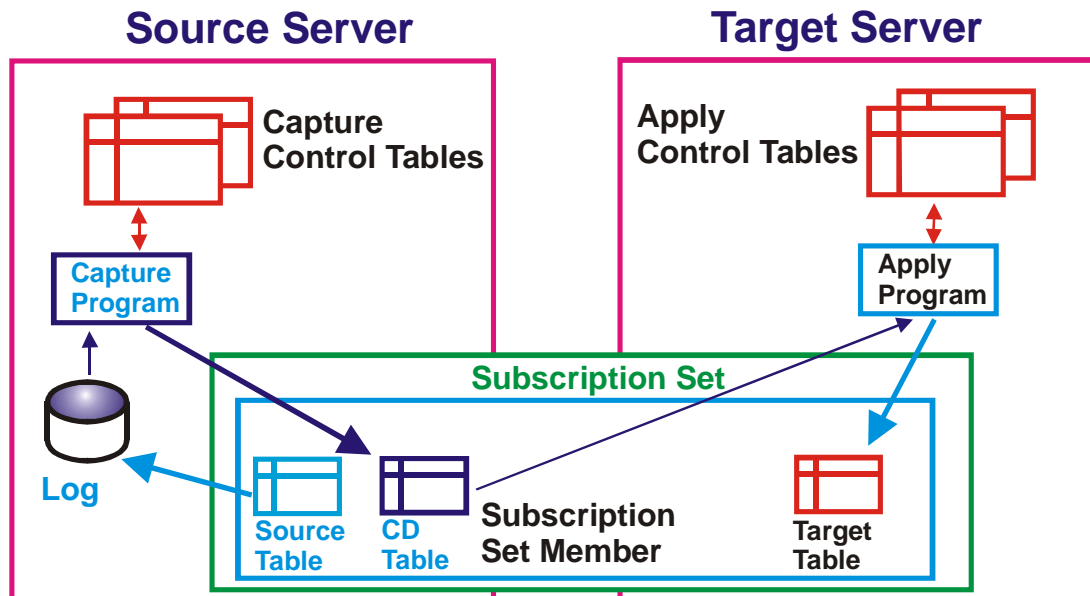
The Capture process can filter data by column.

The captured changes are stored in a table local to the source table and are automatically deleted (pruned) after they have been applied.

When changes are made to the source table, DB2 writes log (journal) records. These log records are used for database recovery and for replication.

The Capture program uses database interfaces to access log records:

- DB2 z/OS and OS/390 IFI 306
- iSeries RCVJRNE command
- DB2 Windows and UNIX asynchronous log read API db2ReadLog



