Data Replication and Integration Applications – in Financial Services

A Gravic, Inc. Article

Payment Authorization

Data is valuable. And the more current the data is, the more valuable it is. For this reason, the resident data in online transaction processing systems is some of the most valuable. However, if this data is trapped in that system, and not available to enable other real-time business intelligence processes, then its full value is not being exploited, and competitive opportunities are missed. To avoid these missed opportunities, companies need to access their online data in real-time. This issue presents a problem. As the demand for real-time data increases, accessing it while the online system is processing transactions can significantly impact the throughput and response times of that system. This article discusses such a situation, and the clever way it was resolved using data replication.

Merchant Operations

The merchant's business grew and was operating two datacenters, Datacenter North and Datacenter South, which were located 1,000 miles apart on the East Coast of the U.S. The transaction authorization switch functions are duplicated in each datacenter, and processing workload is split between the two sites.

The merchant runs its certified management account (CMA) systems on HPE NonStop servers which maintain a database of all cards issued by participating banks. Banks send batch files to the CMA system for new cards they issue, cards they report as stolen or lost, and other card status changes. These batches are received by the Datacenter North CMA system and stored in its card database.

The Merchant's First Replication Architecture: Disaster Recovery

The merchant's distributed CMA system needed to be synchronized. It also needed to replace its legacy Disaster Recovery (DR) plan for a solution that would provide a much faster recovery time. The merchant performed extensive research and evaluations, and discovered that data replication was the best long-term solution.

When Datacenter North receives its batch changes, they are replicated to Datacenter South by the HPE Shadowbase data replication engine. Therefore, both CMA systems always possess the complete and up-to-date data of all the cards that were issued. This architecture provides continuous processing services even beyond a disaster that destroys one of the datacenters or causes a failure of one of the systems. If one system fails, then all batch updates from the issuing banks are routed to the other CMA system, which is responsible for keeping the nodes updated at both sites.

This architecture also allows for rolling upgrades to be made to the CMA systems. One system can be taken down and its operating system, application, database, or hardware upgraded while the other system carries the CMA load. The first system can then be returned to service and, likewise, the second system taken down for an upgrade.

The Merchant's Second Replication Architecture: Active/Active

For the CMA system, all switching data is kept in memory, and the switch configurations are kept on disk in a configuration database, which also contains a map of failover configurations so that workload on the failed nodes can be routed to surviving nodes. Consequently, the configuration database is fundamental to proper switch operation.

After several years of successfully operating in its DR architecture, the merchant upgraded to an active/active architecture. A copy of the configuration database is maintained on each node for local access. These copies are kept synchronized with the Shadowbase bi-directional data replication engine; any change made to a copy of the configuration database on any node is immediately replicated and made available to all other nodes.

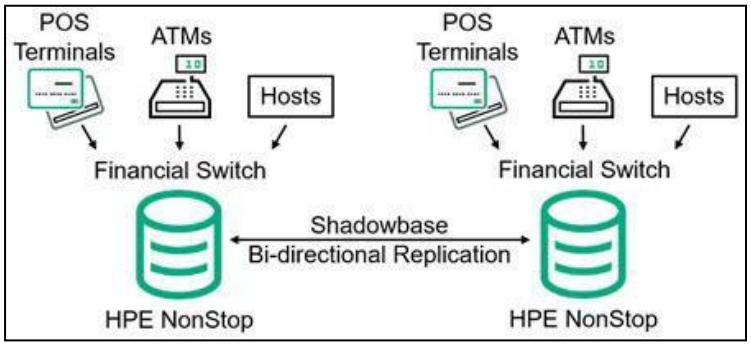

Figure 1 — Active/Active Architecture for a Financial Message Switch Provider[1]

As shown in Figure 1, the nodes are partitioned by card numbers. Each card is assigned one node, which avoids the potential problem of data collisions arising from card updates. All transactions for a card are routed to a single node; the node makes transactional updates to its database, and Shadowbase software replicates the changes to the other database copies in the network. Since no two nodes are making simultaneous changes to a card, data collisions are avoided.

