SYBASE®
An SAP Company

TECHNICAL WHITE PAPER

Database Replication

Technical White Paper

EXECUTIVE SUMMARY

Database replication—the distribution of information across a series of database servers, which may be located in the same place, or scattered across the globe-was once reserved for large, sophisticated organizations. The past few years have seen wider application of database replication solutions across all segments of the industry.

This white paper provides you with a basic understanding of this important technology, along with a series of best practices for you to consider when implementing a database replication strategy.

COMPLEXITY OF A DISTRIBUTED DATA INFRASTRUCTURE

When you consider the complexity of an average IT organization, the following was the typical scenario:

- Information was kept in a central location, under lock and key.
- Data was generally measured in kilobytes; only the largest and most sophisticated applications worked with megabytes.
- Reports were run in batches; two or more hour delays for results were common.
- Single vendor environments were prevalent.

Today, the IT landscape is riddled with new complexity and struggles against the realities of doing business in a global economy.

Running an IT organization in today's corporate landscape is extremely challenging. Information is spread across servers globally, with consumers located at various locations worldwide. The rate at which data is distributed is accelerating, too. As if these challenges didn't present enough of a headache, the quantity of data that consumers use is expanding exponentially, as well.

Budgets, of course, are not keeping pace with the need for IT resources to manage the increase in data quantity and dispersion, and as you can imagine, most organizations need to staff the additional expertise required to face the challenges presented by today's IT landscape.

Considering Business Realities

For IT managers struggling to adapt to these technological complications no relief is in sight on the business front. In fact, the overall enterprise is dealing with pressures that add to the IT burden.

To begin, globalization has mandated 24 x 7 x 365 access to information. Many organizations have operations around the world, running around the clock. Management expects instantaneous access to operational intelligence from these far-flung parts of the enterprise. Stale or static data is no longer acceptable.

The unprecedented financial and corporate governance failures of the past few years are now translating to massive new regulatory requirements. In order to comply with these regulations, accurate information is vital. For those organizations unwilling or unable to toe the regulatory line, civil and criminal penalties will be harsh.

To add fuel to the operational and regulatory fires, most organizations are living with a competitive environment that they could not have imagined even a few short years ago. There are now more rivals in more locations, translating to significant top and bottom-line pressures. Management is under the gun to reduce costs while boosting profitability. Operational efficiencies are often seen as the best way to make this happen, and the IT organization is often first in the cost-cutting crosshairs.

Users themselves are far more sophisticated than ever before. "Generation Twitter" has come to expect instantaneous access to information where and when they want it. To these impatient consumers, latency is not acceptable. Downtime is even harder for them to imagine, let alone tolerate. In fact, they aren't even willing to put up with a brief system outage to allow for application or infrastructure upgrades.

Finally, executives and line-of-business managers are now employing sophisticated business intelligence software that consumes data in unprecedented ways to help make real-time operational decisions. With stale or incorrect information, inaccurate and erroneous decisions are inevitable. Delayed data is often just as bad as no data at all.

DATABASE REPLICATION TO THE RESCUE

Increasing numbers of IT organizations are turning to database replication as a way to deal with the pressures of operating in today's globalized corporate environment. They are doing so largely because with advancements in technological products, moving large volumes of information across networks and onto destination servers is no longer cost-prohibitive.

Understanding Database Replication

At its core, database replication means the process by which information is created or updated at one location, and then distributed to other destinations. Database replication also requires that this information be kept synchronized without the need for user intervention. In fact, from the perspective of the information consumer, well-executed database replication should simply mean that data is available where and when it's needed.

The exact mechanism of database replication is a complex topic. You can use any of several approaches to replicate information across a network. Here are some popular techniques, along with some of the challenges imposed by each approach:

- **Dump and reload.** Data is extracted from one location and loaded into one or more destinations. Hours, days, or even longer gaps may occur between the creation of data and its propagation. This approach represents a one-way replication strategy that doesn't allow for remote updates.
- **Two-phase commit-driven distribution.** The database engine includes destination servers as part of every transaction. This adds significant overhead to the amount of time it takes to complete a transaction, which makes this approach unsuitable for highly time-sensitive applications.
- **Snapshots.** An image of the database is taken at a given point in time. This image is then loaded onto one or more destination servers. With this approach, recipients work with older data, which makes this technique not appropriate when real-time information is needed. This style of database replication is also not cut out for remote updates.
- **Database triggers (non-transactional).** A trigger copies data to one or more remote locations; however, no transactions are invoked to guarantee successful delivery. While this reduces overhead, it's easy to imagine scenarios where databases get out of sync, which can lead to all sorts of problems for the information consumer.
- **Database triggers (with transactions).** Because transactions are a well-proven method for guaranteeing the integrity of data modifications, this approach increases reliability of trigger-driven database replication. Of course, the cost of this increased integrity is reflected in decreased application performance caused by transactional overhead.
- **Log-based.** Employed by Sybase Replication Server, this is the most efficient and effective way of delivering replication. Instead of mandating intrusive and costly interaction with the underlying database, log-based replication uses direct reads from the online transaction logs.