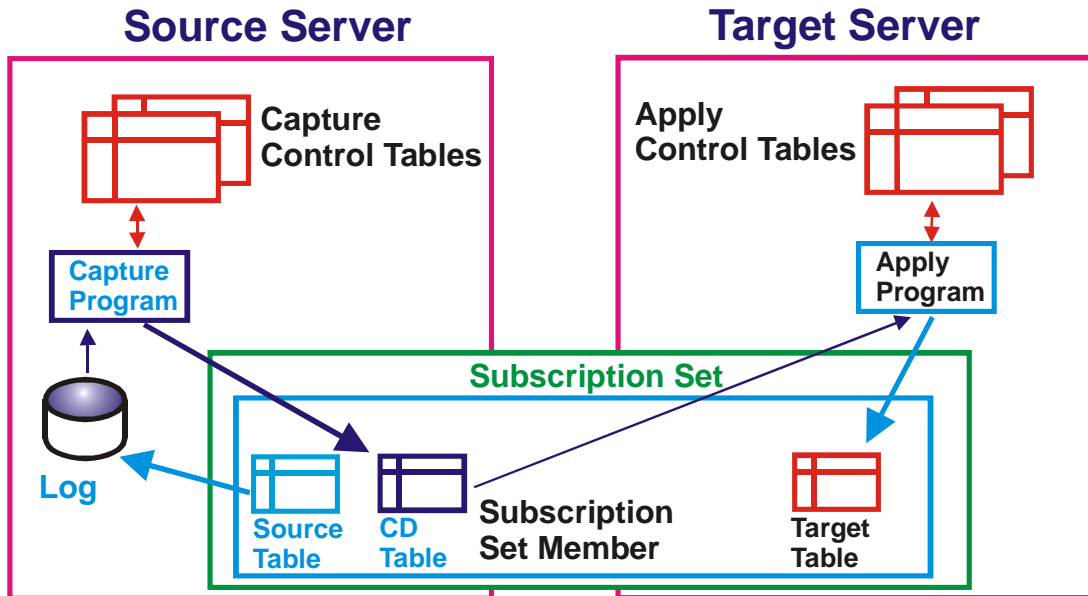
Capturing data from a Source Table

DRA-50

Each source table has a corresponding Change Data (CD) table where the captured changes are stored. You create the CD table when you define a replication source table using the Replication Center.

You can opt to capture a subset of the source table columns. You can also capture the values before the change is made (before-image columns) along with the values after the change is made (after-image columns).

Each change is uniquely identified by its log record sequence number (LSN).



Capturing data from a Source Table

DRA-50

The Capture program holds the changes in memory until a COMMIT is issued.

When a COMMIT is issued for a transaction that involves the replication source tables, Capture inserts the captured changes into the appropriate CD tables and stores the COMMIT information in the Unit of Work (UOW) control table.


If Capture detects a ROLLBACK, it removes the associated changes from memory.

You can improve throughput by running multiple Capture programs on a source server.

Each Capture has its own schema for control tables and its own set of CD tables. The schema is defined when you create the capture control tables using the Replication Center. You specify the name of the schema when you define your replication sources and targets and when you start the Capture program. The default capture schema is ASN.

The CD tables and the UOW table are always located on the server where the source table is located (except when iSeries remote journaling is used).

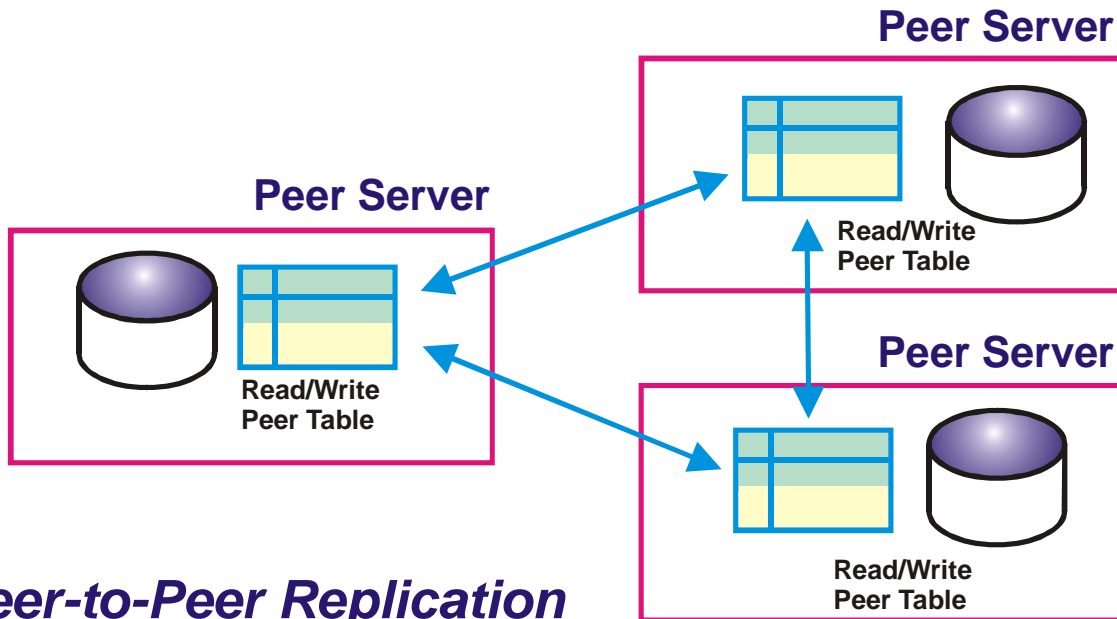
Bidirectional replication uses a Capture program running on both servers with one Apply program running at the target server to process changes in both directions.