The Shadowbase data replication engine performs another important function. As the failover configuration is modified (which must be done for every new card), the update is entered on the node receiving the configuration updates. However, this configuration data is different for each node, since each node takes a different action on failover. For instance, node 1 might handle card A. If node 1 fails, then node 2 might handle card A. If node 2 then fails, then node 3 might handle card A. By applying specific business rules, the Shadowbase engine automatically adjusts the configuration data for each node as it replicates the data.

The company converted its independent, four-node switch to an active/active configuration, which can almost instantly recover from any fault in the system, including a total node failure, with a recovery that is transparent to client company's end users.

The Merchant's Third Replication Architecture: Real-time Fraud Detection

A major function of payment authorization is to detect the fraudulent use of cards. If a transaction looks suspicious – for instance, a new transaction occurs in London while the previous transaction an hour earlier occurred in Boston – it may be rejected.

The merchant faced the difficult and costly problem of detecting and stopping fraud. It would be too difficult to develop and implement a home-grown solution on the application, and taking the application offline was not an option. Each transaction received by the authorization switch needs to be incorporated into a fraud detection system that maintains a log of recent card transactions and compares them with other transactions for suspicious use on the same card.

Communication between the fraud detection system and the authorization switch is achieved by using the Shadowbase data replication engine. When a transaction is received, the request is passed via the engine to the fraud detection application. The fraud detection system then "scores" that transaction against other card activity to determine if the activity appears suspicious. The response from the fraud detection system is then replicated by Shadowbase *heterogeneous application integration* back to the authorization switch for further analysis. The authorization switch nodes are HPE NonStop systems and the fraud detection systems are Unix/Oracle-based Figure 2.

[1]The merchant's architecture actually contains four nodes in an active/active network, but only two are depicted for simplicity.

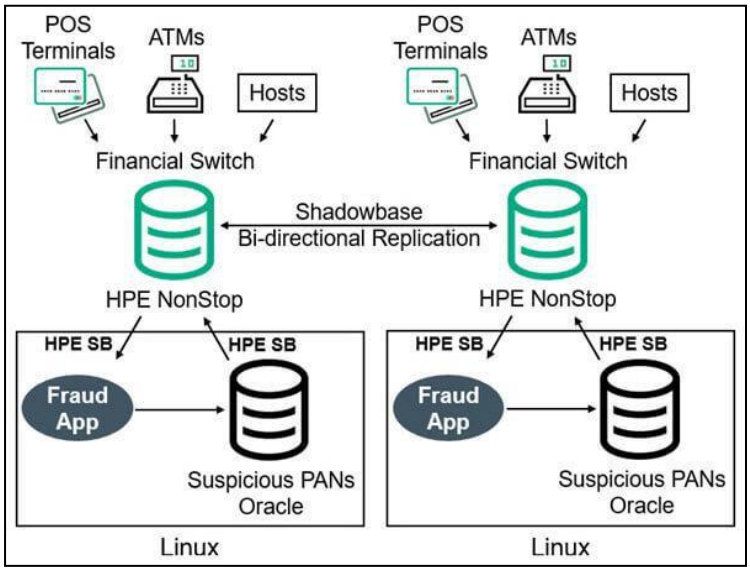

Figure 2 — Active/Active Architecture with Real-time Fraud Detection

Summary

This merchant provider carefully expanded its use of data replication and integration software over the years to add new capabilities and features to its financial switch. First, the company used data replication to provide an active/passive business continuity solution for its CMA system. Second, the company expanded its use of the technology to provide continuous availability for its multi-node switch, via an active/active business continuity solution. Third, to help prevent card fraud, the company used heterogeneous application integration to integrate the switch with a remote heterogeneous application (the fraud detection system). Finally, the company integrated additional data validation and consistency checks to satisfy stakeholders and auditors that the various copies of the data are all in sync. This case study demonstrates some of the many ways HPE Shadowbase data replication software significantly enhances enterprise applications.

But how do I know my backup data is consistent?