Sybase Replication Server®

Sybase Replication Server employs a distributed architecture, with a collection of capabilities to guarantee delivery of data alterations. In addition to Sybase's own Adaptive Server® Enterprise (ASE), Sybase Replication Server is designed to work with a variety of data sources, including

- Oracle
- IBM DB2
- Microsoft SQL Server

SYBASE REPLICATION SERVER SUPPORTS MULTIPLE PLATFORMS, CHARACTER SETS, AND LANGUAGES. THIS FLEXIBILITY GIVES DATABASE ADMINISTRATORS THE OPTION OF REPLICATING INFORMATION FROM ANY DATA SOURCE TO ANY DESTINATION.

Simply enabling multifaceted data replication scenarios isn't sufficient. To help enterprises administer these intricate replication configurations, Sybase provides the Replication Server Manager. This is a powerful three-tier management tool with a graphical user interface where system administrators can create replication environments, monitor the status of the individual servers, and manage the flow of replicated data, all from a single desktop.

Modeling tools have a long history of helping database and development professionals visualize, document, communicate, collaborate and create systems with greater clarity, accuracy and predictable reliability. Sybase PowerDesigner® introduces a new graphical modeling paradigm, the Information Liquidity Model (ILM) that can both graphically diagram the replication environment as well as generate Sybase Replication Server implementation code from these models.

Since database replication is such a performance-sensitive application, Replication Server offers a series of throughput-related features to keep information flowing as rapidly and effectively as possible. These include:

- **Dedicated database replication functionality.** Database replication has the potential to bog down a database server with extraneous overhead. To combat this possibility, Sybase Replication Server is delivered as a specialized add-on solution, rather than as a "bolt-on" to a database server. This architecture is minimally invasive to the source database.
- **Continuous real-time log-based transactional data capture.** This tactic, which reads directly from the online re-do log, is faster than trigger-based database replication.
- **SQL statement replication.** To minimize the amount of information sent over the network, only the SQL statement itself is sent to the destination. This is an ideal complement to the previously described log-based replication strategy implemented by Sybase.
- **Configurable network routing.** Sybase Replication Server gives administrators the option of selecting the network route that replicated information will follow when being propagated to remote destinations.
- **Enhanced transaction performance in the destination database.** Because Replication Server works with products from different providers, it includes vendor-specific throughput optimizations and tweaks for each supported platform.

ORGANIZATIONS OF ALL SHAPES AND SIZES RELY ON SYBASE REPLICATION SERVER TO EFFECTIVELY DISTRIBUTE AND PROTECT THEIR VITAL INFORMATION. IT'S BEEN PARTICULARLY WELL ADOPTED IN HIGHLY CHALLENGING ENVIRONMENTS SUCH AS FINANCIAL SERVICES AND CAPITAL MARKETS. THESE ENTERPRISES ARE EXTREMELY PARTICULAR WHEN IT COMES TO THEIR THROUGHPUT AND LATENCY REQUIREMENTS.

REPLICATION SERVER FACILITATES CARGILL'S TRANSACTION-INTENSIVE, LARGE DATA VOLUME APPLICATIONS.

"IT'S A FAIRLY COMPLEX SYSTEM, BUT THANKS TO SYBASE, IT'S RUNNING SMOOTHLY. FOR EXAMPLE, IN ADDITION TO THE 2.5 TERABYTES OF DATA STORED IN OUR CENTRAL COMPLEX, WE'VE GOT ABOUT 1,600 TABLES IN OUR DATA MODEL, ABOUT 6,000 STORED PROCEDURES AND 4,000 TRIGGERS. THE SYSTEM MANAGES INVENTORIES OF APPROXIMATELY ONE BILLION BUSHELS ANNUALLY WITH AS MANY AS 16,000 DAILY SHIPMENTS. ON A TYPICAL DAY, WE'RE MOVING A GIGABYTE OF DATA BACK AND FORTH ACROSS OUR ENTERPRISE AND ARE RUNNING ABOUT 8,500 REPORTS AN HOUR. WE SIMPLY COULDN'T DO THIS EFFECTIVELY WITHOUT SYBASE."

