Stock Exchange and Clearinghouse Gain Operational and Availability Benefits with HPE Shadowbase Solutions

A Gravic, Inc. Case Study

Stock Exchange and Clearinghouse Improve Operations and Availability with HPE Shadowbase Solutions

Executive Summary

A regional South American stock exchange embarked on a project to decrease settlement time and improve the accuracy of its trading operations. Commonly referred to as "T+3 Settlement" in the industry, the stock exchange commonly ran into issues when it tried to settle its trades within the three allotted days. Additionally, with the industry heading towards the "T+0" immediate settlement paradigm, also referred to as *Straight-Through Processing* (STP) in the industry, the exchange found that it needed to fully automate trading operations to avoid manual re-entry of settlement information.

Previously, a high error rate was reported in trades because traders were required to re-enter trade details manually to the exchange's clearinghouse. The resulting manual reconciliation process required to correct these errors slowed down the clearing to the extent that settlement commitments often could not be met. Therefore, the exchange needed to re-architect its interaction with the clearinghouse to make it fully automated, and did so using HPE Shadowbase data replication.

The exchange also manages several Linux/MySQL data warehouses that provide trading history to its traders. The exchange was previously updating these data warehouses using a periodic batch update process, which unfortunately led to the data in the warehouse being stale and out of date. Consequently, the exchange decided to replace this process with real-time HPE Shadowbase data replication to keep the data warehouses current.

As part of its IT re-architecture process, the exchange also decided to standardize its data replication and business continuity solutions to a common replication vendor. This standardization required heterogeneous uni-directional data replication solutions for the clearinghouse interface and for feeding the data warehouses, as well as bi-directional homogeneous data replication solutions for the exchange's new business continuity architecture. The exchange chose HPE Shadowbase replication from to satisfy these current and future needs. HPE Shadowbase software now plays a major role in integrating the many systems in the exchange's IT infrastructure, as well as providing continuous availability for its mission-critical business services.

Table of Contents

Table of Figures

Stock Exchange and Clearinghouse Improve Operations and Availability with HPE Shadowbase Solutions

A major South American stock exchange found that it could not reliably achieve same-day and next-day securities settlement commitment times due to the manual reporting of trades to its clearinghouse, which often resulted in data entry errors. It then chose to re-architect its interaction with the clearinghouse to make it fully automatic, which eliminated erroneous manual input, and enabled meeting its settlement commitment times.

The re-architected system required data replication between heterogeneous databases. The exchange chose HPE Shadowbase data replication to satisfy this need. HPE Shadowbase software now plays a major role in integrating the many systems in the exchange's new IT infrastructure, as well as providing continuous availability for its mission-critical business services.

The Exchange

The stock exchange is the result of a 2001 merger of three regional stock exchanges in the country, one of which dates back to the late 1920s. The exchange operates a trading platform for equities, fixed income bonds, and standardized derivatives, using an HPE NonStop server pair in an active/passive configuration for disaster recovery. In addition to its trading functions, the exchange provides equities and fixed income clearing, settlement services, market data, and electronic access to accounts. The exchange lists 89 corporations with a total market capitalization of $200 billion USD.

The Clearinghouse

To provide clearing and settlement services, the exchange partnered with a clearinghouse that acts as a security depository. A security depository maintains all of the stocks and bonds of its account holders electronically so that ownership can be easily transferred via a book entry rather than a physical transfer of certificates. The clearinghouse provides the clearing and settlement functions per the settlement cycles provided in the settlement agreements.

At the end of the trading day, the clearinghouse determines which members are due to deliver funds or securities and which members are due to receive funds or securities by the settlement date. Each trade is assigned to a buyer and a seller, and any funds and securities being transferred between them must match. Settlement is a two-way process that involves the transfer of funds and securities on the settlement date between the trading parties.

The clearinghouse does not transfer paper stock certificates or bonds, though paper records of account holders' holdings are prepared. Rather, the record of securities being held by account holders is maintained in its database (so-called dematerialized securities). Because there is no need to physically transfer physical securities, the clearing and settlement process can be very fast.

