

Compliments of SYBASE®

# Replication

FOR  
DUMMIES®

Sybase Edition

**A Reference  
for the  
Rest of Us!**®

FREE eTips at [dummies.com](http://dummies.com)®

Get to know how  
database replication  
can work for you!



Robert D. Schneider



***Replication***  
FOR  
**DUMMIES®**  
SYBASE EDITION

**by Robert Schneider**



John Wiley & Sons Canada, Ltd

These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

## Replication For Dummies®

Published by  
John Wiley & Sons Canada, Ltd.  
6045 Freemont Blvd.  
Mississauga, ON  
L5R 4J3  
www.wiley.com

Copyright © 2010 by John Wiley & Sons Canada, Ltd. All rights reserved. No part of this book, including interior design, and icons, may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording, or otherwise) without the prior written permission of the publisher.

Trademarks: Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

For authorization to photocopy items for corporate, personal, or educational use, please contact in writing The Canadian Copyright Licensing Agency (Access Copyright). For an Access Copyright license, visit [www.accesscopyright.ca](http://www.accesscopyright.ca) or call toll free 1-800-893-5777.

**LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.**

For general information on John Wiley & Sons Canada, Ltd., including all books published by Wiley Publishing Inc., please call our distribution centre at 1-800-567-4797. For reseller information, including discounts and premium sales, please call our sales department at 416-646-7992. For press review copies, author interviews, or other publicity information, please contact our publicity department, Tel. 416-646-4582, Fax 416-236-4448.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

ISBN: 978-0-470-73980-8

Printed in Canada

1 2 3 4 5 PC 14 13 12 11 10



These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

## *About the Author*

**Robert Schneider** is a Silicon Valley-based technology consultant. He has written five books and numerous articles on advanced technical topics such as Service Oriented Architecture (SOA), open source, and relational database design/optimization. Most recently, he's authored *Microsoft SQL Server 2005 Express Edition For Dummies* and *Microsoft SQL Server 2008 All-in-One Desk Reference For Dummies*. He can be reached at [Robert.Schneider@think88.com](mailto:Robert.Schneider@think88.com).

## **Publisher's Acknowledgments**

We're proud of this book; please send us your comments through our online registration form located at <http://dummies.custhelp.com>. For other comments, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

Some of the people who helped bring this book to market include the following:

### ***Acquisitions and Editorial***

#### **Vice-President Publishing**

**Services:** Karen Bryan

**Editor:** Robert Hickey

**Project Manager:** Alison MacLean

**Senior Project Editor:**

Elizabeth McCurdy

**US Project Editor:** Kelly Ewing

**Copy Editor:** Laura Miller

**Editorial Assistant:**

Katie Wolsley

**Cartoons:** Rich Tennant

([www.the5thwave.com](http://www.the5thwave.com))

### ***Composition Services***

**Project Coordinator:**

Lynsey Stanford

**Layout and Graphics:**

Melissa K. Jester

**Proofreader:** Dwight Ramsey

---

### ***John Wiley & Sons Canada, Ltd.***

**Bill Zerter**, Chief Operating Officer

**Jennifer Smith**, Vice-President and Publisher,  
Professional & Trade Division

### ***Publishing and Editorial for Consumer Dummies***

**Diane Graves Steele**, Vice President and Publisher,  
Consumer Dummies

**Kristin Ferguson-Wagstaffe**, Product Development Director,  
Consumer Dummies

**Ensley Eikenburg**, Associate Publisher, Travel

**Kelly Regan**, Editorial Director, Travel

### ***Composition Services***

**Debbie Stailey**, Director of Composition Services

These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

# Contents at a Glance

---

<b>Introduction .....</b>	<b>1</b>
Who This Book Is For .....	1
About This Book .....	2
Icons Used in This Book.....	3
<b>Chapter 1: Today's IT Realities.....</b>	<b>5</b>
Appreciating the Complexity of Today's IT Landscape .....	6
Considering Business Realities .....	6
<b>Chapter 2: Database Replication to the Rescue .....</b>	<b>9</b>
Understanding Database Replication.....	9
Sybase Replication Server .....	11
<b>Chapter 3: Three Database Replication Solutions .....</b>	<b>15</b>
Disaster Recovery and High Availability .....	15
Introducing Evergreen Motors .....	17
Disaster recovery at Evergreen .....	18
Real-Time Reporting.....	20
Identical copy of production database.....	21
Data warehousing.....	21
Evergreen's approach to reporting.....	22

