Two Merged Retailers Integrate, HP NonStop SQL and Oracle RAC

Paul J. Holenstein
Executive Vice President
Gravic, Inc.

Company mergers always present a host of challenges, and prominent among them is integrating the disparate IT systems of the merged parties. Two major retailers that joined forces faced this exact predicament. Each retailer had thousands of stores, and the merged retailer wanted to allow a shopper to return or to exchange goods at any of its stores regardless of the store where the original purchase took place. This challenge required transforming HP NonStop SQL (HP NSSQL) applications and Linux-based Oracle RAC applications so that the disparate databases could be integrated into a common framework.

The two retailers turned to Shadowbase Streams from Gravic, Inc., to integrate the distinctly different applications by mirroring the HP NSSQL and Oracle databases using bidirectional, active/active data replication. The goal was for each store to have local access to all sales, refunds, and exchanges so that any customer could easily be serviced at any store regardless of the retailer.

The Retailers Evolve with HP NonStop BASE24™ Systems

The two retailers in this case study have each been in operation for over a century, and each has significant brand recognition. Their wares have very little overlap, so the merger of the two chains made sense in terms of cross-selling between the brands and increased market competitiveness in an ever-changing retail world. Following the merger, they continued to operate their respective stores under their original brand names, which we will call Brand 1 and Brand 2.

Over the last several decades, the two retailers had independently evolved their IT operations to suit their own needs. When credit cards came into vogue, they each installed their own point-of-sale (POS) terminals in their stores. Over time, both retailers adopted ACI Worldwide, Inc.'s BASE24™ payment processing engine (www.aciworldwide.com) to route credit-card and debit-card transactions from their POS terminals to issuing banks for authentication and authorization, as shown in Figure 1. Each BASE24 system was heavily customized to meet the needs of its retailer.

Both retailers used HP NonStop systems to run their versions of BASE24. BASE24 controls all of the POS terminals in a retailer's stores, sends card transactions to the issuing banks for authorization, and records all sales.

One of the critical BASE24 facilities that each retailer used was the refund/exchange module. Since all of each retailer's stores had access to the retailer's BASE24 system, customers did not have to return items to the specific stores in which the products were purchased. Rather, customers could come into any of the retailer's existing stores with their purchase receipts and purchased products and obtain a refund. Alternatively, customers could exchange the products for other ones.

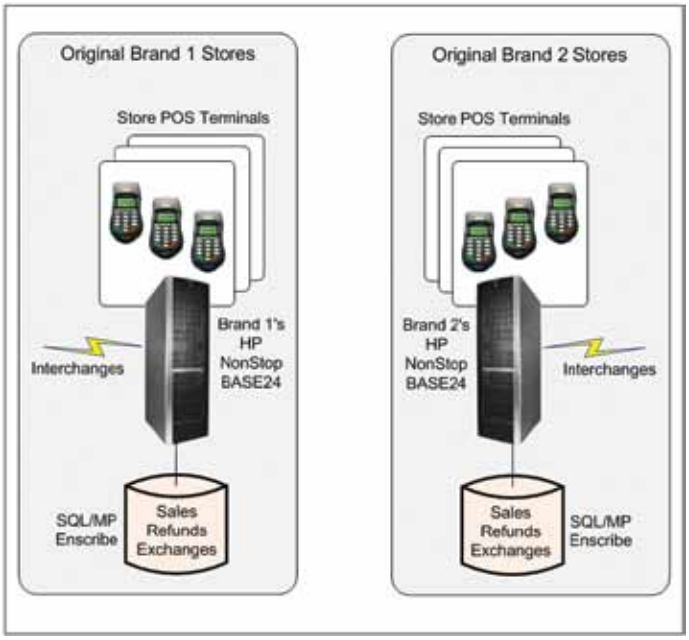

Figure 1: The Original Payment Processing Systems

New Stores Open with an Oracle RAC Database

Following the merger, the retailer wanted to integrate its systems and extend its refund/exchange policy so that a customer could return or exchange a product at a store of either Brand 1 or Brand 2, regardless of where the product had been purchased.