Straight-Through Processing (STP)

It is normal for settlement schedules to be measured in days (typically three business days, or trade-plus-three, "T+3"), since the clearing and settlement process can be highly manual and very complex. It is the exchange's goal to achieve *Straight-Through Processing* (STP), or "T+0," in which settlement can be done on the same day or on the next day, which requires that the entire clearing and settlement process be automated.

The clearinghouse used by the exchange meets this criterion. All securities are maintained electronically so that trades can be matched and assigned to buyers and sellers with no manual intervention. A trade is rejected for manual reconciliation only if there is a mismatch. For instance, Trader A says that he sold 100 shares of IBM to Trader B, but Trader B says that he bought 200 shares of IBM from Trader A.

The problem that the exchange faced with consistently achieving T+0 was that the entry of trade information into the clearinghouse's system was an error-prone, manual data input process. This method resulted in a high rate of trade reporting errors, which led to a high rate of trade rejects requiring manual reconciliation. The result was a high rate of STP settlement violations in which trades could not be settled on the same or next day.

In order to correct this problem, the exchange and the clearinghouse cooperated to allow the electronic entry of trades from the exchange's systems into the clearinghouse system. Since this solution was achieved, the exchange has been able to offer true STP to its traders.

The Original System

The original method for trade entry, reconciliation, and clearing is shown in Figure 1. The exchange's trading system is built upon a pair of HPE NonStop Server systems configured as an active/standby pair for disaster recovery. The trading databases use HPE NonStop Enscribe and HPE NonStop SQL/MP. The primary system is a four CPU NS2100 NonStop Server, and the standby system is a two CPU NS2100 NonStop Server.

