BUSINESS PROCESSING GIANT TURNS TO HENSON GROUP TO ENHANCE SQL DATA MANAGEMENT SOLUTION

CYNERGY DATA, LLC

Microsoft SQL Data Replication

"Using replication, you can distribute data to different locations, to remote or mobile users over a local area network, using a dial-up connection, and over the Internet. Replication also allows you to enhance application performance, physically separate data based on how it is used, for example, to separate online transaction processing and decision support systems, or distribute database processing across multiple servers."

— Alfred McBean, Chief Technology Officer, News America Marketing

Microsoft Gold Certified Partner The Henson Group, Inc. (THG) provided expert services for achieving the Microsoft SQL Server Replication that client Cynergy Data required to make it much easier to copy, distribute, and potentially modify data across the entire enterprise.

Overview

Country or Region: United States
Industry: Business Processing

Customer Profile
Cynergy Data provides electronic credit, debit, and check transaction processing and related services to such customers as retailers, restaurants, e-businesses, and catalog firms.

Business Situation
With the introduction of Microsoft Office SharePoint Server 2007 (MOSS 2007) in 2006, TPG-Axon sought to migrate from SharePoint Portal Server 2003 to leverage the latest features of SharePoint, while also seeking custom functionality.

Solution
Cynergy Data was interested in achieving comprehensive, two-way, real-time replication of an SQL database running a custom-banking application.

Results & Benefits
The Henson Group exceeded Cynergy Data's expectations. The replication achieved allows multiple sites to keep copies of the same data. This is useful when multiple sites need to read the same data or need separate servers for reporting applications.

BUSINESS SITUATION

Founded in 1995 by entrepreneurs John Martillo and Marcelo Paladini, Cynergy Data provides electronic credit, debit, and check transaction processing and related services to such customers as retailers, restaurants, e-businesses, and catalog firms.

Today, Cynergy Data is an industry leader in transaction processing with more than 90 employees and a constantly expanding client base does. Cynergy aims not just to acquire new clients, but to provide them with the tools and support to expand their business as well.

With two offices in New York, Cynergy Data currently processes more than $1.6 billion annually, in credit, debit, check conversion and numerous other forms of electronic payment transactions. Considering the nature of its business, Cynergy Data obviously places a premium on effective data management technology.

(For more information regarding this THG client, please visit: http://www.cynergydata.com/about.)

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com

Recently, Cynergy Data began exploring with THG the benefits of investing in a comprehensive, two-way, real-time replication of an SQL 2000 database running a custom-banking application.

“Replication is the process of sharing data between databases in different locations,” Henson says. “Using replication, you create copies of the Database and share the copy with different users so that they can make changes to their local copy of the database and later synchronize the changes to the source database.”

The Henson Group offers a dedicated Data Management Practice specializing in both small- and large-scale Microsoft SQL Server deployments, providing a comprehensive array of services including logical and physical site architecture, replication, migration, reporting, analysis, security, and much more.

A **Microsoft 2007 Partner of the Year**, The Henson Group is a "Gold Certified Partner" founded by former Microsoft engineers in 2002, specializing in enterprise deployments and .NET development of Microsoft products that overcome today's business challenges.

The preferred solution provider for many US and international corporations, The Henson Group leverages direct ties to the Microsoft product groups, is a preferred sub-contractor of Microsoft Consulting Services, holds a seat on the national Microsoft Partner Advisory Council, and is recognized among the top US consultancies in Microsoft's partner community (**Microsoft’s "Solution Finder"**).

SOLUTION

Replication architecture is fairly extensive, designed to ensure that the architecture is versatile enough to meet the needs of just about any replication situation. However, this versatility also makes replication difficult to configure. To make the replication go smoothly requires effective planning, which involves selecting a specific replication model and performing any necessary preliminary tasks before replication configuration commences.

“The primary objective of replication is data distribution,” Henson explains. “Placing multiple copies of the same data at different locations enables you to load balance or reduce network traffic generated by remote database clients, merge data from various sources into a comprehensive target, or separate databases based on their purpose.”

Copying critical data to a secure location also serves frequently as the basis for disaster recovery or high availability solutions.

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com

The Henson Group performed detailed discovery to understand Cynergy Data's network architecture, server implementation, client operating systems, and specific SQL Server 2000 replication needs and objectives. Once discovery was conducted, a detailed document was prepared revealing all the aspects of Cynergy Data's SQL environment. This included the following:

1) Geographic Profile: documented the geographic locations of the organization's sites, including information about operating system configurations and time zone considerations.