—DON NIELSEN, MANAGER OF THE DATABASE ADMINISTRATOR GROUP FOR THE CARGILL GRAIN DIVISION

DATABASE REPLICATION SOLUTIONS

Disaster Recovery and High-Availability

An increasing number of organizations are running their operations around the clock, around the globe. For these highly sensitive entities, downtime can have a devastating effect in any of the following ways:

- Lost revenue
- Damaged brand
- Relationship issues with partners and customers
- Supply chain problems
- Legal repercussions

Rather than being lumped into one overarching category, it's best to think of downtime as falling into one of three major categories:

- **Managed, planned downtime.** This type of downtime is common when upgrading hardware, software, networks, and so on. Even though it's managed, however, it still has the potential to impede operations.
- **Unplanned downtime.** This type of downtime frustrates users the most and causes the largest amounts of disruption. Unplanned downtime most commonly results from software bugs, system failures, network outages, and other problems.
- **Disasters.** While relatively rare, these outages have the potential to cause significant, long-term damage to an organization.

NO MATTER WHAT KIND OF DOWNTIME YOU'RE FACING, DATABASE REPLICATION HAS THE POTENTIAL TO DRAMATICALLY REDUCE, OR EVEN ELIMINATE, SYSTEM OUTAGES.

Depending on the type of enterprise, its budget, and its exposure to system outages, database replication is typically used in concert with a collection of redundant hardware, software, and network infrastructure.

By engaging in a well thought out strategy of database replication, you can go a long way toward increasing the reliability of your information architecture, while greatly reducing the likelihood of a major system or data outage.

A DATA INTEGRATION & DISASTER RECOVERY SOLUTION FOR SACO

SACO, Saudi Arabia's largest superstore and an authorized dealer for Ace Hardware, needed to find a faster way to replicate data among its 15 retail stores, headquarters and warehouse locations. The company also needed to deploy a more reliable disaster recovery plan. SACO found the answers to both challenges in Replication Server, which copies data in real-time from headquarters and the warehouse to the retail stores providing up-to-the-minute inventory and pricing information. Replication Server also copies data immediately to on-site and off-site data recovery servers.

"THIS DISASTER RECOVERY DESIGN HAS MADE OUR RECOVERY PROCESS A LOT EASIER, WHICH GIVES US PEACE-OF-MIND. WE DON'T USE ANY THIRD-PARTY TOOLS—JUST REPLICATION SERVER WITH ITS HOT-STANDBY REPLICATION TECHNOLOGY."

—RAED AYYOUB, IT DIRECTOR FOR SACO

Introducing Evergreen Motors

Throughout this white paper, we'll use the example of Evergreen Motors, or simply, "Evergreen," to illustrate the major concepts in an easy-to-understand way.

Evergreen Motors' was founded by a group of environmentally focused partners. Its mission is to offer consumers hybrid vehicles for rent at major airports worldwide. The company started small, but in the past two years, business has begun to grow at an accelerated rate.

Unfortunately, not everything has gone smoothly. A series of system slowdowns, outages and other IT-related mishaps have led to unhappy customers and embarrassing press reports.

To get matters back under control, Evergreen has decided to employ database replication technology. We'll describe three scenarios that benefit from this new infrastructure.

Figure 3-1 shows the initial configuration of reservation agents, an application server, and a database server.

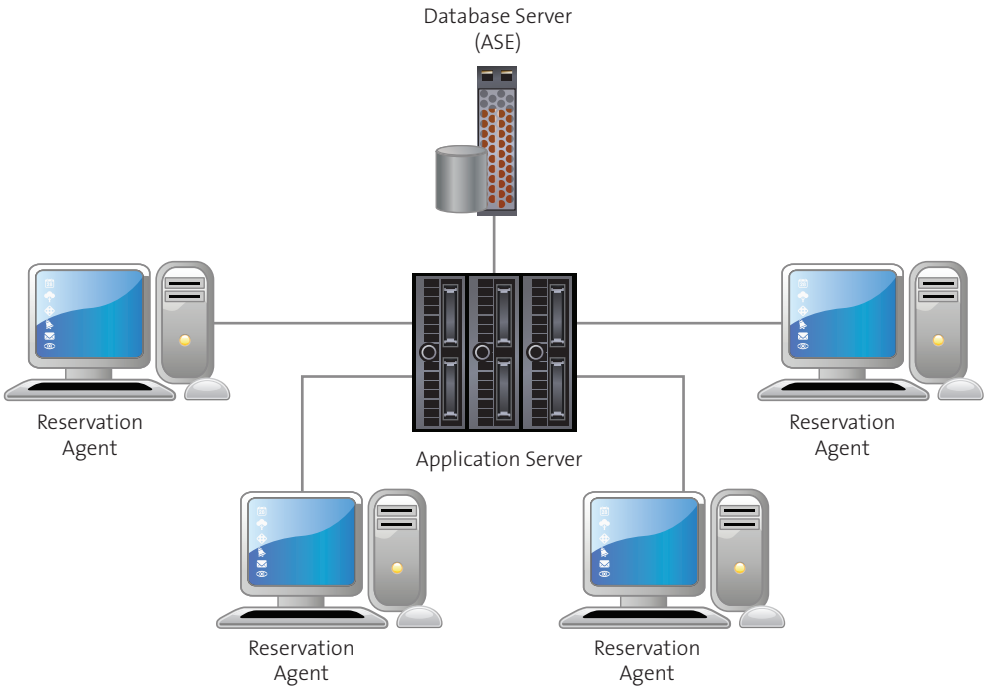

Figure 3-1: The Evergreen Motors initial configuration.