To complicate matters, the retailer implemented a new financial-payment engine with additional functionality for the many new stores that it opened under either of the two brand names or for existing stores that it decided to refurbish/retrofit. It continued to use the HP NonStop BASE24 systems for its other existing stores, so its current POS systems did not need to be replaced in those stores. However, new stores and refurbished existing stores used the POS systems provided by the new financial-payment system vendor.

The new vendor was AJB Software Design, Inc. (www.ajbsoftware.com). AJB's Retail Transaction Switch (RTS) performs real-time authorization for a variety of transaction types from numerous sources. RTS provides the authorization functions for POS transactions with the issuing banks. AJB's Flexible Integrated Payment System (FIPay) is a store-level communication gateway that interfaces a retailer's POS devices to RTS.

RTS runs on an Oracle Real Application Cluster (Oracle RAC) multi-server cluster. Oracle RAC provides the transparent deployment of a single database across a cluster of commodity servers. Applications can run simultaneously on all servers in the cluster, accessing the database via a common, distributed cache, as shown in Figure 2. Should one server fail, the processing load is assumed by the surviving servers. Thus, an Oracle RAC cluster provides application fault tolerance – the application will continue

to function in the presence of any single (or perhaps multiple) server failures.

Oracle RAC clusters are complex and are typically much more difficult to set up and to manage than HP NonStop servers. Failover to the surviving servers following a server failure can take minutes compared to the sub-second recovery time of HP NonStop systems. However, the retailer wanted to take advantage of additional functionality provided by the AJB system.

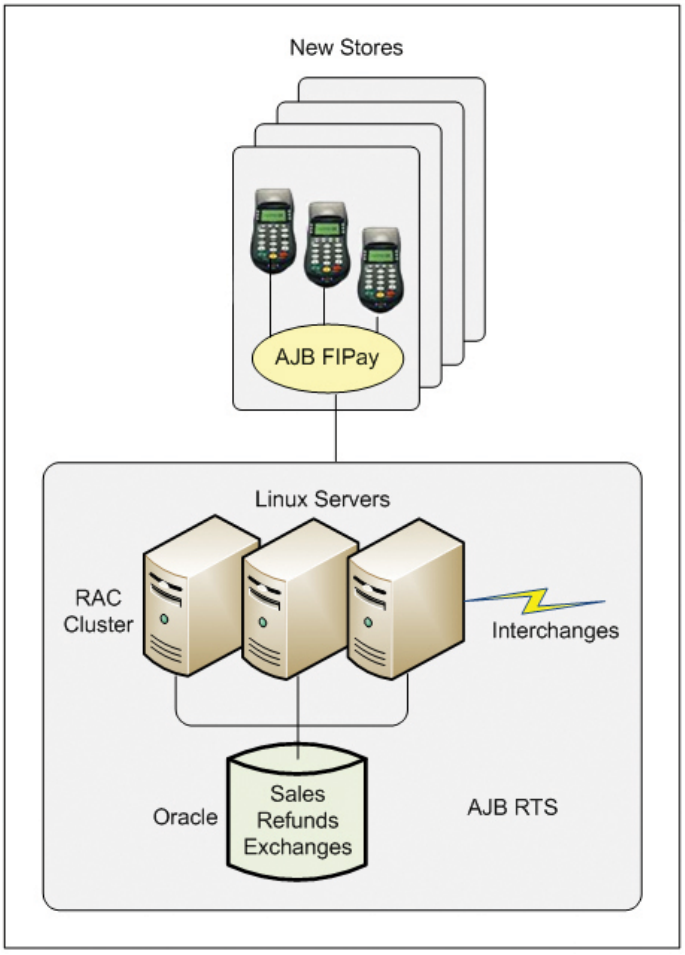

Figure 2: The New Payment Processing System

Integrating the BASE24 HP NSSQL and Oracle RAC Databases

The Retailer's Integration Strategy

The retailer next faced the challenge of having three different payment processing systems – an individually customized HP NonStop BASE24 system for each of the original retailers and an AJB Oracle RAC system for new stores being opened or for existing stores being retrofitted. Each system managed sales, returns, and exchanges for the subset of stores that it served. The retailer wanted to extend the return and exchange functions so that any store, existing or new, of either brand could accept product returns and exchanges from any of the other stores.