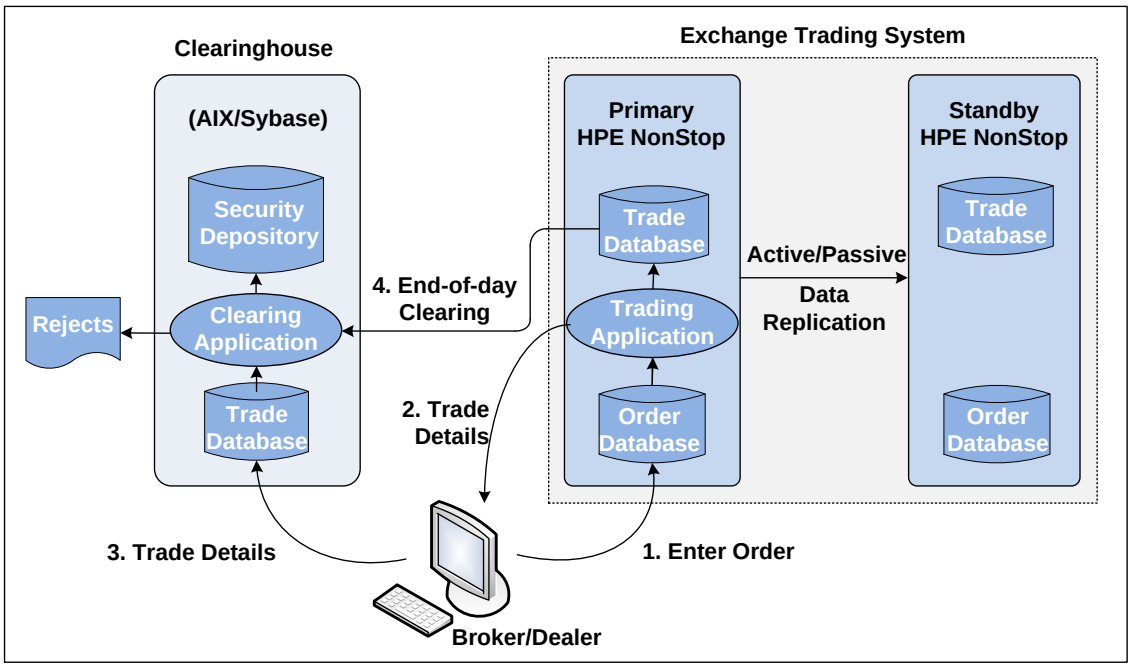

Figure 1 – The Original Trading and Clearing System

The exchange's trading application is active in the primary node of its trading system, and the primary node's databases are replicated in real-time to the standby node in an active/passive disaster recovery architecture. The standby node is thus ready to take over the trading function should the primary node fail. To do so, it must bring up the trading application, mount the trading databases, and switch over the traders to the new primary system.

Traders (the brokers/dealers) enter buy/sell orders via their trading terminals. These orders are stored in the exchange's order database awaiting execution. As the exchange's trading application executes orders, the trade details are stored in the exchange's trade database. The trade details are returned to the originating traders as a trade confirmation.

The clearinghouse maintains the security depository, which is the record of ownership for all securities managed by the clearinghouse. At the end of the day, the day's trading activity is reconciled by the clearing application, and the security depository is updated to reflect the new security ownerships.

The procedure used to trade and clear with this system was originally implemented as follows:

- A broker/dealer entered an order at his trading terminal and submitted it to the exchange (1).
- All pending orders were stored in the exchange's order database.
- When the exchange's trading application was able to execute an order, the result of the trade (successful or not, the price, the broker/dealer on the other side of the trade, etc.) was stored in the exchange's trade database.
- The trade details were returned to the broker/dealer and displayed on his trading terminal (2).
- The broker/dealer then had to manually enter the trade results into a trade database resident in the clearinghouse's system (3).

- At the end of the day, the exchange's trading system transferred its record of the day's trading activity to the clearinghouse's clearing application (4).

Unfortunately, Step (3) proved to be time consuming and highly subject to data entry errors. Typically, a trader was very busy intensively managing multiple trades and his primary focus was not on this manual intrusion into his primary trading activity.

At the end of the day, the clearing application attempted to match each trade that it received from the exchange to the trades that it had accumulated during the day from the broker/dealers. Inevitably, there was a high error rate, and many trades did not match. These trades had to be rejected and manually reconciled. The time that it took to do these manual trade submissions and eventual reconciliations for mismatched trades added overhead and extra cost and prevented the exchange from meeting many of its STP time commitments.

Enhanced Trade Submission and Reporting Using HPE Shadowbase Solutions

To correct this problem so that STP could be achieved, the exchange and clearinghouse cooperated to eliminate the requirement for the manual entry of trade results by the broker/dealer, allowing electronic entry of trades from the exchange's systems directly into the clearinghouse system. The modifications to the original system to achieve this enhancement are shown in Figure 2.

The first step was to incorporate the clearinghouse's validation database to the exchange's trade entry verification sequence. This database contains all of the consistent and correct values for trade parameters, including security CUSIP numbers, symbols, and valid broker IDs and names. Updates to this database are entered by the clearinghouse, and the validation database is replicated to the exchange's trading system in real time in order to keep the two synchronized. The exchange's trading system uses the validation database to ensure that orders entered by the broker/dealers are accurately formatted and that they contain standardized parameters. Thus, all orders and their subsequent trades are guaranteed as consistent and acceptable to the system applications.

One problem to be overcome was that the clearinghouse's validation database was a Sybase database stored on an IBM AIX system, whereas the exchange's copy of that database was a NonStop SQL/MP database stored on a NonStop server. The clearinghouse recognized the need to reconcile these two unlike databases, and chose data replication as the mechanism to perform the data-sharing. Replication between the source and target required heterogeneous (dissimilar) replication. The exchange chose HPE Shadowbase replication, which supports a wide variety of source and target systems and databases.