Disaster recovery at Evergreen

Each night, Evergreen Auto administrators back up the production database. However, after reading reports about a competitor that suffered a disastrous fire that destroyed their data center (including backup tapes), Evergreen decides to use database replication to perform real-time copying of information to a backup server located in a separate data center in a different state. And going forward, Evergreen plans on implementing a fully replicated environment, including duplicate application and database servers. Database replication is a big part of this strategy.

Figure 3-2 illustrates this configuration:

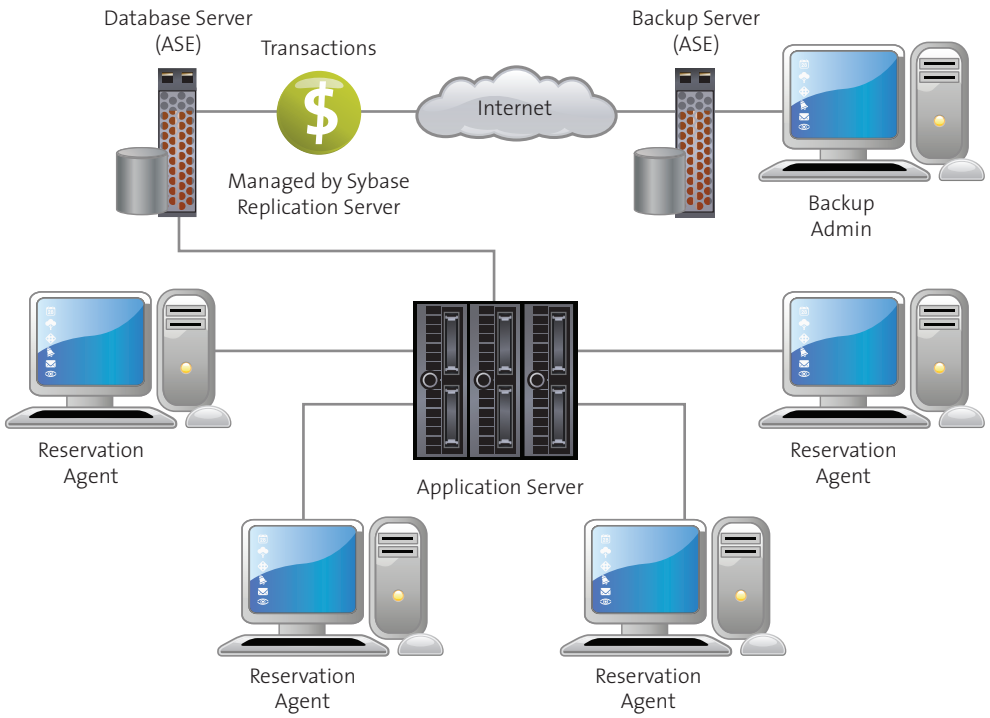

Figure 3-2: Evergreen prepares for disaster recovery.

Real-Time Reporting

User appetite for information is insatiable. Creating, viewing, and modifying real-time reports are increasingly popular ways to work with data. Modern information technology makes these types of queries feasible. High-speed CPUs, large capacity disk drives, fast networks, and powerful business intelligence software have opened up new horizons for users to model and manage their data.

However, all these capabilities have introduced significant risk to production applications. Many IT organizations are learning the hard way what happens when you place real-time analysis and reporting tools in the hands of the average information consumer, and then turn them loose against production servers. The inevitable result is degraded application performance, with all sorts of subsequent negative side effects. Database replication is an ideal solution for these challenges.

When you picture data-replication-driven applications for reporting, imagine two classes of reporting solutions, both of which are greatly facilitated by database replication.

Identical copy of production database

In this scenario, IT administrators elect to use database replication to maintain an exact replica of live, online production data. Users can then construct on-the-fly queries and reports with minimal IT involvement. All the live data that could possibly be available is present on the replicated servers, in near real-time fashion.

