

SalesLogix Data Synchronization Technology for Disconnected Sales Automation Users



SalesLogix Data Synchronization Technology for Disconnected Sales Automation Users

Table of Contents

Introduction	4
An Overview of the Data Synchronization Process	5
Keys to Successful Synchronization	6
Efficient Data Distribution and Transport	6
Standard Communications Platforms	6
Custom Conflict Resolution Rules	7
Automatic Distribution of Changes to the Application and Database	7
Scalability, Flexibility, and Performance	8
Powerful Error Handling and Recovery	8
The SalesLogix Data Synchronization Process	9
SalesLogix Synchronization Features	9
An Overview of SalesLogix Synchronization	9
Minimizing Synchronization Connection Times	10
Connecting to the Host	11
Securing Account Data	11
Determining Accounts to Synchronize	11
Resolving Conflicts	12
Automatic Synchronization of Customization and Schema Changes	12
Recovering from Errors	12
Administering Synchronization	13
SalesLogix Synchronization in Technical Detail	13
SalesLogix Synchronization Software	13
The SalesLogix Synchronization Server Cycle	13
Transacting	14
Applying	14
Sending	14
Cleaning	15
Running Agents	15
The Synchronization Client Cycle	15
Sending	15
Receiving	16
Applying	16
Recording Transactions	16
Error Logging	16
Error Recovery	16
Summary	17

Introduction

Mobile computing and the geographic dispersion of sales organizations have created new challenges for database systems. Web-enabled devices, handhelds, and laptop computers are now important tools for mobile salespeople, but they operate disconnected from corporate data processing systems most of the time. Internet access, while important, is often unavailable to the remote salesperson while traveling or at customer locations. Yet, companies that manage valuable customer and sales information must distribute it to field salespeople to help them close sales faster.

That makes the synchronization of data between host corporate systems and remote users one of the most critical technical issues in sales automation today.

Data synchronization enables remote salespeople to access up-to-date corporate data on their “disconnected” mobile devices. They maintain a subset of the master database and update that local data at the same time others are working with the same information, on either their own local databases or on the master database. When remote users synchronize their local data with the master database, the synchronization system coordinates all the updates so everyone receives accurate, up-to-date data. This is beneficial not only in making sales, but also because it helps corporate managers and sales teams share information created by field salespeople, such as meeting notes, schedules, and the opportunity data that feeds forecasts.

Key trends in selling make data synchronization even more important:

- Salespeople are spending more time out of the office with customers and prospects.
- Many salespeople are telecommuting—working from their homes rather than in corporate offices.
- Mobile salespeople and managers need access to real-time customer information and up-to-the-minute reports.
- As products become more complex and technical, salespeople operate more as members of selling teams, so the need to share information grows.
- Entire sales and marketing organizations are using computer-based customer, sales, and project information to sell more effectively.

While most sales automation systems offer synchronization capabilities, they vary widely in their capabilities and quality. Good synchronization is so critical to the efficiency and good operation of a sales automation system that its evaluation should be a primary part of any sales automation software assessment. It is more than a simple checklist item.