Data Consolidation and Synchronization .....	24
Selecting a configuration .....	26
Choosing what to replicate .....	26
Evergreen consolidates data.....	27

## **Chapter 4: Database Replication Technology**

### **Selection and Implementation Guidelines ..... 31**

Operational Best Practices .....	31
Administration .....	33
Extensibility and Upgrading .....	34
Letter from Sybase.....	37



# Introduction



**D**atabase replication, which is the distribution of information across a series of database servers that may be located in the same place or scattered across the globe, was once reserved for large, sophisticated organizations sporting the big IT budgets. Thankfully, for a variety of reasons, the past few years have seen wider application of database replication solutions.

This book 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.

## *Who This Book Is For*

This book is aimed at two communities. The first group of readers is IT and business leadership. Therefore, I steer away from detailed technical discussions of database replication and focus on the bigger picture, especially regarding how database replication can be used to solve the business and technological challenges that most organizations face today.

The second community encompasses database administrators who are frequently on the front lines of implementing database replication-driven solutions. While I don't have space in this book to make comprehensive recommendations, the configuration examples I use can serve as a helpful starting point for this group of readers; the collection of best practices also provide a solid foundation.

## *About This Book*

In order to clearly and concisely explore database replication, this book is divided into four chapters.

In Chapter 1, I look at several trends and factors (from both a technical and business perspective) that place stress on IT organizations. We then move on to examine how database replication can help make these challenges more bearable. While replication isn't a universal solution, it can add real value and rapid return on investment (ROI).

In Chapter 2, my goal is to help you get a better understanding of database replication. I supply a high-level overview, as well as cite several techniques used to implement database replication.

Because database replication is a multi-faceted, highly customizable approach toward improving IT operations, Chapter 3 highlights several of the most popular replication use cases. I also provide customer examples to illustrate in a concrete way how database replication can benefit an organization.

In Chapter 4, I suggest a series of best practices to follow when implementing database replication in your organization. Note that these guidelines are vendor-neutral: they're good practice no matter what technology you use.

## Icons Used in This Book

This book uses the following icons to indicate information that you may find especially helpful as you familiarize yourself with database replication technology.



These helpful hints point out ways you can get the most from database replication technology, perhaps even saving time and money as you do.



You don't want to forget this important information.



Watch out for these common pitfalls and you'll be cruising your way to success in no time.



You may find this information interesting, but you can skip it, if you like. You can always return to read it later.



# Chapter 1

## Today's IT Realities

- ▶ Understanding the complexities of running IT
- ▶ Accepting the business realities

**W**hen you consider the complexity of an average IT organization, you'll find that contrasting it with organizations from a few years back is worthwhile. In the past, for example, the following was true:

- ✔ 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. This chapter describes the technological challenges IT faces, as well as the business context within which it deals with those challenges.

## ***Appreciating the Complexity of Today's IT Landscape***

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 the technological complications described in the preceding section, 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 will be 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.





## Chapter 2

# Database Replication to the Rescue

- ▶ Knowing the techniques for database replication
- ▶ Understanding Sybase Replication Server

**I**ncreasing 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 (see Chapter 1). 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.

This chapter describes how database replication works and some of the more popular techniques for performing it.

## *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. The next section covers this approach in detail.

## *Sybase Replication Server*

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 employs a distributed architecture, with a collection of capabilities to guarantee delivery of data alterations.



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.** As we describe earlier, 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.

### **Testimonial — Cargill Grain Division**

As the following customer quote illustrates, Replication Server facilitates today'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

## Chapter 3

# Three Database Replication Solutions

- ▶ Exploring disaster recovery and high availability
- ▶ Knowing how to implement real-time reporting
- ▶ Consolidating and synchronizing data

**T**his chapter focuses on three scenarios where database replication can add significant value to every IT organization.

### *Disaster Recovery and High Availability*

An increasing number of organizations are running their operations around the clock, around the globe (see Chapter 1). 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.

## ***Introducing Evergreen Motors***

Throughout this chapter, I 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. This chapter describes 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.



For space and clarity considerations, these examples are as concise as possible. In addition, each example builds on the previous scenario. Finally, each scenario closes with a concrete example of how Sybase Replication Server is being used to solve similar challenges for real customers.