The retailer considered different options. One was to provide a new system that could act as a central repository of all sales transactions – in effect, an operational data store (ODS). Returns and exchanges would be managed by this new central system. However, the acquisition and management of yet another system as well as the integration of existing systems with it via messaging adapters was not a cost-effective move for the retailer.

Another option was to integrate the databases of the three systems so that they presented a single-system data view containing all sales, refund and exchange transactions across all stores. In this way, a customer could show up at any store with a purchased product for return or exchange and would be properly serviced.

The retailer chose the second option. The strategy used a bidirectional data replication engine replicating all of the sales, refund, and exchange transactions to each of the databases. The replication engine would transparently handle transforming the data from the source schema/data types to that of each of the destination databases, including identifying and resolving any data collisions that might occur.

To simplify the transition, the retailer initially provided this capability primarily for its new stores. Any product purchased at a new store could be returned to or exchanged at any store, existing or new, of either brand. Likewise, a product purchased at any existing store could be returned to or exchanged at any existing store of the same brand or at any new store of either brand. The only limitation was that products purchased at existing stores could not be returned or exchanged at existing stores of the other brand. (As the existing stores are retrofitted with the new AJB POS terminals, this restriction will disappear.)

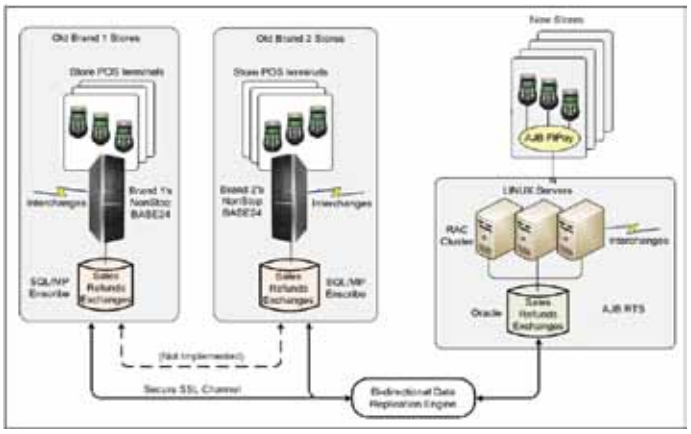

Figure 3: Integrated Refunds and Exchanges

Sales, return or exchange transactions were not replicated between the two BASE24 systems, leading to the restriction that product purchases at an existing store of one brand could not be returned to an existing store of the other brand. This limitation could have been avoided if the retailer had included bidirectional replication between the two original BASE24 databases, but it decided to defer this option to a later date.

Data Protection

The retailer faced a major problem dealing with the data-protection requirement imposed by the Payment Card Industry Data Security Standard (PCIDSS) because the BASE24 systems and the AJB system address this requirement differently.

According to the standard, sensitive data items such as credit-card numbers (primary-account numbers, or PANs) and social

security numbers must be protected from unauthorized access. The BASE24 version used by the retailer did not employ encryption in its database (however, newer versions of BASE24 do support this capability). In contrast, the AJB RTS encrypted credit-card numbers in-place and in-flight.

Powerful user exits provided by the bidirectional replication engine integrated the two protection technologies, as shown in Figure 3. Whenever a data item is sent from a BASE24 database to an AJB database, it is read in plain text from the BASE24 database and is sent over a secure, encrypted SSL channel to the AJB system. There, using an API furnished by the AJB system, the replication engine encrypts the data before storing it in the AJB database.

Conversely, when replicating from AJB to BASE24, the replication engine reads the encrypted data item from the AJB database and uses an AJB API to decrypt the data item. The data item is then sent to each of the BASE24 systems over a secure, encrypted SSL channel and is written in plain text to the BASE24 databases.

