# Informix Cloudscape Application Synchronization

*Informix® Cloudscape™ Application Synchronization goes beyond traditional data replication. Its synchronization of data, schema, and business logic uses patent-pending mechanisms to provide application-level consistency that is ideal for disconnected operations and the deployment and maintenance of remote applications.*

Unlike data replication, which is typically limited to the migration of data and data-related objects, synchronization allows coordinated movement of data, schema, and business logic between central and distributed databases in a distributed application architecture. This provides an ideal mechanism to support occasionally connected users, by enabling them to synchronize with a central site not just the data that they need, but also the schema that describes its structure, and the business logic that manipulates it. Cloudscape Application Synchronization has been used to solve a wide variety of business problems, including maintaining remote caches of news information, smart distributed product catalogs, deployed data marts, and supporting multilevel supply chains.

Cloudscape Synchronization refreshes database objects such as tables, views and indexes, as well as schemas, and Java™ code (JAR) files. It manages both the object definition and the data that the object contains. Cloudscape Synchronization manages the synchronization of objects between a *source* database, which contains the master copy of the object, and multiple *target* databases. The unit of synchronization is the *publication*, which lists the objects to synchronize, such as complete tables or only selected columns and rows. A publication is created using a statement very similar to a view definition.

Cloudscape Synchronization implements a hub-and-spoke topology. The hub or *source* database holds the "master" copy of the published objects while the *target* databases on the spokes holds their own local copies of the these objects. A *source* database can have many *targets*, but a *target* database has only one *source* database.

## LUCID™ Synchronization

Cloudscape uses an innovative synchronization technique that deploys data out from the central source, but moves business events back from the remote targets. This technique is called LUCID, meaning Logic Up, Consistent Information Down.

### Consistent Information Down

The Cloudscape Synchronization cycle begins when a target database is created using a publication from a source database. The remote application can read and modify the data in the target database autonomously using calls to database-hosted Java methods. For example, a salesperson in the field could enter orders into an application on a disconnected laptop. The application would make a call to a method that performs the order entry and updates the target database on the laptop. For purposes of Application Synchronization, the method called by the remote application is called a *Work Unit*.
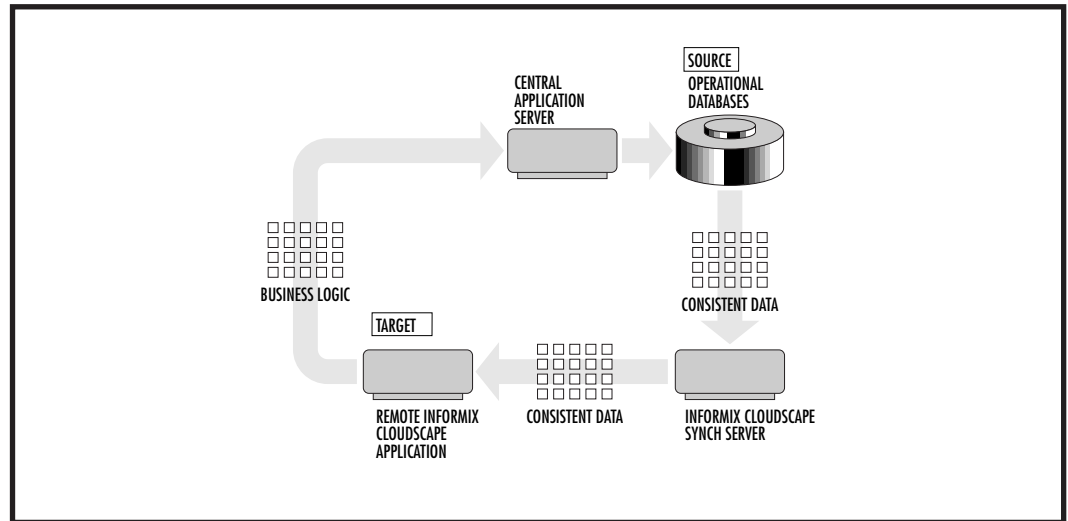
CENTRAL
APPLICATION
SERVER