An Overview of the Data Synchronization Process

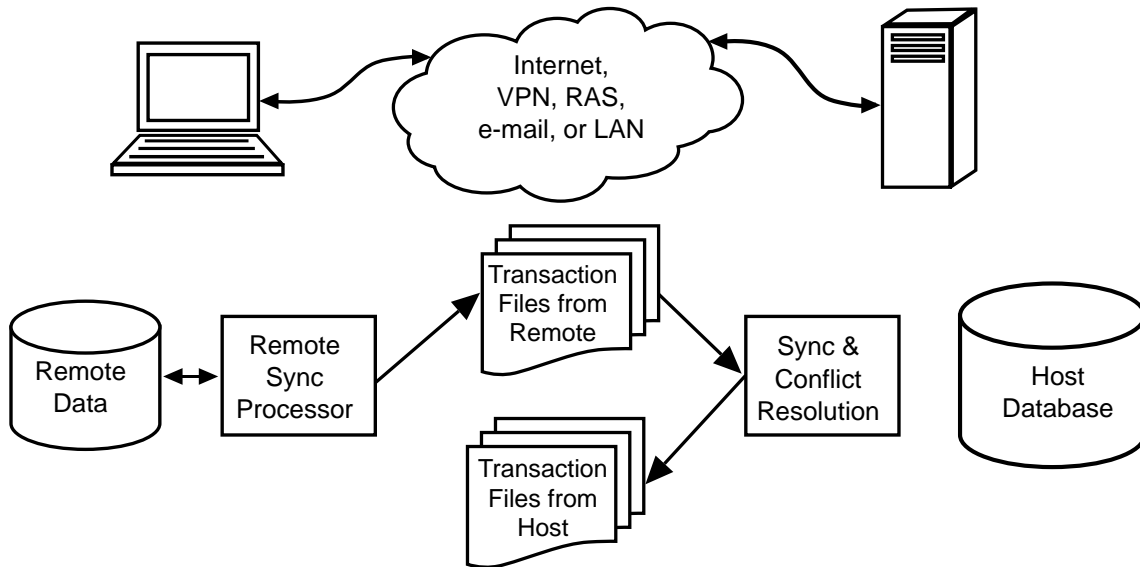


Figure 1. The synchronization process.

In most sales automation systems, field salespeople with laptop computers download pertinent subsets of the host database, manipulate and update the data, then reconcile their changes with the host database. The typical process of synchronizing data between the remote and host systems involves several basic steps:

1. Remote databases are created for mobile salespeople and branch offices. Each database is a relevant subset of the corporate database. For example, a field salesperson may need data only about his or her accounts or territory.
2. The field salespeople update their remote databases, while others make changes directly to the host database. Because different versions of the data now exist and are being updated at the same time, the synchronization system tracks changes and updates in both the remote databases and the host database. These changes, also called transactions, are stored in log files.
3. Remote salespeople connect to the host database using Internet connections (FTP) or remote network connections (RAS, VPN). Salespeople and managers who are based at the main office may connect via the local area network (LAN).
4. During the connection, the log files containing updated information from the respective databases are exchanged.
5. After the connection is completed, new data is applied to each database, providing each with current information.
6. When changes to identical data don't reconcile, data conflicts can occur. This happens when more than one user changes the same data at the same time on different computers. A method of conflict resolution is required to maintain consistency.

Keys to Successful Synchronization

Because remote database synchronization technology is not included with database management systems, most sales automation solutions include some synchronization features of their own. But they are not all equal. There are different methods of synchronizing data, different ways of logging and posting changes—and different levels of success. Many sales automation implementations fail because their data synchronization technology either did not work or became so inefficient and impractical that field salespeople refused to use it.

An efficient and useful synchronization system must have several key characteristics: efficient data distribution and transport; standard communications platforms; custom conflict resolution rules; automatic distribution of changes to the application and database; scalability, flexibility, and performance; and full error handling and recovery.

Efficient Data Distribution and Transport

Successful synchronization systems minimize the amount of data distributed to remote users, thus minimizing expensive connection and processing times. To make data distribution as efficient as possible:

- **Synchronize small subsets of the database.** Salespeople should automatically receive only the data that pertains to their territory, sales team, leads, or accounts—not the entire database.
- **Transfer only changes in data.** Some synchronization systems have remote users send all their database files or records to the host, then compare the data on the host system. This requires long connection times. Efficient systems log only changes to the most granular level of data—the field level—to minimize the size of data transfers. Data compression and encryption should also be used when sending data via modems or the Internet.
- **Apply transactions offline.** To minimize connection time, the synchronization process should not apply changes to the database while the remote computer is still connected. The connection should transfer files then hang up the modem so the user can apply the changes when it is convenient. This also enables the host system to post changes during off-peak hours.
- **Enable users to subscribe to other data as they need it.** Remote users, such as sales managers, may be authorized to access large amounts of data, but only need to work with a small subset of that data to get the information that is timely and important to them. An efficient synchronization system has flexible rules that allow users to subscribe to critical data as they need it.

Standard Communications Platforms

Field salespeople using laptops must connect with the host system to send and receive new data, as must disconnected branch offices. Practical synchronization systems should support standard remote networking and Internet connections to transport that data. Proprietary communication methods add cost and complexity for information system administrators and require remote users to learn yet another method of connecting. Using standard communications platforms also makes the connection process as easy as possible for remote salespeople.