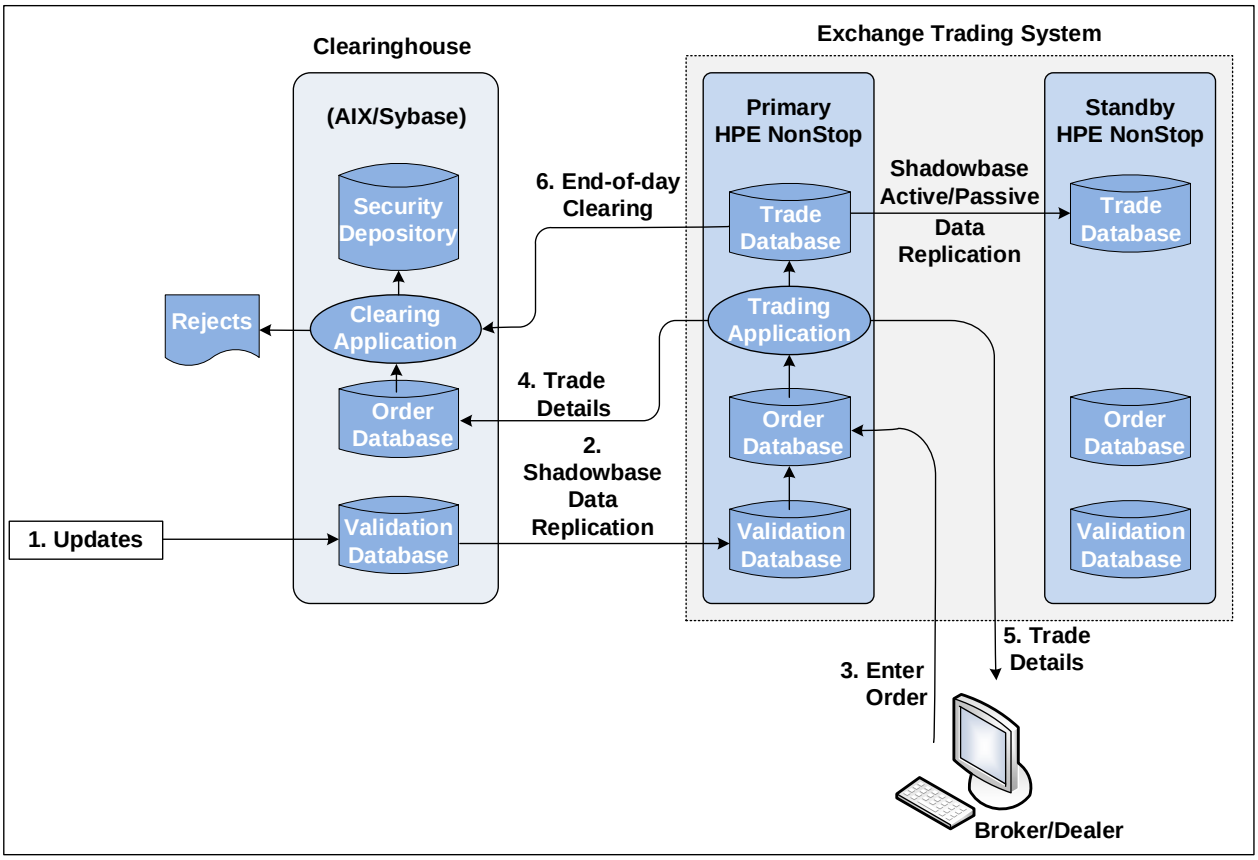

Figure 2 – Enhanced Trade Reporting System

With reference to Figure 2, the trading procedure is now as follows:

- Updates to the clearinghouse's Sybase validation database are entered by the clearinghouse (1).
- These updates are replicated in real-time to the exchange's NonStop SQL/MP validation database by HPE Shadowbase replication (2).
- A broker/dealer submits an order to the exchange's trading system (3).
- The order parameters are verified against the exchange's validation database.
- If there is a parameter error, the order is rejected.
- If the order is formatted properly, it is placed into the exchange's order database.
- The exchange executes orders in the order database as soon as it is able, based on the order parameters.
- On execution, the trade details are sent by the exchange to the clearinghouse, which stores the trade details in its trade database (4). (This step replaces the earlier manual re-entry of trade data by the brokers/dealers).
- The results of the trade (the executed order) are returned to the broker/dealer (5).
- At the end of the trading day, the exchange sends its entire trading history for the day to the clearinghouse (6).
- The clearinghouse compares these trades to the trades that it has received during the day from the exchange that are stored in its trade database.
- If there is a mismatch, the trade is rejected for manual resolution.
- For matching trades, the trades are cleared and the security depository is updated to reflect the new ownerships.