A DISASTER RECOVERY & DATA INTEGRATION SOLUTION FOR SACO

SACO, Saudi Arabia's largest superstore and an authorized dealer for Ace Hardware, needed to find a faster way to replicate data among its 15 retail stores, headquarters and warehouse locations. The company also needed to deploy a more reliable disaster recovery plan. SACO found the answers to both challenges in Replication Server, which copies data in real-time from headquarters and the warehouse to the retail stores providing up-to-the-minute inventory and pricing information. Replication Server also copies data immediately to on-site and off-site data recovery servers.

"THIS DISASTER RECOVERY DESIGN HAS MADE OUR RECOVERY PROCESS A LOT EASIER, WHICH GIVES US PEACE-OF-MIND. WE DON'T USE ANY THIRD-PARTY TOOLS – JUST REPLICATION SERVER WITH ITS HOT-STANDBY REPLICATION TECHNOLOGY."

—RAED AYYOUB, IT DIRECTOR FOR SACO

Data warehousing

An organization may elect to selectively consolidate large quantities of data into one or more centralized data warehouses. Users are then supplied with business intelligence software and encouraged to analyze this subset of the organization's information using new or existing queries and reports. This data may also be updated by users and fed back into the live production system.

Regardless of the selected reporting approach, the overall effects of these techniques are a reduced workload for the production application and related server(s).

Evergreen's approach to reporting

As Evergreen's business continues to grow, the organization hires a collection of well-paid marketing and financial analysts to help optimize revenue and profit. Armed with spreadsheets and business intelligence tools, these analysts run a series of "what-if" scenarios.

Although interpreting these simulations is always somewhat subjective, one thing is clear: Running these queries against the production database server has caused dramatic slowdowns for reservation agents (and customers, too).

Because these slowdowns begin to affect revenue, Evergreen decides to add a dedicated server for the marketing and financial analysts, as shown in Figure 3-3. Database replication keeps the reporting server up-to-date, which allows the analysts to run all the queries and simulations they like without impacting production users.

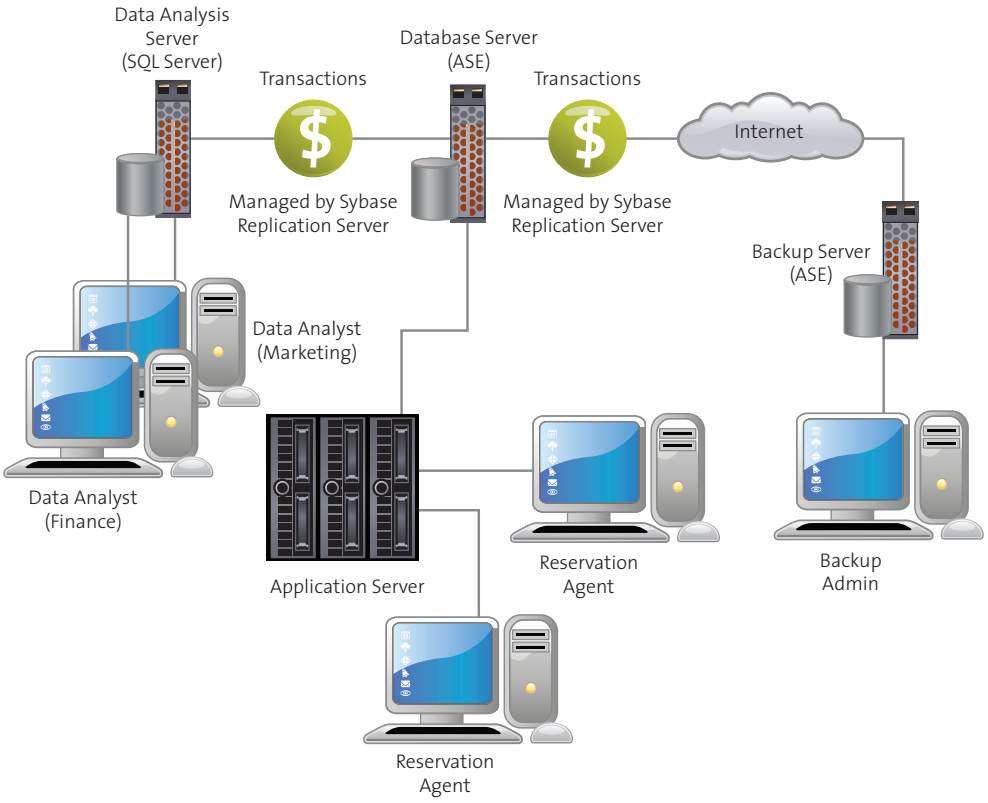

Figure 3-3: Evergreen's dedicated server for marketing and financial analysts.

SPICE TELECOM FOR SYBASE REPLICATION SERVER OPTION FOR ORACLE

Profit margins are thin in the highly competitive telecommunications industry and even small degrees of efficiency improvements add up and fund investments in services and basic infrastructure.