To make the connection process as efficient and useful as possible, a sales automation system should:

- **Support standard remote network and Internet connections.** Remote connections should be automatically invoked to connect to the host. When users connect to the Internet or network, synchronization files should be automatically exchanged.
- **Allow LAN and WAN connections.** Occasionally, salespeople have access to high-speed LAN and WAN connections. The synchronization system should have a user-selectable transport method so they can choose the most efficient connection available.

-
- **Provide a simple and automatic connection for end users.** The process for connecting and exchanging transaction files should be simple and automated. Users should not have to wade through multiple screens or configure complex settings. The software should provide for scheduled connection and transfer of files with no user intervention.

Custom Conflict Resolution Rules

Conflicts occur when the same record or field on both the remote and master databases are updated and those changes are not identical. Rules must be written to automatically resolve conflicts in the data transferred to the host.

While most data is synchronized without conflicts, the likelihood of conflicts increases in team-selling and support environments, where many people deal with the same customer or prospect data at the same time.

For example, a field salesperson may work intensely with an account, trying to close a big sale. After an important meeting, he updates information about the opportunity in his computer, increasing the likelihood of closing the deal to 80 percent. But let's say that earlier that same day, the sales manager had reduced the same close percentage to 30 percent, based on information she received from different people at the same account. When the changes made by the field salesperson are posted to the master database, the different percentage numbers for this opportunity will conflict. Whose change should take priority and be retained?

Most synchronization systems for sales automation software perform extremely basic conflict resolution, if conflicts are resolved at all. Changes are posted based on who synchronized most recently. That is, the most recent change posted "wins." This scheme often creates data that does not reflect reality and can cause data integrity problems, especially when remote salespeople synchronize infrequently, as is often the case. A truly useful conflict resolution system should:

- **Recognize conflicts at the field level.** Because different users may provide different data at the same time on the same record, data conflicts should be reconciled at the field level, rather than at the record level. This ensures maximum data accuracy and efficiency.
- **Provide conflict resolution rules.** Most synchronization systems date-stamp transactions and post changes with a single "last-in wins" rule, which may not fit the way any particular company does business. A synchronization system should enable its administrator to configure conflict resolution rules that test changes and post them automatically. For example, when a conflict occurs, the changes from the account manager or remote salesperson could be set to take precedence, based on the conflict resolution rules which are, in turn, based on the needs of the business.
- **Allow for infrequent updates from remote users.** Synchronization systems should not require remote salespeople to dial in daily, or even regularly, to maintain data consistency. Salespeople may not have access to phone lines for their modems or may simply not dial in for several days. The system should recognize when updates from remote users contain "old" data and not apply out-of-date information.

Automatic Distribution of Changes to the Application and Database

While maintaining up-to-date information about customers, leads, schedules, inventory, and orders is critical, it is not the whole picture. For administrators to respond quickly to changing business requirements, a synchronization system must also automatically update remote users' applications to reflect changes made at the home office. Without this capability, distributing a new report to the field or adding new fields to the database, for example, require manual reconfiguration on each remote system. To automatically handle these kinds of changes, the synchronization system should:

- **Deliver remote application changes.** Typical sales automation applications include customizable database screens, reports, field pick lists, menu and toolbar settings, automation agents, custom programming, and more. When the system administrator modifies these components at the home office, they should be delivered to the field and automatically configured on the remote system during the synchronization process.
- **Distribute and apply database schema changes.** In response to changing business requirements, administrators often need to modify the database schema. This includes changes such as adding new database tables and fields to track information. When the host database schema is modified, remote databases should automatically be updated during synchronization without user intervention.

Scalability, Flexibility, and Performance

A synchronization system designed to support a business's needs for years to come should be capable of supporting large-scale field implementations with hundreds of users, even if the remote sales force is currently small. Some synchronization systems perform fine in small test environments but become impractical in real-world situations, especially with large groups. Besides using the most efficient means to distribute and post data, as discussed above, flexible support for modern client/server databases is critical to meet the demands of data synchronization for large numbers of users. The synchronization system should include:

- **Relational database support.** High-performance synchronization requires robust database capabilities and the kind of performance available only with relational databases such as Microsoft® SQL Server™ and Oracle®. The synchronization system should take advantage of standard SQL database technology and should support the database standards of the company.
- **Database independence.** The synchronization system must be database independent so different database systems can reside on the remote and host systems.