2) Organization's Structure: documented the divisions or departments within Cynergy Data's organization and their associated managing and reporting structures.

3) Network topology: documented Cynergy Data's network infrastructure, including LAN and WAN architectures, physical topology, network size, bandwidth, usage, traffic patterns, network protocols, and subnets.

4) Client Environment: ascertained the number of clients at each location, software applications and operating systems in use, including logon scripts, client mobility and type of network connectivity (dial-up, wireless, LAN, etc.).

5) Domain Model: documented Cynergy Data's forest, domains, and Active Directory sites in the client's Active Directory site structure, if applicable, and organizational units.

6) Server Environment: documented the locations of the core servers on Cynergy Data's network, indicating their primary functionality and operating system version level.

7) Data Replication Objectives: defined the desired results of replication, including the amount of data, type of data (snapshot, merge, and transactional), and performance goals.

8) IT Organization: acquired knowledge of Cynergy Data's IT organization, including the support areas defined for IT staff.

Cynergy provided The Henson Group with the following information during discovery:

- There were 2 data centers located approximately 40 miles apart
- There was reliable T2 network connection between the sites

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com

- Each site hosted a SQL Server 2000 database located on equally configured Windows 2000 servers with 2 CPU and 2 GB of RAM
- The size of database is about 25 GB, and it is slow growing (200-500 MB per month)
- Both sites could equally manipulate data, including Update/Delete the same record
- Conflicts between entered information on both sides were possible, but do not happen often
- There were 4 types of tables in a database:
    - Lookup tables (infrequently updated by developers on either sides; conflicts can be avoided by using only one site to make changes)
    - Tables with historical information (regularly loaded during batch process on either sides; conflicts can be avoided by using only one site to upload data)
    - Processing tables (users can do everything on either site; conflicts are possible)
    - Old tables that can be removed from the database
- There was relatively low load (10-50 transactions per min). Application mainly is in use from 7am to 9pm EST
- There was a remote possibility of adding the third processing site
- Cynergy's DBA did not have practical experience of supporting SQL Server replication

Based on preliminary discussions, Cynergy reported the following requirements:

- Configure run-time replication between both sides
- Changes made on either site must be replicated to another site with minimum delay
- Minimize impact of hardware and (or) software failure on one site to the operations on another site
- If possible, consolidate the changes made by the users on different sites in different columns of the same row
- Create flexible replication topology that can add additional site with minimum efforts
- Make configuration and support of replication as simple as possible

Based on information gathered during the Discovery phase, THG provided Cynergy Data with a document that divulges the logical and physical replication strategies that will enable THG to achieve Cynergy Data's expectations. Based on THG's proprietary process for project management, this project plan was structured according to detailed deployment steps that were anchored with specific milestones and delivery dates. Reporting the successful achievement of each milestone

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com

ensured Cynergy Data's IT professional are kept apprised of major developments.

THG sought Architecture Validation from key professionals inside Microsoft's SQL Server 2000 product group, with specific knowledge of replication, to review the project architecture and plan. This ensures that the architecture and features deployed are compatible with future versions of the product, accommodating a migration strategy when, and if, Cynergy Data decides to implement such a migration plan.

Based on information gathered during the discovery phase, THG will execute the replication process implementing snapshot, merge, or transactional replication between the two designated office locations.

This was followed by Distributor Configuration and Replication. As the name indicates, distributors are used to distribute replicated data. THG configured the the core set of distributor administration tasks including setting up a new distributor, updating an existing distributor, and deleting distributors.

Next, THG creates the distribution databases to be used to store the information being distributed to subscribers. Each publisher that uses a distributor is assigned a distribution database to which it can connect. Publishers can share distribution databases, and create additional databases as necessary.

Distributors can work only with servers and databases that are enabled for their use. THG therefore enabled publishers for the replication process.

Once THG configured a distributor and enabled publishers, publication databases, and subscribers, it worked with Cynergy Data to create publications. Afterward, Cynergy Data need only to manage the publications as it would any other SQL Server resource.

THG then worked with Cynergy Data to leverage Data Transformation Services (DTS) to manipulate publication data before sending it to subscribers. This technique is useful if individual subscribers need different views of the data or if there is a need to map information from one table or data format to another table or data format.

Once the solution was deployed successfully, THG and Cynergy Data met to review all aspects of the deployment. As part of this process, THG documented all of the concerns of Cynergy Data's project leaders, including requests for modifications and/or programming enhancements that were subsequently addressed.