Bi-directional Replication

DRA-90

Peer-to-peer replication has a Capture and an Apply program running at each server.

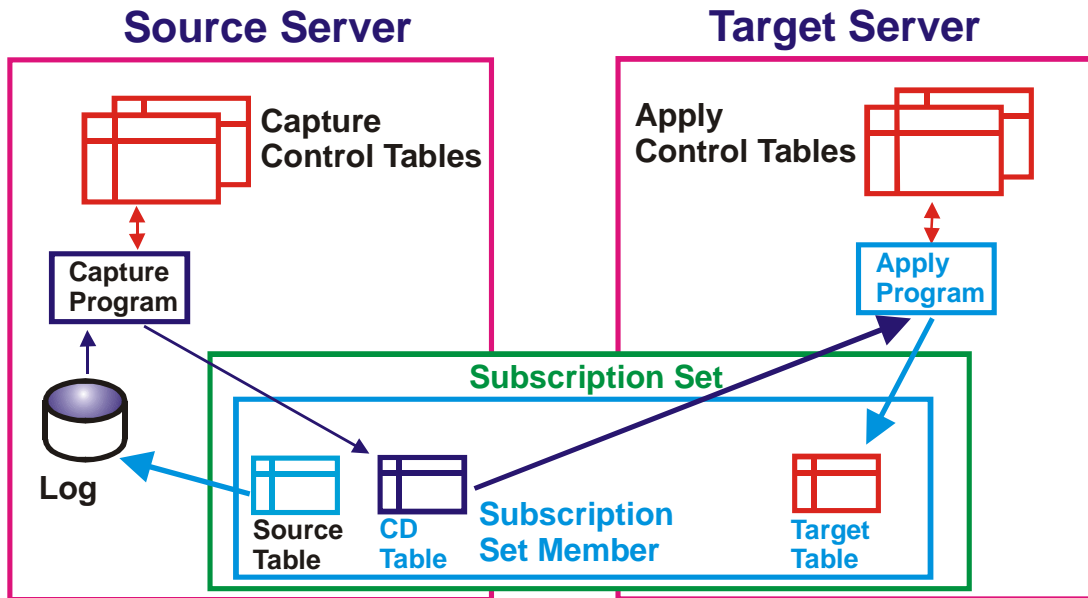


Peer-to-Peer Replication

DRA-100

Apply

The captured changes are applied to target tables by Apply programs.