Powerful Error Handling and Recovery

Remote users rely on their sales automation systems for urgent information about orders, schedules, leads, and more. The sales automation system is critical to their success—and the company's. When hardware or software problems arise, administrators must be able to respond quickly to keep busy salespeople up and running. Strong synchronization includes:

- **Backup and restoration of remote databases.** If a laptop is dropped, a hard drive damaged, or data files accidentally deleted, the administrator must be able to restore the remote user's data and entire environment completely, based on the most recent synchronization with the host. This includes user settings and document files such as quotes, presentations, and correspondence.
- **Automatic error checking.** Remote communication with modems can sometimes cause problems. If a disconnection occurs in the middle of a synchronization session, some log files may not be transmitted to the host and others may be incomplete. To avoid database errors, the system should apply only those files that are complete, while automatically detecting missing and incomplete log files. Log files that were not applied to the host database because of errors should be automatically transferred and applied during a subsequent connection.
- **Complete administration tools.** Administrators should be able to manage the synchronization system from a central console. Tools should support central management of user settings such as passwords, security, and other synchronization profiles. For efficient resource usage, processing of log files on the server should be scheduled to run automatically during off-peak times.

The SalesLogix Data Synchronization Process

Data synchronization is designed as a core capability of SalesLogix. SalesLogix combines modern relational database technology, an efficient and flexible design, and centralized administration tools to deliver useful synchronization that supports changing business needs.

SalesLogix Synchronization Features

- **Easy for end users.** Simple “one-click” or automatic background synchronization with a “What’s New” view of all changes made to the Remote user’s data.
- **Efficient data distribution and subscription.** Users synchronize a minimum subset of data and can subscribe to other data when needed.
- **Custom conflict resolution rules.** Administrators can configure multi-level conflict resolution rules that test changes at the field level and post them automatically.
- **Support for multiple databases.** Microsoft SQL Server, MSDE, and Oracle database systems are all supported.
- **Open communications transport.** SalesLogix uses standard remote networking protocols such as remote VPN connection, Microsoft RAS, Internet FTP, WAN, and LAN.
- **Automatic distribution of customization and schema changes.** Customizations such as database screens, reports, field pick lists, and custom programming, and database schema changes are automatically distributed and configured on remote systems.
- **Automatic territory realignment.** Accounts, contacts, deals, and activities can be redistributed automatically by the system administrator.
- **Complete administration tools.** Configuration and maintenance tools include server process scheduling, archiving, and user administration.
- **Documents automatically synchronized.** Documents in the SalesLogix Library and other attached documents are centrally stored and synchronized automatically only to users who require the documents.
- **Server-based agents.** Efficient maintenance and processing are supported with scheduled running of reports, SQL scripts, Basic programs, and more.

An Overview of SalesLogix Synchronization

The following figure provides an overview of the SalesLogix synchronization process between remotes and the main office. The main office contains the host database, the central set of shared synchronization folders, and the SalesLogix Synchronization Server. It supports all users who connect to the host database.

SalesLogix remotes include both Remote users and Remote Offices. Remote users work with their individual databases on their own machines. Remote Offices can support workgroup users who connect directly to the Remote Office database via a network. All remotes (users and offices) have a subset of synchronization folders and either the Synchronization Client or the Remote Office Synchronization Server.

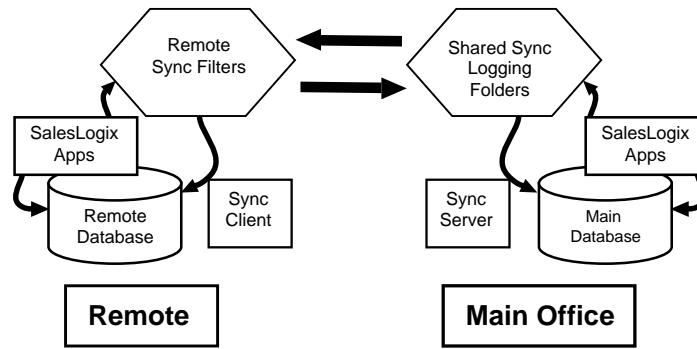


Figure 2. The SalesLogix synchronization process.