SOURCE
OPERATIONAL
DATABASES

BUSINESS LOGIC

CONSISTENT DATA

TARGET

REMOTE INFORMIX
CLOUDSCAPE
APPLICATION

CONSISTENT DATA

INFORMIX CLOUDSCAPE
SYNCH SERVER

**Figure 1** Informix Cloudscape Application Synchronization.

**Logic Up**

Any time the target application is connected to the network, it can perform a *refresh* operation to synchronize its database with the source. During refresh, the business events that were executed at the remote target are forwarded to the source database. This is the Logic Up phase.

The source database re-executes the Work Units performed at the targets, but with one crucial difference: the code that executes at the source may be completely different from the code for the same Work Unit at the target. For example, in a sales-automation application, the target sale entry Work Unit may be concerned only with creating a well-formed order. At the source, however, the Work Unit may make calls to ERP applications, interface with other systems such as fulfillment and billing, and perform additional integrity checks that are not possible using the subset of corporate data deployed to the remote target.

As part of its processing, this central Work Unit will update the publication being used by the target. To complete, the refresh cycle, the target database is then refreshed, in a Consistent Information Down step, with the new information from the source publication.

**Conflict Resolution**

Any distributed system is subject to conflicts, because, by definition, disconnected operations are performed without up-to-date knowledge of the state of the central system. The classical example of this is two disconnected sales-persons, each of whom sells the last of a particular item. Each believes that their transaction is valid, but when combined, they are seen to be in conflict. Careful application design can reduce, but rarely completely eliminate, the chance of conflicts. Conflicts present a challenge to the distributed application designer because they undermine one of the common assumptions of database processing: that a transaction is durable. In a distributed system, committed transactions may have to be undone. Synchronization technologies can be judged on how well they deal with the issue of conflict resolution.

Data replication technologies lack the concept of work unit, and instead rely on propagating and merging database changes. This technique makes conflict resolution very difficult because the business context of the conflict has been lost. Instead, the application designer is faced with the problem of interpreting data row and column changes in order to determine an appropriate outcome.

The Cloudscape Work-Unit model is superior because it applies logical operations to the source database, rather than simply merging changed remote data into the central database. This allows the central application to make business decisions about remote events and the possible conflicts between them, as well as enabling it to perform additional integrity checks and interface with other applications and data sources. The Cloudscape technique ensures total central control of data integrity and the application of business rules, even over remote operations, and thus simplifies the code that needs to be deployed at the remote targets. Without the need for complex conflict resolution code on the target, application development and maintenance become much simpler, dramatically reducing the total cost of ownership of the application.

**Guaranteed Agreement**

In data replication systems, because of the conflict resolution problem, it is extremely difficult to guarantee that all copies of a given object are in agreement. Cloudscape Application Synchronization avoids these problems by ensuring that the data on the target exactly matches the source at the time of the refresh. There can never be data inconsistencies at the target after a refresh because its data is overridden by data originating at the source that is correct by definition.

**About Informix**

Based in Menlo Park, California, Informix Corporation specializes in advanced information management technologies that help enterprises in the i.Economy get to market quickly, generate new revenue, build a unique strategic advantage, and solve their most complex business problems. Informix offers customers a complete software infrastructure for the Web that delivers highly scalable transaction processing, personalized content management, integrated business intelligence, full multimedia capabilities and complete e-commerce solutions. For more information, contact the sales office nearest you or visit our Web site at www.informix.com.

4100 Bohannon Drive
Menlo Park, CA 94025
Tel. 650.926.6300
www.informix.com

**INFORMIX REGIONAL SALES OFFICES**

| | | | |
|---|---|---|---|
| Asia Pacific | 65 298 1716 | Japan | 81 3 5562 4500 |
| Canada (Toronto) | 416 730 9009 | Latin America | 305 591 9592 |
| Europe/Middle East/Africa | 44 208 818 1000 | North America | 800 331 1763 |
| Federal | 703 847 2900 | | 650 926 6300 |