Now that there is no manual re-entry of trades required of the broker/dealer, the number of trades rejected for manual resolution is minimal, and the exchange is able to meet its STP obligations.

The Exchange Chooses HPE Shadowbase Data Replication for Improved Service Availability

The exchange was using an existing replication product to provide data replication between the NonStop servers in its active/passive disaster recovery pair, in order to provide continued business service in the event of an unplanned outage of the primary system. However, the existing replication product ran into various issues, and, after choosing HPE Shadowbase data replication, the exchange realized that it did not make sense to have multiple replication products in production. Therefore, it moved to HPE Shadowbase replication for its active/passive HPE NonStop server replication. In future project phases, the exchange is also considering upgrading to an architecture with even higher levels of availability, such as HPE Shadowbase *Sizzling-Hot-Takeover* (SZT) or fully active/active architectures, which offer higher and extreme levels of business service availability.[1]

The Exchange's Business Continuity Architecture Imposes Asymmetric Failover Processing

Note that the exchange is using a four CPU primary system and a two CPU standby system for the exchange applications. This mode of operation lessens the cost of software licensing and hardware. However, it can lead to load shedding (or reduced response times/processing power) in the event of a failover to the standby system. The exchange has identified various non-critical applications (e.g., certain reporting subsystems) that will be taken offline in the event of a failover to reduce the overall system load. Although this approach saves costs, anyone considering a similar configuration should carefully weigh the benefits vs. the costs and potential failover issues, to determine if it is a suitable solution. Of course, regardless of the architecture, everyone should periodically test the failover and processing functions on a standby system to ensure it meets acceptable criteria. Selecting this architecture also requires a data replication engine that can handle heterogeneous configurations, a challenge HPE Shadowbase technology is particularly adept at managing.

The Exchange Chooses HPE Shadowbase Data Replication for Data Warehouse Feeds

In addition to its trading system, the exchange provides a series of query nodes that act as data warehouses for trading history, as shown in Figure 3. These data warehouses are useful for the brokers and dealers to track a security over a period of time, an analysis that helps them make purchase and sales decisions. The exchange's data warehouses use MySQL databases running on Linux servers.

At the current time, the exchange has deployed ten data warehouse systems. Seven are co-located with and are connected to the primary exchange system's NonStop server, and three are co-located with and are connected to the standby server.[2]

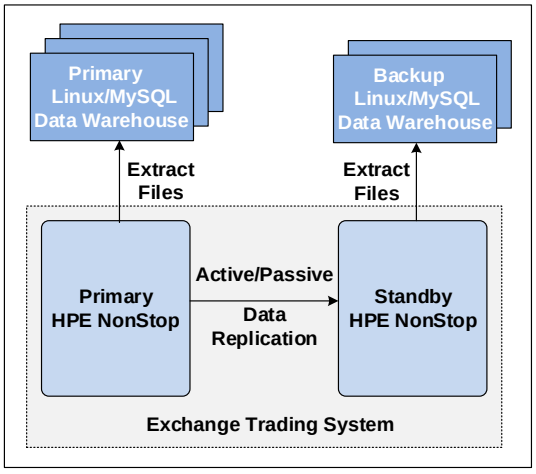

Figure 3 – The Exchange's Original Data Warehouse Architecture

Since the previous replication product did not directly support heterogeneous replication between HPE NonStop SQL/MP databases and MySQL databases, the exchange was forced to periodically create mini-batches of trading activity and replicate the activity as flat files via the existing replication product from the exchange's HPE NonStop servers to the Linux data warehouses, as shown in Figure 3. These files were then periodically loaded into the target database. This approach meant that the data warehouses were often lacking the latest data for supporting broker queries.