After planning, installation, testing, and creating a working business continuity plan, including the backup systems, network interconnect and databases, the inevitable happens. It is not much use switching to a backup system if its data is not an accurate reflection of the data on the production system (garbage in, garbage out). But how do you *know* that it is accurate and a consistent copy of the now-lost production database?

Long before the inevitable happens, an integral part of your business continuity plan is to verify that your backup databases accurately reflect your production databases. Periodic and continuing data validation satisfies this requirement, and also plays a vital role in helping satisfy audit and regulatory compliance requirements. As part of performing periodic full failover tests, another check must be made to verify data consistency.

Large European Financial Institution

For these reasons, a large European financial institution wished to validate the consistency of its backup database. This institution manages a message switch that generates more than a terabyte of audited data per day, on a database size of about 0.5 terabytes. Additionally, the customer base and subsequently, transaction workload demand is growing.

The Old Method

The institution ran its production and backup comparison processing overnight using a homegrown data comparison script (every row for both databases were “brute-force” compared, and the process took a long time and was very inefficient). As customer growth continued, the overnight window used for the comparison shrank rapidly, forcing the institution to search for alternatives for better throughput and overall efficiency.

HPE Shadowbase Compare

Fortunately, HPE Shadowbase Compare provides a much more efficient way of performing comparisons leveraging *data compression*. Rather than sending every row to the backup system, blocks of data are run through a data compression algorithm, which generates a hash key representing that chunk of data. This hash key is significantly smaller than the entire block of data it represents. The hash key for each block of data is sent from the production system to the backup system. Using the same algorithm, the backup system computes the hash key for the same block of data in the backup database. If the two hash keys match, then it is known that the two blocks of data are consistent, and the process moves on to the next block. If the hash keys do not match, the backup system sets aside the information regarding the offending block for an additional re-check.

Once the entire database is compared in this way, the backup system processes the set-aside blocks. There are several ways these blocks can be handled, but most simple is to just retry the comparison, e.g., sometimes there may be a delay in the application of backup data into the database. Alternatively, the production system sends across the entire block uncompressed, which can then be compared on the corresponding data in the backup system to find the specific row/column inconsistency.

Once the inconsistencies are identified (rows in the production database not in the backup database, rows in the backup database not in the production database, rows in both with columns that do not match, etc.), they can be processed in several ways. One method is to start with a Shadowbase Compare report to detail the inconsistencies found, and then manually resolve them.

More and more frequently, customers use the Shadowbase Repair function provided with Shadowbase Compare. Repair is an out-of-the-box function for resolving data discrepancies. Since the production system is always assumed to be correct, this function typically entails overwriting the data on the backup system, however this frame of reference can be swapped if the customer decides to reverse the repair function.

By switching to Shadowbase Compare, the financial institution dramatically shortened the time required to fully validate its backup, easily completing the comparison within the overnight window with time to spare. An important part of the solution is the usage of multiple parallel comparison operations, each against a distinct subset of the database. The institution is now prepared to failover to its backup system and prove to auditors that their production and backup systems are consistent.

International Partner Information

Global

Hewlett Packard Enterprise
6280 America Center Drive
San Jose, CA 95002
USA
Tel: +1.800.607.3567
www.hpe.com

Japan

High Availability Systems Co. Ltd
MS Shibaura Bldg.
4-13-23 Shibaura
Minato-ku, Tokyo 108-0023
Japan
Tel: +81 3 5730 8870
Fax: +81 3 5730 8629
www.ha-sys.co.jp

Gravic, Inc. Contact Information

17 General Warren Blvd.
Malvern, PA 19355-1245
USA
Tel: +1.610.647.6250
Fax: +1.610.647.7958
www.shadowbasesoftware.com
Email Sales: shadowbase@gravic.com
Email Support: sbsupport@gravic.com

Hewlett Packard Enterprise Business Partner Information

Hewlett Packard Enterprise directly sells and supports Shadowbase Solutions under the name ***HPE Shadowbase***. For more information, please contact your local HPE account team or visit our website.