Sanjay Srivastava, General Manager – IT, Spice Telecom, explains part of the need for analytics:

"THE TELECOMMUNICATIONS INDUSTRY IS ONE OF THE MOST CHALLENGING AND DATA-INTENSIVE MARKETPLACES IN EXISTENCE. COMPETITION IS RELENTLESS, AND WE HAVE TO CONSTANTLY MAKE ON-THE-SPOT DECISIONS BASED ON THE INFORMATION WE HAVE AVAILABLE. WE NEED TO BE READY TO COUNTER OUR COMPETITION'S OFFERS AND BE ABLE TO PROVIDE MORE OPTIONS WITHOUT ENDANGERING PROFITABILITY AND KEEP OUR COSTS TO A MINIMUM. WITH MORE DATA WE CAN MODEL DIFFERENT SCENARIOS."

Spice uses Sybase Replication Server Option for Oracle along with Replication Server to immediately capture data changes in the Oracle system and move them into Sybase IQ. Replication Server's flexibility also lends itself to a number of other uses, including extending the life of the Online Transaction Processing (OLTP) database.

Data Consolidation and Synchronization

In the minds of many IT professionals, data consolidation and synchronization (whereby information from multiple locations is coalesced and distributed in an orderly manner) represent the most common scenario where replication is appropriate. However, several classes of replication lie beneath this umbrella term. Here are some of the most popular subcategories:

- **Data synchronization.** Data is distributed among multiple computers, and each computer maintains an up-to-date copy of the information. Generally, every server is permitted to make changes to its information; these alterations are then propagated to all other servers transparently.
- **Data integration.** Information from multiple sources is combined and then distributed onto a collection of destination computers.
- **Data consolidation.** A collection of remote, diverse data is extracted and merged onto one or more replication servers. In many situations, this approach facilitates data warehousing or other business intelligence usage.
- **Data migration.** In this case, database replication moves information from legacy sources onto new destination servers. The alternative to replication in this context is to unload the data from the source and load it onto the destination.

Employing more than one of these types of replication at the same time is a common practice. For example, you might use data integration to tie disparate applications together, and then turn to data synchronization to distribute this information onto several replication destinations.

Selecting a configuration

You must consider two primary classes of database replication when consolidating or synchronizing information. In these classes, information has either

- A master location along with a collection of subordinate destinations.
- No particular master location but is instead distributed onto a collection of equal peers, all of which are free to make alterations.

THE DATABASE REPLICATION TECHNOLOGY YOU SELECT MUST BE ABLE TO MANAGE DATA CHANGES FROM ANY LOCATION. THIS FUNCTIONALITY IS ESSENTIAL, EVEN IF YOU DON'T INITIALLY PLAN FOR THESE TYPES OF MODIFICATIONS; THE EVENTUAL USAGE OF DATABASE REPLICATION TYPICALLY FAR EXCEEDS THE ORIGINAL PLANS.

Choosing what to replicate

To support these data distribution scenarios, replication management software gives administrators a choice of what should be replicated. Administrators may decide to replicate specific columns, rows, or both, depending on the replicated needs of their users. Business-driven rules may also come into play.

In the past few years, the rise of low-cost, commoditized hardware has had significant repercussions throughout the entire technology industry. In particular, data consolidation and synchronization are especially well-suited for leveraging low-cost hardware. This has lowered the price barriers for many new applications that are based on the availability of database replication.

Some of the most sophisticated database replication technologies can now support information hosted by multiple database vendors. This is particularly important given the fact that most IT organizations struggle under the weight of disparate platforms acquired for many different reasons. A database replication solution that does not address this reality is insufficient for the majority of modern enterprises.

Evergreen consolidates data

Evergreen continues to grow. In fact, business is going so well that Evergreen acquires a competitor ("Recycled Wheels," with a primary presence in Europe). Their new acquisition uses a series of completely different systems, running on different database technology. Evergreen is faced with the challenge of integrating these two disparate environments without disrupting either business.

Database replication can help address Evergreen's needs. Figure 3-4 shows a configuration in which Evergreen Motors and Recycled Wheels have a live, bi-directional database replication between their primary servers.