Data Collisions

One important challenge facing bidirectional replication is that of data collisions. A data collision occurs if the same data item is updated on two different systems almost simultaneously – within the replication-latency interval during which the changes are propagating between the two systems. If this collision should happen, and if no provision is made to handle it, the change made at one system will overwrite the change made at the other system and vice versa. Now both systems will have different values for that data item, and both will be wrong.

In the return/exchange application of the retailer, collisions are highly unlikely since the customer is supposed to present the original sales receipt, and there is only one of them. However, during batch processing, it is possible that a batch update will make a change to a data item in one system at the same time that the returns/exchange application is doing so in another system.

The bidirectional replication engine chosen by the retailer provides data-collision detection and correction. In this case, the replication engine chooses a winning update based on rules established by the retailer. For instance, the rule may be that the more recent timestamp wins (such timing rules generally require that the clocks of the systems be synchronized). In the unlikely event of a data collision, the replication engine resolves the collision by selecting a winner and reports on the collision loser for later manual review and reconciliation.

Shadowbase Streams Chosen for the Bidirectional Replication Engine

After exhaustive evaluations, the retailer chose Shadowbase Streams from Gravic, Inc., (www.gravic.com/shadowbase) to keep the databases synchronized. Shadowbase Streams provides bidirectional data replication between heterogeneous databases.

As a heterogeneous replication engine, Shadowbase Streams uses Change Data Capture (CDC) technology to automatically map database updates from one schema to another and to replicate between databases from different vendors (Figure 4). This capability includes field-format conversion, aggregation, filtering, data cleansing, and schema translation.

Shadowbase Streams also includes powerful user exits that execute bound-in procedures for providing customized function extensions for the replication engine, such as the need to support the AJB encryption facility when dealing with the non-encrypted BASE24 database contents.

Figure 4: The Shadowbase Streams Data Replication Engine

Whenever a sales, return, or exchange transaction is executed on an existing store's BASE24 system, it is replicated by Shadowbase Streams to the AJB Oracle database serving the new stores. Likewise, a transaction executed at a new store is replicated from the AJB Oracle database to both BASE24 SQL/MP databases serving the existing stores. In this way, purchases made at any of the retailer's stores, regardless of the brand, may be refunded or exchanged at any of its other stores.

Summary

Following a major merger of two retailers, the combined company faced the challenge of integrating the IT operations of the two original retailers and adding new functionalities. These companies used different applications running on different platforms to perform the same functional job. Disparate databases containing the same classes of data needed to be integrated into a common database in order to provide better customer service. This integration task was daunting, especially when the integration needed to occur with little or no application downtime.

In this case study, the IT integration problem was solved more easily than by replacing multiple systems with a single integrated system. The merged retailer elected to keep the original payment systems of each of the original retailers, despite their differences. The merged retailer went even further and elected to upgrade to a new payment system for new stores that it was opening.

The databases of the three payment systems are kept synchronized via bidirectional data replication provided by Gravic's Shadowbase Streams. In this way, any system can properly process any transaction, regardless of where it originated. In particular, the retailer is able to extend its customer service by allowing a customer to return or exchange a product at any store regardless of the store brand purchased (with some restrictions).

Integrating applications and databases via data replication is a powerful technique for managing the differences between diverse IT infrastructures and for extending their functionalities. Shadowbase Streams successfully integrated the distinctly different applications of this retailer by mirroring the HP NSSQL and Oracle databases using bidirectional, active/active data replication.

Paul J. Holenstein is Executive Vice President of Gravic, Inc. He is responsible for the Shadowbase® suite of products. The Shadowbase replication engine is a high-speed, unidirectional and bidirectional, homogeneous and heterogeneous data-replication engine that moves data updates between enterprise systems in fractions of a second. It also provides capabilities to integrate disparate operational application information into real-time business intelligence systems. Shadowbase Total Replication Solutions® provides products to leverage this technology with proven implementations..