A Bank's Crisis Migration to a New Data Replication Solution

Paul J. Holenstein
Executive Vice President
Gravic, Inc

During a system upgrade project, a major bank found itself squeezed between either paying a large increase in license fees for its current data replication engine or having those licenses terminated. With only two weeks to go, the bank turned to Gravic, Inc., for help. The Gravic technical team configured, installed, and tested its Shadowbase data replication product in time to replace the bank's existing replicator before the licenses expired. The bank is now extending its use of the Shadowbase product suite to satisfy all of its data replication needs.

The Bank's Heterogeneous Online Banking Systems

RAK and BASE24™

The bank is recognized as one of the most important global systemic banks, one whose operations are a major underpinning to the world's financial community. It serves 50 million clients in 40 countries. The bank has been a user of HP NonStop systems for decades, dating back to the Tandem days before Tandem Computer's ultimate acquisition by HP. Figure 1 depicts the bank's online banking applications as well as the bank's ACI BASE24 environment for managing its ATMs.

The Real-time Authorization Kernel (RAK) is a home-grown application that furnishes online customer services including account-balance queries, fund transfers between internal accounts, and fund transfers between customer accounts and external accounts. RAK also provides online authorization services for the bank's credit cards and debit cards. The RAK database is primarily HP's SQL/MP.

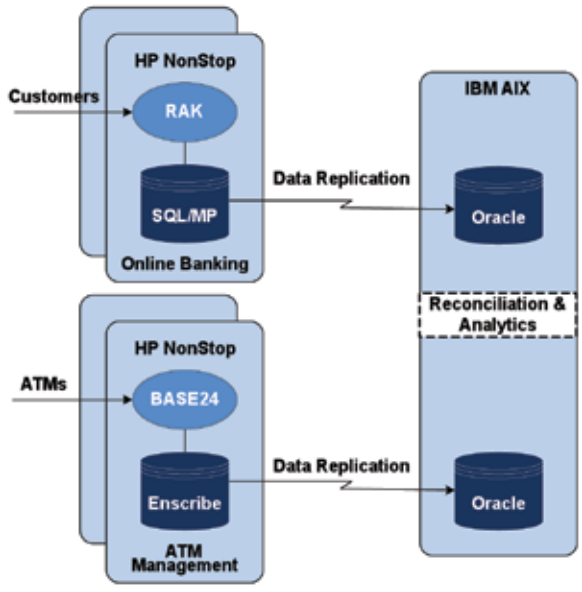

Figure 1 – The Bank's NonStop Systems

The ACI BASE24 Classic system administers the bank's ATMs. It receives and manages the authorization of ATM withdrawals by sending transactions to the banks issuing the cards that are used at the ATMs. The BASE24 environment is primarily an Enscribe environment and employs HP's AutoTMF to audit and protect the Enscribe data files. All ATM transactions are recorded in an Enscribe log file.

RAK and BASE24 run on their respective HP NonStop systems, which are configured as active/passive pairs for business continuity purposes. One system in each pair is the production system that performs all of the processing, while the other serves as its backup. The database of each backup is kept synchronized with its production counterpart via data replication. (Replication is not shown in Figure 1.) In this way, the backup system is available to take over processing if the production system fails.

The AIX/Oracle Reconciliation System

The bank's RAK and BASE24 systems interoperate with an IBM/AIX Unix system that uses an Oracle Real Application Clusters (RAC) database. Online banking transactions and payment card transactions must be sent from the NonStop RAK system to the AIX system for reconciliation. Likewise, completed ATM transactions must be sent from the BASE24 system to the AIX system for analytical processing. The AIX system supports fraud detection and in-depth business analysis and intelligence as well as many other offline functions.

The Bank's Use of Data Replication

The bank uses data replication for several purposes. HP NonStop RDF replicates changes in an active/passive architecture from a NonStop production database to its backup database. RDF is utilized both by the RAK system and the BASE24 system to keep their backup systems synchronized with their production systems. Data replication is also employed to replicate data from the RAK and BASE24 systems to the AIX/Oracle environment. RAK uses SQL/MP tables, and BASE24 uses Enscribe files. Changes to these databases must be replicated in real time to the Oracle relational databases on the AIX system.