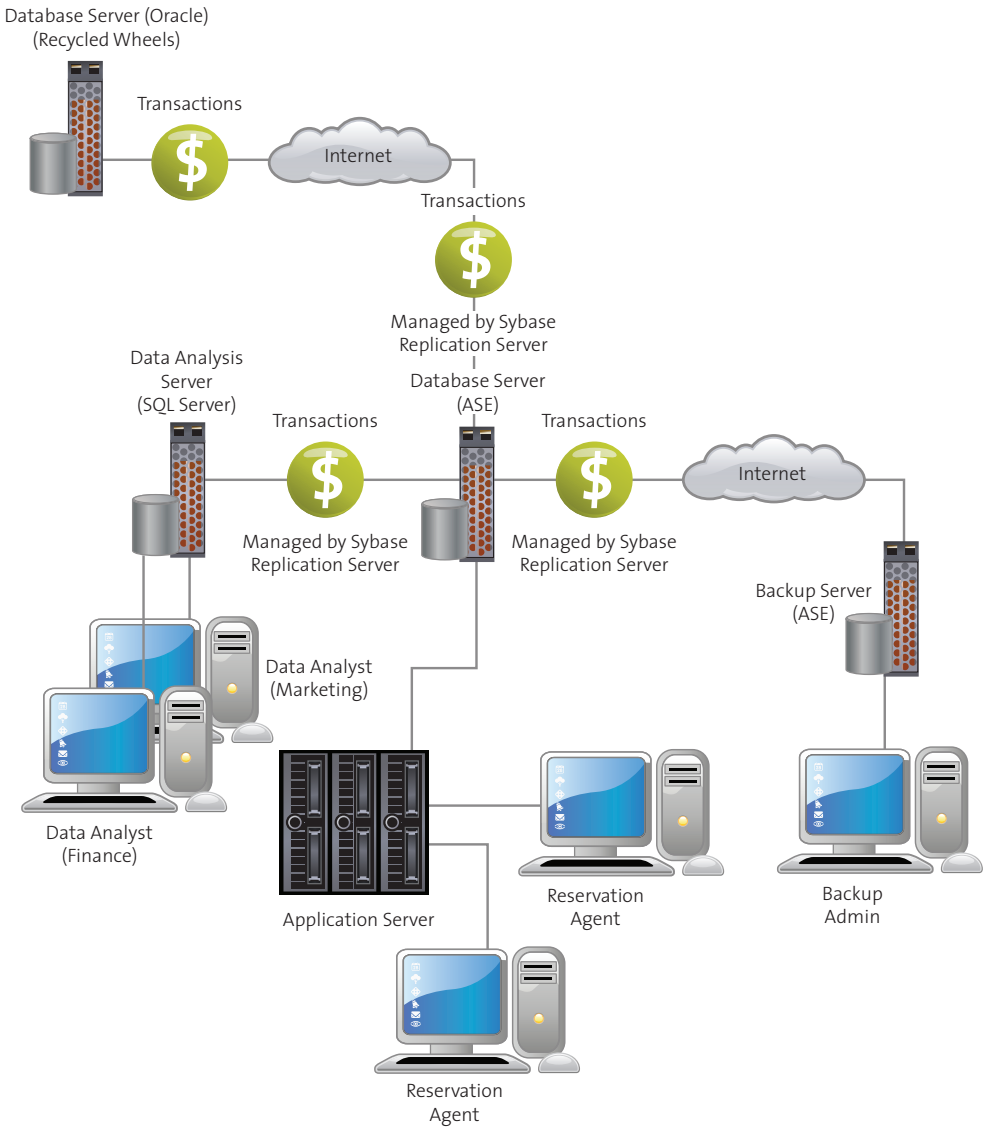

Figure 3-4: Live, bi-directional database replication.

THE COMMONWEALTH OF PENNSYLVANIA COMMON PLEAS CASE MANAGEMENT SYSTEM

The Commonwealth of Pennsylvania Common Pleas Case Management System (CPCMS) is a comprehensive, statewide system for criminal case management developed by the Administrative Office of Pennsylvania Courts (AOPC) under a mandate from the Pennsylvania Supreme Court. The system, built using Sybase Technology, brings the 60 Pennsylvania Judicial Districts in 67 counties together in a single, uniform approach to criminal case processing.

"THE ABILITY TO QUICKLY ACCESS THE MOST CURRENT COURT INFORMATION THROUGHOUT THE STATE—INFORMATION INCLUDING PRIOR CONVICTIONS, BAIL HISTORY, PENDING CHARGES AND OUTSTANDING WARRANTS—IS INVALUABLE TO JUDGES, LAW ENFORCEMENT OFFICERS, AND OTHER CRIMINAL JUSTICE AGENCIES."

—RALPH J. CAPPY. CHIEF JUSTICE OF PENNSYLVANIA

DATABASE REPLICATION TECHNOLOGY SELECTION AND IMPLEMENTATION GUIDELINES

The tips here fall into three categories, based on the capabilities of the database replication platform, how you administer it, and how you plan for the future. Sybase Replication Server supports each of these guidelines and recommendations.

Operational Best Practices

When you implement technology in support of a database replication strategy, use the following guidelines to ensure a successful rollout.