**Figure 3-1:** The Evergreen Motors initial configuration.

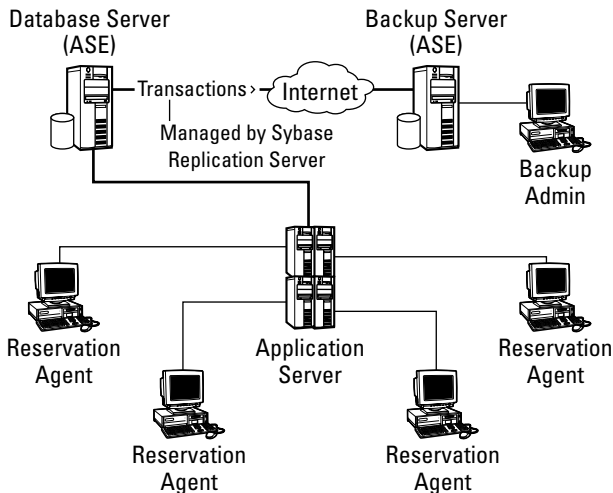
---

### *Disaster recovery at Evergreen*

Each night, Evergreen Motors 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. Figure 3-2 illustrates this configuration:

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:** Evergreen prepares for disaster recovery.

## **SACO for Replication Server**

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

## ***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.

### ***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).

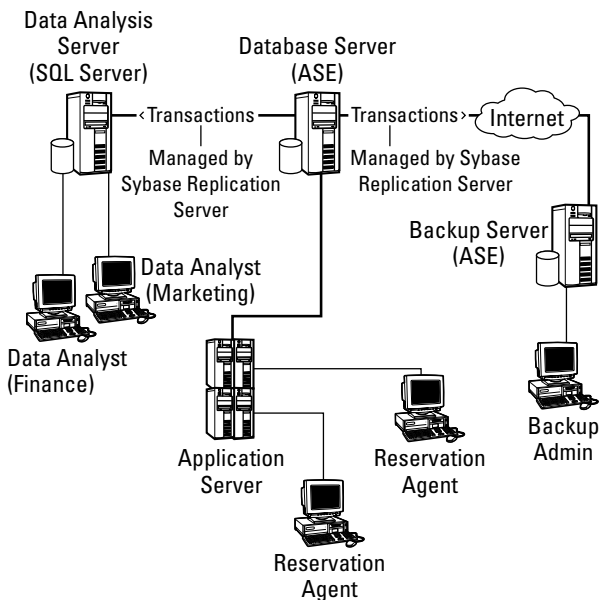
### ***Evergreen's approach to reporting***

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.

Database Server (Oracle)  
(Recycled Wheels)



Transactions



Transactions

Managed by Sybase  
Replication Server

Data Analysis  
Server  
(SQL Server)



Transactions

Managed by  
Sybase Replication  
Server

Database Server  
(ASE)



Transactions

Managed by Sybase  
Replication Server




Backup Server  
(ASE)



Data Analyst  
(Marketing)

Data Analyst  
(Finance)



Application  
Server



Reservation Agent



Reservation  
Agent



Backup  
Admin

**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 state-wide 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



## Chapter 4

---

# Database Replication Technology Selection and Implementation Guidelines

- 
- ▶ Getting a handle on operational best practices
  - ▶ Knowing the administration guidelines
  - ▶ Grasping extensibility & upgrading
- 

**T**he tips in this chapter 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 the guidelines and recommendations we list in this chapter.

### *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 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 administrators 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 destinations — 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*

In this section, we delve into several important factors that you should bear in mind regarding how database replication might interact with other strategic IT initiatives.



While these topics might not be on your radar screen right now, you might find that the IT and business trends we described earlier, combined with falling hardware and bandwidth costs, will make these technologies more compelling sooner than you think.

- ✓ **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 vital 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.



## *Letter from Sybase*

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.

Please visit our Web site for access to our demos, datasheets, and Webcasts ([www.sybase.com/replication](http://www.sybase.com/replication)), or contact us to set up a call with one of our Replication Server experts. We are happy to answer your questions as you take the next steps in creating your own data management solutions.