Thus, data replication is highly heterogeneous. The source databases and the target databases are from different vendors. With respect to BASE24, replication must occur from nonrelational Enscribe files to relational SQL tables. Data must be cleansed, filtered, validated, and transformed as it is being replicated. Data aggregation is also necessary when data from multiple source databases is replicated to a single target database, which requires the combining of fields and columns from differing files and tables into a single target row. Likewise, data deaggregation is used to send data changes from a single data source to multiple target

databases. Significant data normalization is employed to redefine data formats between the source and target databases and to convert the use of arrays and redefines between the databases. For example, for the BASE24 system, the primary task is to convert and replicate the variable length and format of Enscribe ATM transaction log records to the table schemas used in the AIX/Oracle analytics and reconciliation database.

The transaction and I/O rates of the SQL/MP tables (RAK) and the Enscribe files (BASE24) are quite high, and they can spike to several times the normal load during peak periods (for example, holiday season). The replication engine must be able to handle these high data replication volumes and to scale as the bank's business grows.

The Licensing Crisis

The bank planned to upgrade its NonStop systems to the new NB54000 NonStop BladeSystems. The plan was for RAK to run on a pair of eight-CPU NonStop NB54000 BladeSystems and for BASE24 to run on a pair of ten-CPU NonStop NB54000 BladeSystems. All systems would be dual-core, though the NB54000 NonStop servers easily could be upgraded to quad-core without requiring an application outage. Because of the migration to the new version of NonStop servers, the bank had to obtain updated licenses for its replication products.

For many years, the bank had been using a third-party data replication engine to replicate data between its RAK and BASE24 NonStop systems and its AIX/Oracle system. Although the bank periodically had issues with the third-party's offshore support organization, the bank had no immediate intention of moving off its current replication engine. The bank originally expected to migrate its existing licenses to the new NonStop hardware when the upgrade occurred. As it turned out, the renegotiation process stalled, and significant licensing issues arose as time was running out.

In the end, the bank was able to negotiate a one-year extension of its BASE24 to AIX/Oracle replication licenses under its previously similar terms. Unfortunately, the bank discovered that the data replication vendor required a substantial increase in its license fees for the data replication engine needed for the RAK system. The fees were partly based upon the use of quad-core NB54000 blades; there was no price break for the bank using dual-core blades. The bank considered the new license fee proposal to be cost-prohibitive and rejected it.

The Bank's Options

The bank was faced with limited options to continue the mandatory operation of its RAK system, and it did not have much time to spare. The last-ditch alternative was to renew the RAK license with the existing data replication vendor for the prohibitive license fee.

To avoid this unacceptable option, the bank initiated an intense development effort to build its own file-transfer facility so that RAK periodically could refresh the AIX/Oracle database with new data changes made to the RAK SQL/MP database. However, the transformations required to map the SQL/MP source database to the Oracle target database introduced significant complexity, and the batch nature of the data-refresh process meant that the target environment would be working on stale data most of the time.

With time running out, the bank turned to HP for help. HP's response was to bring in Gravic and its Shadowbase data replication product. The Shadowbase replication engine supports SQL/MP, Enscribe, and Oracle (among many other databases) and comes with a broad range of data cleansing, filtering, and transformation functions. In addition, Shadowbase user exits allow the rapid creation of custom transformations that are not already in the Shadowbase repertoire.

The Shadowbase license fees were well within budget for the bank. The bank decided to give Gravic the go-ahead to install the Shadowbase replication engine, provided that the installation and testing could be completed before the current RAK licenses expired.

The Race Against Time

By this point, only two weeks remained until RAK license termination. Gravic assembled a team of its experienced software engineers and began the installation effort. In order to configure the Shadowbase transformation facilities, Gravic had to know the transformations that were needed, requiring close coordination with the bank's technical staff. Major challenges quickly emerged:

- a bank requirement that all testing be physically performed at the bank's central European facilities;
- coordinating access to the key bank personnel, who were heavily involved in their own day-to-day responsibilities;
- and implementing and testing the myriad functions that perform the actual data transformations during replication of events from the NonStop server to the Oracle target environment.

Starting with the bank's development environments, the needed functions were quickly implemented and were tested with customer test data. The effort then moved into the bank's User Acceptance Testing (UAT) facility. However, this environment was significantly scaled back from the full-blown production environment, meaning that data loading and testing could be simulated but could not be completely performed until production roll-out. The team validated the Shadowbase implementation in the UAT environment by running it in parallel with the existing data replication solution. The target database tables between the two systems were compared to verify that they did indeed match and were processing the source data in the same way for the same types of application events.