- **Support for large volumes of data.** Data volumes are growing by exponential rates, and this expansion shows no sign of slowing down anytime soon. For this reason, the database replication technology you choose must be able to deal with large sets of information in a timely manner without bogging down the data source servers with overhead.
- **High availability.** Because availability and disaster recovery are the most common uses for database replication, it's essential that the selected replication technology itself be able to maintain a very high uptime rate.
- **Heterogeneous data source access.** IT and data administrators typically face a hodgepodge of information sources, all of which need to participate in replication. Thus, any database replication technology that you choose must factor in these disparate data sources and treat them all equally well.
- **Consistent information delivery.** Even the most intelligently designed database replication architecture will be an utter failure if it is not able to deliver consistent information across all replicated nodes. The prospect of making a business decision based on incorrect or out of date or replicated data is not one that most IT leaders would care to consider. For this reason, be sure that your chosen database replication solution has built-in transactional integrity capabilities.
- **High performance.** Because database replication implies additional processes and resource demands, the database replication technology you choose must place a minimal load on your production systems.
- **Local autonomy.** In the past, most database replication topographies sported a master server working together with a number of subordinate destinations. Generally, data was never modified on these secondary locations. That restriction is no longer applicable in many modern computing environments. In fact, many organizations are finding that they no longer run a master-subordinate configuration but instead work in a peer-to-peer configuration.

REGARDLESS OF WHETHER YOU CHOOSE TO USE MASTER/SUBORDINATE OR PEER-TO-PEER, WHATEVER DATABASE REPLICATION TECHNOLOGY YOU CHOOSE SHOULD BE ABLE TO DEAL WITH EITHER OF THESE TWO TYPES OF CONFIGURATIONS IN A MANNER THAT IS TRANSPARENT TO THE USER.

Administration

The most sophisticated database replication solution will be seriously hampered if the solution is difficult to configure or administer. Here are several administrative recommendations:

- **Easy administration from a central location.** Because modern IT organizations are dealing with complex information topographies, the administrator must be issued the proper set of tools to help them make better database replication decisions faster and more accurately.
- **Support for multiple data sources.** Given today's reporting and other requirements, chances are good that your replication technology will need to accommodate several (or possibly many more) data sources, all of which must be factored into your database replication plan and be supported by your chosen technology.
- **Support for multiple targets.** Just as you will commonly need to distribute information from more than one source, it's quite likely that you will need to send this information on to many different in a graceful, efficient manner, ideally in as visually-friendly a way possible.
- **Selective replication.** Your database replication solution should give you the ability to select rows, columns, or other filtering criteria when determining what and how to replicate data.

Extensibility and Upgrading

How database replication might interact with other strategic IT initiatives might not be on your radar screen right now, however falling hardware and bandwidth costs, will increasingly make these technologies more compelling.

- **Message bus support.** Enterprise Service Buses (ESB) offer a collection of well-integrated capabilities that provide reliable messaging, guaranteed delivery, and other middleware features necessary for mission-critical applications. ESBs are available in proprietary and open-source implementations. Many organizations are deploying this technology as the foundation of their Service Oriented Architecture (SOA). Regardless of your chosen database replication technology, it's important that it be capable of integrating with the real-time messaging features found in an ESB. This partnership increases the reliability of your database replication while also boosting performance.
- **Mirroring interoperability.** Applying hardware-based alternatives to database replication, such as standby servers, redundant networks, and disk mirroring, helps ensure high availability for your valuable information.

IDEALLY, YOUR CHOSEN DATABASE REPLICATION SOLUTION WILL EASILY INTEGRATE WITH DISK MIRRORING INFRASTRUCTURE. THIS APPROACH PROVIDES YOU WITH A HOT STANDBY DATABASE THAT IS KEPT IN SYNC WITH YOUR PRIMARY DATABASE SERVER. AS AN ADDED BENEFIT, DURING NORMAL OPERATIONS YOU'RE FREE TO USE THE STANDBY DATABASE FOR REPORTING PURPOSES, THEREBY REMOVING PROCESSING LOAD FROM YOUR PRIMARY SERVER.

CONCLUSION

Today's IT organizations are under increasing amounts of pressure. Globalization, around-the-clock operations, increased regulatory requirements, and higher competitive pressures are translating into relentless demands to trim costs. At the same time, IT organizations struggle to manage diverse vendor platforms that are distributed around the world. Data is growing rapidly, yielding even more serious implications of an outage or other unplanned downtime.

Database replication solutions such as Sybase Replication Server can help organizations deal with these challenges. By distributing data to multiple destinations from multiple sources, IT organizations gain flexibility when meeting user needs. Information consumers can now construct real-time queries and reports that work with live data. The organization is protected from the effects of a system outage, data corruption, or other business-threatening event.

Sybase, Inc.
Worldwide Headquarters
One Sybase Drive
Dublin, CA 94568-7902
U.S.A
1 800 8 sybase

www.sybase.com

SYBASE® An SAP Company