Applying data to a Target

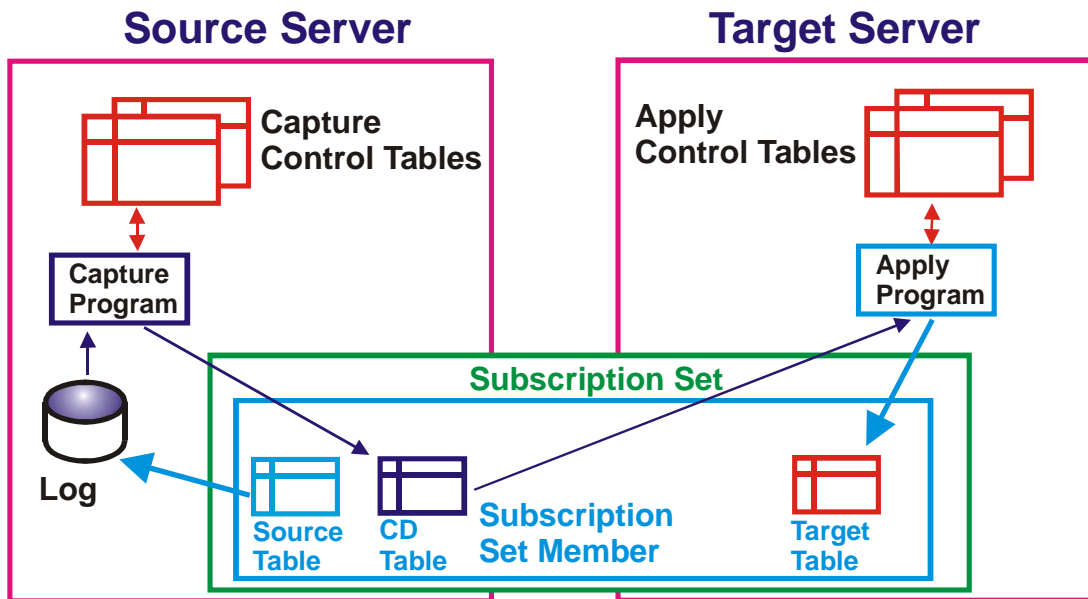
DRA-60

The Apply program can run on any server and must have connectivity to both the source and the target servers.

The data can be filtered by column, filtered by row, joined with other data (using views), and transformed with SQL expressions during the Apply process.

You will use the Replication Center to map a source table or view to a target table.

You do this by defining a subscription set which is a group of one or more target tables (subscription members) that will be processed as a unit by Apply.



Creating a Subscription Set

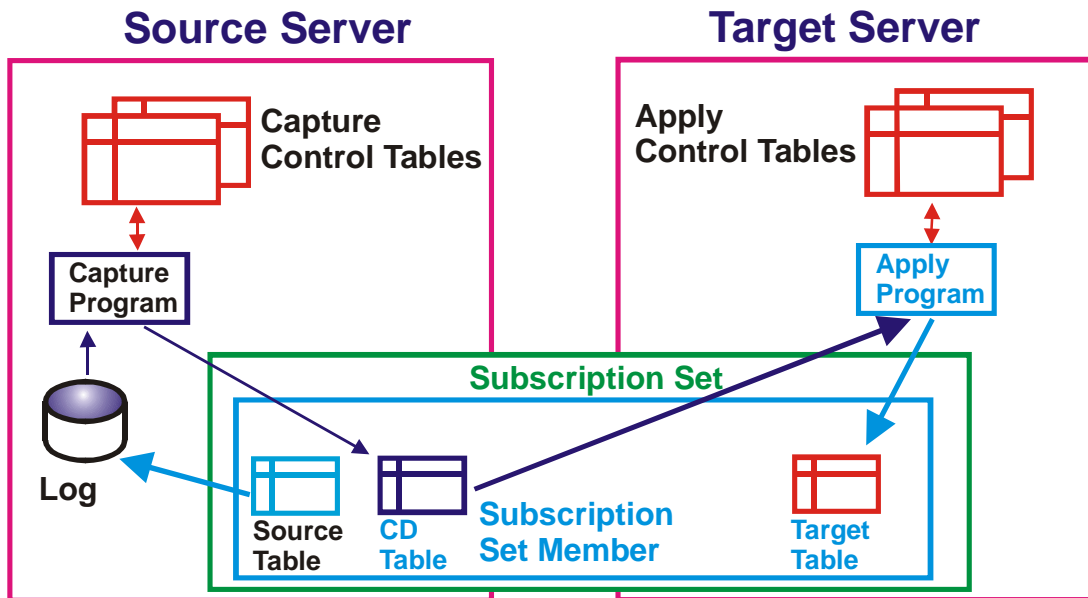
DRA-40

The changes from the CD tables are applied to each target table separately, in the same order in which they occurred at the source.

A single COMMIT is issued after the last member in the set is processed.