Once the Shadowbase UAT environment was validated, the bank scheduled the production roll-out. A Friday afternoon was selected to allow sufficient time for monitoring the new solution over the weekend, during which periodic full-daily processing cycles were performed, including load scale up/down functions. By the end of the weekend, all functions had been confirmed, and full production processing continued the following week.

The result included long hours by the bank and Gravic staff to obtain the information they needed, to configure the Shadowbase replication engine to meet the replication requirements, and to thoroughly test and deploy the Shadowbase solution. These steps demanded considerable onsite effort from the Gravic team.

With great relief on the part of the bank, the Shadowbase effort was successful. In just two weeks, the Shadowbase replication engine was installed and was working in production, replicating data from the RAK system to the AIX system. The bank avoided having to purchase an expensive license for the upgrade and began its efforts to consolidate and base its replication solutions on the Shadowbase product suite.

Lesson Learned

Of course, performing a migration from one product to another can be a risky endeavor even under the best of circumstances. Typically, projects such as these should be undertaken when there is sufficient time to fully plan the effort, fully test the new solution, and then fully deploy the replacement solution on your schedule and not an artificial one imposed by a nearly impossible-to-meet license expiration deadline. Unfortunately, not leaving enough time or allocating sufficient resources to the replacement project is an all-too-often barrier to success, subsequently forcing the customer to continue along with what it has done before, working under less than desirable circumstances. The obvious lesson here is to start the planning process as early as possible with sufficient management support to see it through to the end.

The Next Steps

BASE24 Replication

The bank was still left with the one-year data replication engine license for its BASE24 system. Should the bank extend that license or switch to the Shadowbase replication engine to replicate data from the bank's BASE24 system to its AIX/Oracle system? It made sense to have only one replication engine product to maintain. Besides, the Shadowbase license fee cost was substantially less than the existing data replication license fee. The bank therefore decided to switch to the Shadowbase replication engine for BASE24 data replication. With only three months to go on the existing data replication license, the bank authorized Gravic to proceed with configuring the Shadowbase replication engine for the BASE24 Enscribe-to-Oracle replication task.

Though more time was allocated than for the original RAK installation, a new challenge arose. The data structures for the ACI BASE24 Enscribe files required considerable scrubbing and cleansing to transform the data into the required target SQL formats. The Gravic team once again worked diligently with the bank staff to implement the conversion functions, to test the new solution, and to deploy it into production before the existing licenses expired. The bank was now completely off of the previous data replication vendor's solutions and successfully onto Shadowbase technology.

Disaster Recovery Replication

The bank still uses an active/passive architecture for its disaster recovery processing. This architecture actively runs the application on one node, while the other node sits idle receiving the database changes. If a failover needs to occur, the database on the standby node must be brought into a consistent state, the application on the standby node must be started, and the network must be rerouted so that user requests can be sent to the standby node's applications. In addition, the replication engine must be reconfigured to reverse replicate new database changes to the failed node to eventually recover it.

All of this effort takes time and can be risky if one or more of the failover sequences faults. How can that occur? It turns out that failover faults, where the failover process does not go according to plan and an extended outage occurs, can happen much more frequently than expected, especially if the standby environment is not thoroughly, successfully, and periodically tested. Since testing often has to take down the production application environment, this function is usually slated for off-hours and infrequent time frames, which leads to incomplete testing when the failover does not complete within a preapproved outage window. Without complete testing, how can the configuration of the backup system be ensured to remain identical to that of the production system? Otherwise, the failover may fail. Configuration drift is a leading cause of failover faults, in which changes made to the production system fail to be made to the backup system.

The way to improve on this model and to improve the bank's overall application availability profile is to look to the more advanced business continuity architectures, including the Shadowbase Sizzling Hot Takeover (SZT) architecture and the Shadowbase active/active architecture. In an SZT architecture, the application is up and running on both nodes, although only one node is typically receiving database change requests. (The other node can be receiving and processing read-only/reporting or query requests.) The application on the "standby" node has the data files and tables open for read/write access and has made all external connections. The data replication engine is configured for bi-directional replication.