As each user makes changes to his or her database, SalesLogix tracks the field-level changes and stores them in transaction exchange files (TEFs). SalesLogix tracks the time and date of the change, the user and site that performed the change, and other details. Tracking field-level details about each transaction makes possible sophisticated conflict resolution, efficient processing, and data roll back so erroneous updates can be undone.

Unlike synchronization systems based on index scanning or stored procedure schemes, SalesLogix logging makes it possible for updates to be applied across different database platforms. Transaction log files generated by a SQL Server or MSDE database can be applied to an Oracle database with no conversion required.

The SalesLogix Synchronization Client for the Remote user or Remote Office transfers TEFs to the shared Sync Logging folders at the Main Office. It also moves files from the shared Sync Logging folders to the remote synchronization folders. These TEFs are then applied to the remote database, so it contains the same information as the host database. Note that the remotes never connect directly to the Host Database during this process, but only to the shared Sync Logging folders at the Main Office.

The SalesLogix Synchronization Server at the main office processes the TEFs received from the Remote user or Remote Office and applies them to the host database. If other remotes also need this updated information, it generates TEFs for them. Additionally, the Synchronization Server processes the TEFs generated by users at the main office, creating TEFs for the Remote users or Remote Offices that require those updates. This ensures that all the databases will eventually contain the same information as the host database.

Minimizing Synchronization Connection Times

TEFs can be transported via Internet (FTP) or remote network (RAS, VPN) protocols. Users can choose the most efficient connection method available. For example, a field salesperson could choose to connect directly over the local network when she is at the home office.

Because Remote users typically connect with low-speed, low-bandwidth modems from the field, minimizing expensive connection time is critical. Before a connection is made, the SalesLogix Remote Client creates TEFs containing only database changes. These TEFs are automatically compressed and sent to the host database during the connection. After the remote and host exchange TEFs, the communication session ends, and the Remote Client processes the transactions offline.

In contrast to the offline store-and-forward transport method used by SalesLogix, some sales automation products require the data exchange and transaction processing to occur during the online communication session. According to Gartner Group, "users should be aware that online transaction posting and conflict resolution will take longer than simply transferring files and disconnecting, and could potentially 'steal' processing capacity from online transaction processing if many mobile users dial in simultaneously." (Gartner Group, February, 1997)

The result is delays for everyone.

Connecting to the Host

The process of synchronization is simple and fast. Users may choose to synchronize from a menu or click a toolbar button to display the SalesLogix Synchronize With Host dialog box. SalesLogix can also automatically detect an Internet or network connection and synchronize in the background without user intervention.

Once a connection has been established with the main office, TEFs and documents are exchanged. All files are compressed and encrypted before sending for fast, secure transport. Users can choose to apply changes immediately or wait until later. The Remote user can continue to use the database while changes are being applied.

Securing Account Data

The centrally administered account security system used in SalesLogix determines which accounts any particular user has access to and which accounts that user can synchronize. If a user can access a specific account record, he or she also has access to all the information associated with that account: contacts, opportunities, notes, histories, activities, and so on. SalesLogix supports several methods of sharing and securing information. Data can be shared with:

- **No one.** A user can own accounts so only he or she has access to them. By default, managers automatically have access to related information about all accounts owned by individuals on their teams, except for private contacts and activities.
- **Team members.** Information can be shared among members of sales teams. All members of a sales team automatically share information assigned to their team. Certain members of the team may be designated owners and have special rights to add and delete account data and change security rights for the team.
- **All SalesLogix users.** Account information can be shared with everyone using the SalesLogix system, enabling all users to access the information. When Remote users (or those who are team owners) synchronize, new information about all the accounts they own are automatically synchronized. This includes new leads added to the database. The administrator controls which data is automatically synchronized to minimize processing and connection times.

Determining Accounts to Synchronize

A Remote user may have access to a large amount of customer information, but should only synchronize a subset of that information to minimize processing and connection times. SalesLogix enables Remote users to subscribe to the additional information they need and synchronize it to stay up to date. Using this feature, Remote users can create basic selection rules to tailor what additional data is synchronized.