You can specify transactional replication for subscription members that are user copies (read only), point in time (PIT), or replicas and changes will be applied in the same order that they occurred at the source across all the members in the set.

If the target tables have DB2 referential constraints or other relationships that must be preserved, you should choose transactional replication when you define the subscription set.



Applying data to a Target

DRA-60

For the first initialization of the target tables, Apply selects all the rows from the source tables. If you have defined any transformations or filters, they will be processed. This is known as a full refresh.

You can use the Replication Center to perform a Manual Full Refresh. This will update the apply control tables when you do the initial population of the target tables outside of replication.

Following the full refresh, Apply selects changes from the CD tables and applies those changes to the target tables.

Some target table types may require a join of the CD and the UOW table during the Apply process. The join is required for bi-directional copies. However, the join is not required for user copy targets that do not have any column in the UOW table in their predicates.

The Apply program can be run as a batch process or as a task that runs all the time. You specify the schedule details when you define the subscription set. The schedule can be time-based, using an interval from zero seconds to one year.

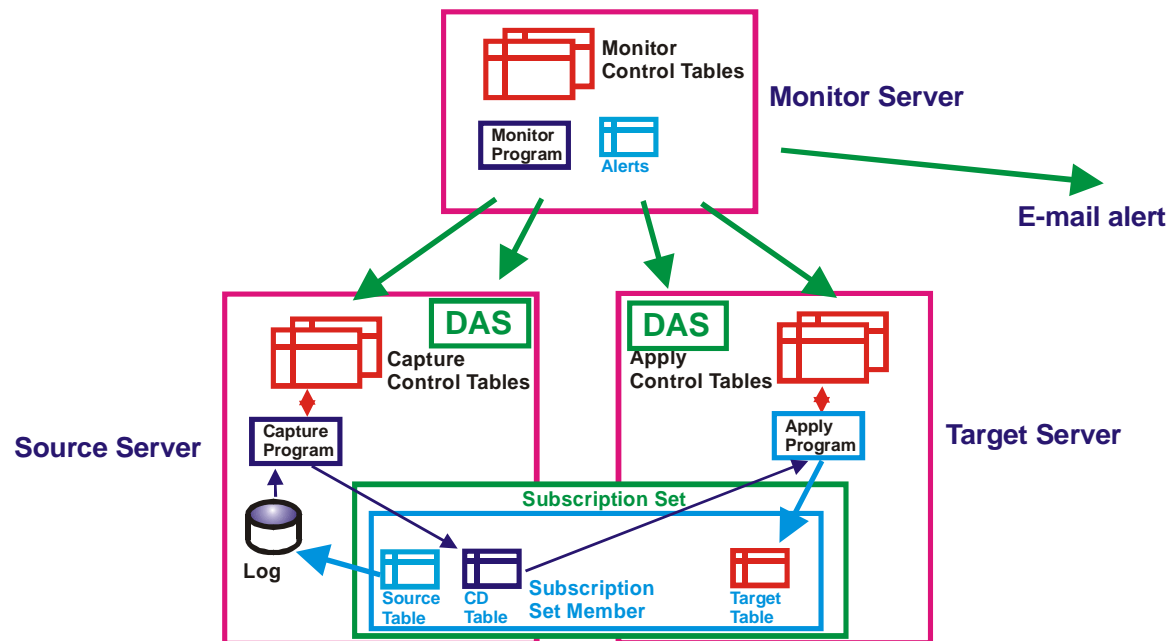
You can also schedule the Apply program using an event. You name the event and then insert a row into the Apply events control table, `ASN.IBMSNAP_SUBS_EVENT`, whenever you want Apply to start copying.

Alert Monitor

The Replication Alert Monitor is used to monitor SQL replication. It uses its own set of control tables, which are created with the Replication Center. You can define thresholds and events for Capture and Apply servers.

You can also define users or groups of users who will receive e-mail notification when an alert occurs.