Jenn Buchanan and Bill Zhang

Sybase Replication Server Team


[www.sybase.com/replication](http://www.sybase.com/replication)

1-800-8SYBASE



## DATABASES

---



978-0-470-04652-4


### Also available:

- ✓ Data Warehousing For Dummies  
978-0-470-40747-9
- ✓ Database Development For Dummies  
978-0-7645-0752-6

- ✓ Microsoft SQL Server 2008 For Dummies  
978-0-470-22465-6
- ✓ SQL All-in-One Desk Reference For Dummies  
978-0-470-11928-0

## NETWORKING

---



978-0-470-45060-4

### Also available:

- ✓ Networking All-in-One Desk Reference For Dummies  
978-0-470-17915-4
- ✓ RTLS For Dummies  
978-0-470-39868-5

- ✓ VoIP Deployment For Dummies  
978-0-470-38543-2
- ✓ Windows Server 2008 All-in-One Desk Reference For Dummies  
978-0-470-18044-0

Available wherever books are sold. For more information or to order direct: U.S. customers visit [www.dummies.com](http://www.dummies.com) or call 1-877-762-2974. U.K. customers visit [www.wileyeurope.com](http://www.wileyeurope.com) or call 0800 243407. Canadian customers visit [www.wiley.ca](http://www.wiley.ca) or call 1-800-567-4797.



These materials are the copyright of Wiley Publishing, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

## PROGRAMMING

---



978-0-470-04652-4


### Also available:

- ✓ C++ For Dummies  
978-0-470-31726-6
- ✓ JavaScript and Ajax  
For Dummies  
978-0-470-41799-7

- ✓ iPhone Application  
Development  
For Dummies  
978-0-470-48737-2
- ✓ Microsoft SQL Server  
2008 For Dummies  
978-0-470-22465-6

## WEB DEVELOPMENT

---



978-0-470-14928-7

### Also available:

- ✓ Drupal For Dummies  
978-0-470-55611-5
- ✓ HTML, XHTML & CSS  
All-in-One Desk  
Reference  
For Dummies  
978-0-470-18627-5

- ✓ Joomla! For Dummies  
978-0-470-43287-7
- ✓ Microsoft Silverlight  
3.0 For Dummies  
978-0-470-52465-7
- ✓ Wikis For Dummies  
978-0-470-04399-8



# DATA CENTER TO DEVICE AND EVERYWHERE IN BETWEEN

It's your vital information and it's ready to work. In the office. On the road. At a customer site. Even a coffee shop. So manage it. Analyze it. Mobilize it. It's yours to use. Because in the Unwired Enterprise, your data is accessible anytime, anywhere.

Learn more at [www.sybase.com](http://www.sybase.com)

## SYBASE®

**WEST COAST APPLICATIONS ARE ON ORACLE.  
EAST COAST ARE ON SYBASE.  
YOUR REPORT is  
DUE NOW.**

**Replication Server® delivers proven database replication for Sybase, Oracle, Microsoft and IBM**

Globalization, round-the-clock operations, increased regulatory requirements, and higher competitive pressures are translating into relentless demands on IT. At the same time, organizations are struggling to manage diverse vendor platforms that are distributed around the world. Data is growing by leaps and bounds and real-time access is no longer a luxury—it is a critical requirement.

Call **1-800-8SYBASE** to request a free trial of Replication Server or visit [www.sybase.com/replicationfordummies](http://www.sybase.com/replicationfordummies)

**SYBASE®**



Add value to your IT organization  
with database replication

## Understand how database replication can solve your organization's business and technological challenges

Organizations today have a growing need to move and synchronize data across the enterprise in real-time without the risk of system downtime or performance degradation. Database replication is the solution, but up until now, it's been costly and difficult to administer. This friendly guide explains how database replication will add value to your distributed organization in a cost-efficient and simple way.

THE  
DUMMIES  
WAY<sup>™</sup>

Explanations in plain English  
"Get in, get out" information  
Icons and other navigational aids  
Top ten lists  
A dash of humor and fun

ISBN: 978-0-470-73980-8  
Part No.: L03227  
Not for resale

## Discover how to:

*Understand the different techniques for database replication*

*Reduce the likelihood of a major system or data outage*

*Do real-time reporting without degraded application performance*

*Consolidate and synchronize data from multiple data sources through replication*

*Follow best practices to ensure your replication rollout is effective*

## Get smart!

@ [www.dummies.com](http://www.dummies.com)

- ✓ Find listings of all our books
- ✓ Choose from many different subject categories
- ✓ Sign up for eTips at [etips.dummies.com](http://etips.dummies.com)

For Dummies®  
A Branded Imprint of