For example, a company's Western Regional Sales Manager may have access to several thousand accounts in her territory. But it is likely not necessary for her to keep data about all of them on her laptop and get updated information every time she synchronizes. She can use supplied rules or set up ones that enable her to receive updates about all accounts with specific characteristics, say, opportunities of more than \$25,000, as well as all national accounts. Then, any accounts that meet these criteria will automatically be synchronized to her laptop.

In addition to rule-based subscription, Remote users can browse a list of available accounts and select individual accounts to subscribe to. When they no longer need those account updates, they can unsubscribe to minimize the subset of data synchronized.

Resolving Conflicts

In addition to containing a field's updated value, a SalesLogix TEF also has the "old" value held by that field in the master database. Before applying a database change, SalesLogix looks at the database field it is about to update and compares its current value with the master database's "old" value. If the two do not match, there is a conflict. If SalesLogix did not retain the field's "old" value, conflicts could only be resolved by a simple date/time test. The result would be more errors.

The SalesLogix administrator configures the system so that, when SalesLogix identifies a data conflict, it automatically applies as many as three conflict resolution rules to determine whether the current transaction should be applied.

Any three of the following conflict resolution rules can be applied in any order:

- A workgroup user made the current change
- A Remote user made the current change
- The record's owner made the current change
- The current change was made most recently

If none of the three qualifications are sufficient to enable SalesLogix to decide whether to apply the change, then the most recent change is applied.

Automatic Synchronization of Customization and Schema Changes

In addition to changes in customer data, SalesLogix synchronizes information so that customizations from the main office can be seamlessly distributed to Remote users. With this capability, there is no need to reinstall software or create new remote databases to make most customizations available to all users.

All database screens, reports, templates, processes, pick lists, custom programming, menu settings, saved queries, and other application components are stored as files in the database and automatically sent to Remote users when they synchronize. This allows for quick and seamless modification of the application without the time-consuming need to provide manual updates to each individual user. The SalesLogix administrator can configure the distribution of these components, called plug-ins, to different groups of users, as needed.

Changes to the host database's schema are also automatically sent to all remote databases during synchronization. The SalesLogix administrator can, therefore, add and delete database tables and fields in the host system in response to changing business needs, then automatically and rapidly roll out the changes to the field.

Many user settings such as window defaults, calendar preferences, and other options are also stored in the database, rather than in the Windows registry or INI files outside the database. Users can, therefore, log in to SalesLogix from different computers, yet retain their personal settings. Also, if a user has hardware or software problems in the field and his or her database needs to be rebuilt, this simplifies the process of restoring the user's settings.

Recovering from Errors

SalesLogix supports a roll back capability to re-process lost data if connection or hardware problems occur. A permanent log file on the server makes this possible. The permanent log file is a chronological accumulation of all changes produced by SalesLogix users. Should a problem occur with the transaction file transfer, SalesLogix will automatically request a new copy of the missing data and apply it to the database.

Administering Synchronization

SalesLogix Administrator is the central tool used to control user configuration, process scheduling, data security, conflict resolution, and default settings for synchronization and subscription. The SalesLogix Monitor Console may be used from any computer with a network connection to monitor the synchronization process and the SalesLogix Synchronization Server.

Distributing territory realignments can be an arduous manual task with some sales automation systems. With SalesLogix, the administrator can transfer accounts from one user or sales team to another based upon geography or other business rules. Realignment updates are sent to Remote users during synchronization. Thus, when territories change, data can be removed from one user's laptop and downloaded to another's computer automatically. All account, contact, opportunity, and calendar information is reassigned to the new user.

SalesLogix Synchronization in Technical Detail

SalesLogix Synchronization Software

SalesLogix includes two software products required for data synchronization:

- **Synchronization Server.** The Synchronization Server is software that distributes data and manages data exchange between Remote Clients and the host system. It runs on a workstation connected to the database server and is the host system software that supports Network and Remote Clients, as well as Remote Offices.
- **Synchronization Client.** The Synchronization Client is used on remote computers that are occasionally connected to the host server. It utilizes a local SQL database and includes support for data synchronization and subscription. Remote Offices use the same Synchronization Client software and are treated much as Remote users.