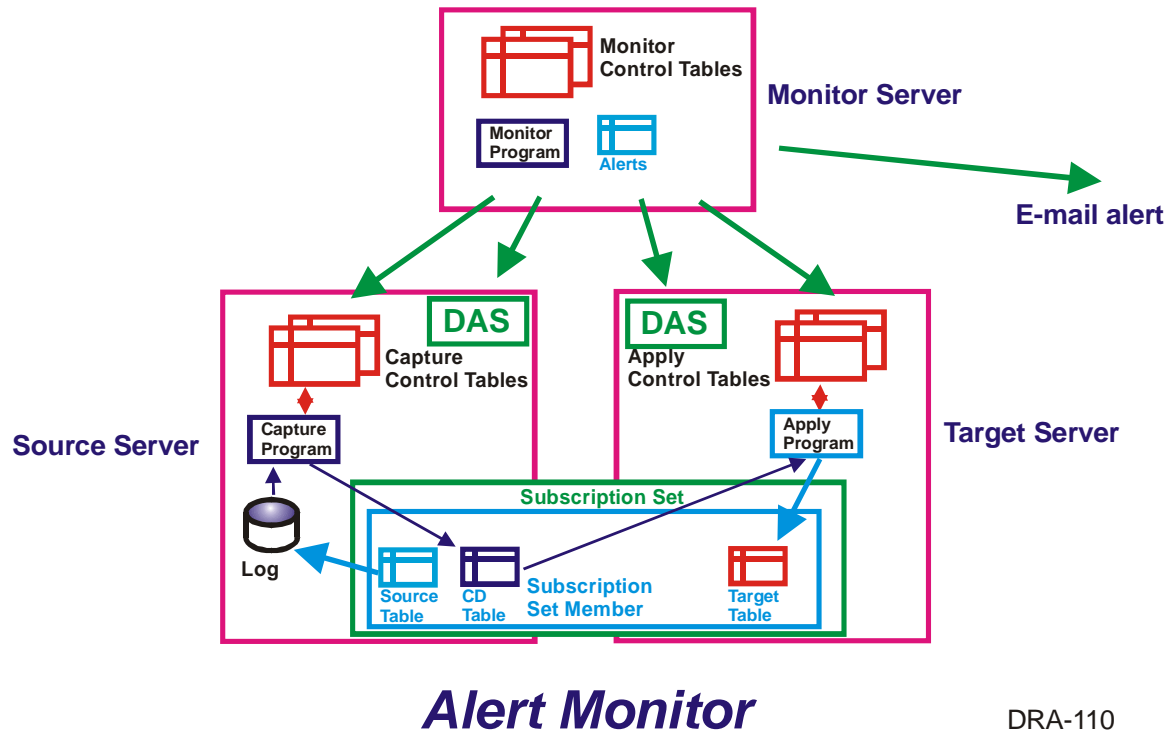
The server where the monitor runs is called the Monitor Server. A monitor server can monitor one or more local and/or remote servers. The Alert Monitor does not monitor the Capture triggers on non-DB2 source servers.



Alert Monitor

DRA-110

The Alert Monitor program collects information from the capture and apply control tables. It uses the Database Administration Server (DAS) installed on the Capture and Apply servers to receive remote commands and supply system information.



Alert Monitor

DRA-110

The DAS does not have to be installed on the Monitor Server unless you intend to use the Capture-Status-Down or Apply-Status-Down alert condition.

DAS is not required when monitoring iSeries Capture and Apply servers.

You can run the Alert Monitor from the Replication Center or by issuing a command.

Use the monitor interval parameter to specify how often the Alert Monitor checks the events and thresholds.

When a monitored event occurs (e.g. an error message is issued, or a monitor threshold is exceeded), the Alert Monitor inserts an alert into the ASN.IBMSNAP_ALERTS table. It also sends e-mail notification to any contacts you have defined.

Notifications are sent using an SMTP server such as Lotus Notes, Microsoft Outlook, and the sendmail program (UNIX).