HPE Shadowbase software provides real-time, transactional replication of updates from the exchange's source HPE NonStop databases into the target Linux databases (data warehouses). Since the exchange experienced success before with HPE Shadowbase data replication, they chose it again for this data integration function. Trading activity is now replicated to the data warehouses in real-time by Shadowbase replication, and the information in the data warehouses is kept current with the exchange's HPE NonStop source databases.

[1]For more information, please see: *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.

[2]Although this approach saves costs, anyone considering a similar configuration should carefully weigh the cost benefits vs. the potential failover issues (load shedding, reduced response times/processing power, etc.), to determine if it is a suitable solution for them.

The Exchange Discovered and Leveraged HPE Shadowbase Data Replication for Multiple Uses

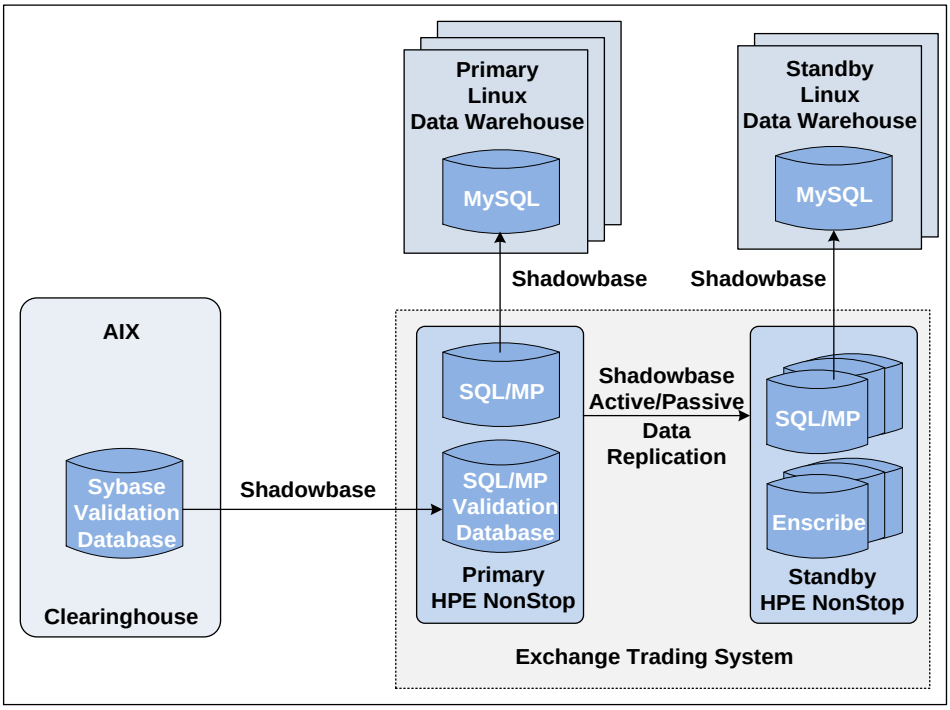

Figure 4 – The Exchange's Many Uses of HPE Shadowbase Data Replication

Once the exchange chose HPE Shadowbase software to meet its original heterogeneous replication requirements, it subsequently discovered several other key product features (discussed below and summarized in Figure 4).

HPE Shadowbase solutions:

- Provide disaster recovery with active/passive replication between the exchange trading system's NonStop servers. Both NonStop SQL/MP and Enscribe databases are replicated.
- Replicate the validation data from the clearinghouse's Sybase database running on an IBM AIX to the NonStop SQL/MP validation database running on NonStop servers (the original requirement to solve the exchange's manual data entry problem).
- Replicate real-time trading data from the exchange's NonStop SQL/MP databases to the data warehouse MySQL databases running on Linux systems (keeping the data warehouse information synchronized with the current state of the exchange's primary trading database).