The SalesLogix Synchronization Server Cycle

The Synchronization Server cycle has four phases:

- Transacting
- Applying
- Sending
- Cleaning

At the end of a complete synchronization cycle, any other programs that have been queued to run, such as agents or remote tasks, are executed.

Note that the synchronization server discussed here is the primary synchronization server at the main office. It is possible to set up multiple synchronization servers at the main office. In fact, this may be a necessity to improve performance when dealing with several dozen Remote users. Any additional synchronization servers are secondary synchronization servers. Remotes can be assigned to secondary synchronization servers. This helps balance each server's processing load. Network users at the main office are always assigned to and supported by the main synchronization server.

Transacting

1. When the SalesLogix synchronization cycle begins, the first step is to process the files in the WGLogs folder of the Sync logging path at the main office. These files are generated by any SalesLogix application (Architect, Administrator, Sales Client, or Support Client) that connects to the host database. The Synchronization Server can only process WGLogs that have been closed. WGLogs are closed when SalesLogix is closed normally—that is not via Ctrl+Alt+Delete—or when the user refreshes SalesLogix.
2. Because these files reflect changes that have already been made to the host database, there is no need to apply them to that database. The Synchronization Server opens these files and all the information in them is written to the Confran.stm file for use later in conflict resolution. The same process is repeated for the Shared Logs folder, which is used only if the system has more than one Synchronization Server.

Applying

1. Synchronization Server now processes the Infiles folder of the Sync logging path. Remote users and Offices send their TEFs to this folder when they synchronize with the main office. TEFs with invalid keybase values are placed in the Old folder. All TEFs with invalid source or target location (site) codes are placed in the Rejected folder.
2. TEFs with a CMD prefix are decompressed and compared to the Archives folder. Once this comparison confirms the validity of a TEF file, it is deleted from the Archives folder.
3. Next, requests from Remote users for Library documents are applied so the Synchronization Server can send requested files to the appropriate Remote user.
4. TEFs with a FILE prefix are placed in the Documents folder.
5. TEFs with a ZIP prefix are decompressed and compared to the Confran.stm transaction log for conflict resolution. Confran.stm stores a record of all changes made to the host database for a specified time interval, along with the rules that govern conflict resolution. Both the interval and the rules are specified in Administrator. If a conflict is discovered, the conflict resolution rules determine which transaction is applied to the host database. Confran.stm then is updated with the change and a log file is placed in the WGLogs folder so that changes are sent to the appropriate Remote users.
6. The Synchronization Server now processes the subscription rules used to minimize the number of accounts sent to Remote users. Subscription and unsubscription rules can be based on any field in the database. A TEF is generated if the subscription rules specify that the Remote user requires additional accounts or if a Remote user asks that an account be removed from his or her database.

Sending

1. SalesLogix now opens the log files in the WGLogs folder and writes them to the Write Cache folder on the Synchronization Server as it compresses them into TEFs. Once the log files are converted to TEFs, the logs are deleted. Log files in the SharedLogs folder are also processed this way.

-
2. Any missing TEFs which were requested by the remotes are copied from the Archives folder into the Outfiles folder in the Sync Logging directory. All confirmed TEFs are deleted from the Archives folder.
 3. Any Documents or Library items are compressed into TEFs.
 4. Once all these TEFs are processed, they are placed into the Archives folder for confirmation and into the Outfiles folder for transmission to Remote users the next time they synchronize.

Cleaning

Once subscription is complete, the SalesLogix Synchronization Server begins cleaning up. Synchronization Server processes Conftran.stm and removes any TEFs that are older than the time specified for conflict resolution in Administrator.

Running Agents

Agents are automated tasks which are executed after the SalesLogix synchronization cycle is complete. Any synchronization files resulting from these activities are then copied into the Outfiles folder so they are ready for remotes to pick up as they synchronize with the host database.

The Synchronization Client Cycle

This process begins only after a Remote user establishes a connection to the main office and initiates synchronization.

The Synchronization Client cycle has four phases:

- Sending
- Receiving
- Applying
- Recording Transactions

If errors are detected during synchronization, they are logged and corrected after these phases are complete.