THG also demonstrated to Cynergy how to monitor SQL Server 2000 replication two ways: Use replication monitors at the server level or use replication monitor groups at the

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com

enterprise level. Once monitors are enabled, Cynergy Data can use the monitor to view the status of replication agents and troubleshoot potential problems with distributors. All replication agents have options that you can set to customize their behavior.

RESULTS & BENEFITS

At the outset of the engagement, The Henson Group discussed three replication types:

> Snapshot Replication: With snapshot replication, the data being replicated is copied in full to data files on the distributor. Normally these snapshot files are the same size as the data you're replicating and are stored in the SQL Server Repldata folder by default.
>
> Transaction Replication: Because transactional replication builds on the snapshot replication model, you'll want to prepare for both snapshot and transactional replication. With transactional replication, an initial snapshot is sent to the distributor and this snapshot is then updated on a periodic basis, such as once a week. In between snapshots, transactions are used to update subscribers. These transactions are logged in the distributor's database and are cleared out only after a new snapshot is created.
>
> Merge Replication: With merge replication, all published tables must have primary keys. If a table contains foreign keys or is used in validation, you must include the reference table in the publication. Otherwise update operations that add new rows will fail because SQL Server can't find the required primary key. Additionally, merge replication affects timestamp column usage. Timestamps are generated automatically and are guaranteed to be unique only in a specific database. Like transactional replication, merge replication has a few limitations when it comes to text and image columns.

Upon further review, The Henson Group endorsed implementing Merge replication as the main replication topology, based on various considerations.

“Merge replication has the most robust conflict resolution mechanism which can be easily configured using GUI,” Henson says. “In the simplest case no scripting will be involved at all.”

Merge replication provides out-the-box mechanism of resolving conflicts on column level. Due to the fact that conflicts were rare at Cynergy and that only two sites were involved in data processing, THG determined that transactional consistency of the operations should be very high.

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com

Merge replication allows replicating text data. Currently there are more than 10 tables in a database with text columns. Continuously running Merge Agents provides data transmission with acceptable latency.

Cynergy also expressed willingness to accept adding a unique identifier column to each replicated table as required by design of Merge replication. *(It will be easy to add another server to the existing topology at later time, if necessary.)*

Several publications can be created to speed up the process of delivering data to the opposite site. Note, all tables linked with Foreign Key constraint must be published together. Therefore, THG recommended creating at least three publications:

a. Major processing tables and linked lookup tables
b. All historical tables
c. Remaining lookup tables

The Henson Group exceeded Cynergy Data's expectations. The replication achieved allows multiple sites to keep copies of the same data. This is useful when multiple sites need to read the same data or need separate servers for reporting applications.

Greater autonomy is achieved, as users can work with copies of data while disconnected and then propagate changes they make to other databases when they are connected.

There is now increased scalability of data, with scale out of data to be browsed, such as browsing data using Web-based applications.

The solution also brings data closer to individuals or groups, helping reduce conflicts based on multiple user data modifications and queries because data can be distributed throughout the network, and you can partition data based on the needs of different business units or users.

Database replication now also augments Cynergy's disaster-recovery plans by duplicating the data from a local database server to a remote database server. If the primary server fails, cynergy can switch to the replicated copy of the data and continue operations.

"As part of our proprietary project framework, we also provided the extensive documentation and on-site training and knowledge transfer to ensure that Cynergy Data achieved the optimum return on its investment for this initiative," Henson says.

About The Henson Group

A Microsoft 2007 Partner of the Year, The Henson Group is a "Gold Certified Partner" founded by former Microsoft engineers in 2002, specializing in enterprise deployments and .NET development of Microsoft products that overcome today's business challenges.

The preferred solution provider for many US and international corporations, The Henson Group leverages direct ties to the Microsoft product groups, is a preferred sub-contractor of Microsoft Consulting Services, holds a seat on the national Microsoft Partner Advisory Council, and is recognized among the top US consultancies in Microsoft's partner community (Microsoft's "Solution Finder").

About Microsoft

Founded in 1975, Microsoft (NASDAQ: MSFT) is the worldwide leader in software for personal and business computing. The company offers a wide range of products and services designed to empower people through great software any time, any place and on any device.

Document published April 2008.

The Henson Group, Inc. 1375 Broadway, Third Floor New York N.Y. 10018 USA
www.HensonGroup.com 800.980.1130 Info@HensonGroup.com