The benefit of this architecture is that the application is fully running on both nodes at all times. If a failover occurs, no delay is needed to bring the database into consistency, nor to bring the standby application up. Additionally, the standby application is in a known-working state as it is already running. A best practice is to send periodic test transactions to it against test accounts. These test transactions will ensure that the application on the standby node is functional for end-to-end processing. Hence, no production application outage needs to occur to test the standby node's application processing; and the testing can be continuous, performed at any time of the day or night.

With bi-directional replication configured, the reverse replication path also validates that it is functioning. If a failover occurs, no change to the replication environment is needed, and the backup system will start to queue the database changes for the reverse replication to resynchronize the original production database once that node is recovered.

Once the active/passive architecture has been replaced with an SZT architecture, a final step will be for the bank to migrate from the SZT configuration to an active/active system, one in which both nodes share the transaction load. Each system replicates its database changes to the other database so that the applications on both systems have the same view of the application state. Failover is rapid, measured in seconds, and is reliable since it is known that both systems are working properly. Both are processing transactions. Furthermore, when a failure occurs, fewer users and data are affected, as only those users connected to the failed node actually have to fail over.

Summary

The bank was caught off guard by a large increase in license fees for its RAK replication engine. With little time to act, it had to develop multiple contingency plans to continue in operation. These plans included relicensing the current replication engine at a significant increase in license-fee cost, building its own replication facility, or moving to another replication engine.

To avoid the substantial increase in license fees, the bank first decided to build its own replication utility as a fallback plan. This option used a micro-batch refresh approach to periodically load the source database changes into the target database on a set schedule. Unfortunately, choosing this option meant that the data in the target was immediately stale after each cycle, and the application Service Level Agreements (SLAs) required current data at all times. Clearly, a real-time data replication solution was needed.

Hence, the bank initiated an aggressive plan to migrate to another data replication engine. Fortunately, this effort succeeded. With only two weeks to act, the Shadowbase team of software engineers configured the Shadowbase replication engine to properly transform and replicate RAK SQL/MP data to the AIX/Oracle system. With an intensive effort constrained severely by time, Gravic was able to help the bank avoid the costly relicensing of its previous data replication engine.

The bank is in the process of deploying Shadowbase replication solutions for its other data replication needs. Shadowbase software now performs the BASE24 to Oracle replication function, and the bank is investigating the enhancement of its business continuity solutions to a Shadowbase SZT model as an interim step to ultimately achieving an active/active implementation.

The Shadowbase Data Replication Engine

The Shadowbase data replication engine provides homogeneous and heterogeneous data replication between diverse databases and applications. Shadowbase data replication can take place between any supported source database and any supported target database. Either database may be a relational database or a non-relational database.

Shadowbase business continuity solutions span the active/passive architecture to the Sizzling-Hot-Takeover architecture, to a fully active/active architecture. Whereas these solutions help eliminate *unplanned* application downtime, the Shadowbase Zero Downtime Migration (ZDM) solution eliminates *planned* downtime for complex system, site, database, and application upgrades and conversions.

In addition, Shadowbase solutions provide data integration and synchronization, as well as application integration. In these cases, data changes typically need to be replicated from one environment to another, for example to feed operational database changes into a data warehouse. Similarly, using Shadowbase technology, real-time business intelligence systems can be built by combining the output of one application with the input of another application, for example feeding a real-time fraud detection system with transactional activity flowing across a financial message switch and returning the results to flag suspicious activity.

The Shadowbase data replication engine includes powerful transformation facilities that map data between the source database structures and the target database or target application structures. Shadowbase user exits allow special transformation customization functions to be embedded into the replication engine for transformations that are not directly supported.

Attributes of Shadowbase data replication are low latency, high capacity, heterogeneity, powerful data transformations, flexible end points, and continuous availability. Integrating heterogeneous data resources is a formidable challenge, a challenge that is solved by Shadowbase software.[1] These Shadowbase solutions are available from HP under the *HP Shadowbase* product name.

[1] For more information, visit Gravic's website, www.gravic.com/shadowbase/whitepapers, to see the white papers: *Shadowbase® Streams for Data Integration and Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.

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. For further information regarding Shadowbase data integration and application integration capabilities that can assist in solving big data integration problems, please refer to the companion documents Shadowbase Streams for Data Integration and Shadowbase Streams for Application Integration, or visit www.Gravic.com/Shadowbase for more information. To contact the author, please email: SBProductManagement@gravic.com.