Sending

1. Remote users insert, update, and delete data directly to their local remote database. These changes are recorded as TEFs and placed in the Outfiles folder.
2. The TEFs in the Outfiles folder are copied to the Archives folder for error checking.
3. The TEFs in the Outfiles folder are then transmitted to the Infiles folder of the Sync Logging directory at the main office.

Receiving

1. The TEFs in the Outfiles folder of the Sync Logging directory at the main office are transmitted to the Remote user's Infiles folder. Only TEFs matching the Remote user's location code are transferred.
2. Any received TEFs which are determined to be outdated are placed in the Old folder.
3. Any TEFs received with an invalid site code are placed in the Rejected folder.
4. Any TEFs received which are confirmations from previously received TEFs are sent to the Archives folder. All confirmed TEFs are deleted from the Archives folder.
5. Any TEFs "looking" for a missing TEF are sent to the Archives folder. Once the missing TEF is located, it is transferred to the Outfiles folder.

Applying

The Synchronization Client now processes all the files in the Infiles folder.

Recording Transactions

Information is captured about each Remote user during the synchronization cycle and is displayed in SalesLogix Administrator. The SyncSequencing table is updated with the Last Send Date, Last Send File, Last Apply Date, and Last Apply File information. Synchronization Client and Synchronization Server determine the next TEF file to apply using information in the ApplyIndex column of the table and determine the next TEF file to create based on information in the SendIndex column in the table.

The SyncSequencing table contains the following fields:

- SourceSite—source location of the log files
- TargetSite—target location of the log files
- SendIndex—next sequential log file to be created by the Synchronization Server or Synchronization Client
- LastSend—date and time of the last log file sent by the Synchronization Server and Synchronization Client
- ApplyIndex—last sequential log file applied by the Synchronization Server and Synchronization Client
- LastApply—date and time of the last log file applied by the Synchronization Server and Synchronization Client

Error Logging

If errors occur during a synchronization cycle, they are written to the file SyncErrors.txt in the synchronization folder, which can be opened with any text editor. These errors are not critical, but they can prove useful when troubleshooting.

Error Recovery

If Synchronization discovers a missing TEF, a record is created in the SyncFileTracking table with a FileStatus of 0 (zero). The TEF is then requested by Synchronization and sent to the host or remote database with a CMD prefix. Once the request has been received

and applied to the database, the requested TEF is recovered from archives and placed in the Outfiles folder to be sent back. Synchronization then applies the requested TEF and any subsequent files, and updates the FileStatus column in the table to 2.

If the request is received, but the TEF is missing, a TEF bearing the name of the missing log file is created. Synchronization then applies the missing and any subsequent TEFs, updating the FileStatus column in the SyncFileTracking table to 1.

The SyncFileTracking table contains the following fields:

- FileTrackID—unique ID of a record in the file-tracking table
- Tranfile—name of the requested log file
- SourceSite—source location of the log files
- TargetSite—target location of the log files
- FileStatus—each missing TEF will be identified with a file status of either 0, 1, or 2:
 - 0—missing TEF has been requested
 - 1—missing TEF has not yet been recovered
 - 2—missing file has been recovered

Summary

Data synchronization is a critical technology which enables salespeople to share up-to-date customer and sales information in the field, even when disconnected. Approaches to synchronization in sales automation solutions vary widely—as do the amount of success they achieve. Evaluating a solution's synchronization system can help avoid problems that have doomed many sales automation projects to failure.

That evaluation should include several important questions:

- Is data distributed and transported efficiently?
- Are standard communications platforms used?
- Can conflict resolution rules be customized?
- Are application and database changes automatically delivered?
- Is synchronization flexible?
- How well are errors and error recovery handled?
- How well does it perform overall?

The judgment as to which synchronization system best meets a particular company's needs is, of course, a business decision. It must support the overall goals of a sales automation system, including providing salespeople with more time to sell and real-time access to customer information, fostering a more efficient sales process, and yielding up-to-the-minute and accurate forecasts. The synchronization system that best meets these needs best serves its users and owners.

NOTES



Majestic Cosulting Group Inc.

3001 Lincoln Dr. West Suite F

Marlton, NJ 08053 U.S.A.

☎ 1 . 800 . 403 . 9400

www.majestictechnology.com

13SY60XXWP01 09/02