It is the flexibility of the HPE Shadowbase homogeneous and heterogeneous data replication capabilities that allow it to provide this wide variety of functions, in this case demonstrating its suitability for both business continuity and data/application integration purposes.

Upgrading to Continuous Availability

With the current active/passive architecture being used by the exchange for its trading system, a failure of the primary system could bring down the exchange for a significant period of time – several minutes to perhaps several hours. The downtime happens because at the point of failure, management must first be apprised of the situation and approve the failover (often leading to extended decision time), the failover must then be initiated (network switching, interface failovers, etc.), and the trading applications must be brought up on the standby system with the databases opened for read/write access. The system must then undergo testing before it is put into service. All of these steps take time, and there are many subtle problems that could lead to

failover faults which would prevent the standby system from being brought into service, resulting in extended downtime.

Gravic and the exchange discussed various architectures to achieve higher levels of service availability, including those that eliminate the risk of failover faults, including Shadowbase bi-directional replication for continuous availability. The proposed architecture is HPE Shadowbase SZT, where the backup system is not idly standing by. Instead, all applications are loaded with all databases mounted for read/write access, but not for processing any UPDATE transactions. (It can process read-only query transactions.) If the primary system fails, traffic is routed to the standby node, which is a *known-working* system, so that failover is accomplished in just a few seconds. With this approach, the standby node is continuously tested by periodically sending test transactions to the node to ensure that it is operational end-to-end. There is no need to take an outage of the primary node for testing, which is the case for active/passive architectures. Therefore, there is no possibility of failover faults. Application recovery is fast and reliable.

HPE Shadowbase SZT is possible because it allows the applications to operate and the databases to open for read/write access on all nodes. Each update made to the primary database is replicated across the application network, keeping all databases in synchronism. Going forward, the exchange now has the opportunity to move to the more reliable Shadowbase SZT configuration to eliminate downtime from failover faults even in the presence of a primary system failure.

Summary

A stock exchange using a trading system built with HPE NonStop servers faced a serious problem with high data entry error rates and the resulting manual reconciliation process required to correct these errors, which prevented it from complying with same-day or next-day clearing commitments.

Working with its clearinghouse, the exchange re-architected its system to eliminate manual data re-entry. Orders were pre-validated to ensure that their format was consistent and correct; and the resulting trades were sent directly and electronically from the exchange to the clearinghouse, avoiding the manual data re-entry of that information. A necessary part of this new architecture was the requirement to replicate a Sybase validation database from the clearinghouse's AIX system to a NonStop SQL/MP database located on the exchange's NonStop trading system. The exchange chose HPE Shadowbase data replication for this task, which provides replication between heterogeneous systems and databases.

The exchange also manages several Linux/MySQL data warehouses that provide trading history to its traders. The existing data replication tool that the exchange was using could not replicate from the HPE NonStop databases to the MySQL databases in real-time. Therefore, the exchange was limited to sending periodic batch updates from its trading system to the data warehouses, causing stale information in the data warehouses. The exchange decided to utilize HPE Shadowbase heterogeneous replication, and the batch updates were replaced with real-time transactional data replication from the NonStop SQL/MP database to the MySQL databases, meaning that the information in the data warehouses was now current with the primary trading system.

Given that the exchange had now made a commitment to use HPE Shadowbase replication for several purposes, it also replaced its current replication product being used for disaster recovery of its trading system (providing active/passive data replication between the primary and standby NonStop servers), with Shadowbase solutions. Because Shadowbase replication supports bi-directional replication, the ability to open the target database for read/write access, and extreme availability with its active/active architectures, this support gives the exchange an opportunity in the future to move to a continuously available HPE Shadowbase SZT or fully active/active architectures, thereby improving its overall application availability.

This case study demonstrates the capabilities and flexibility of HPE Shadowbase software for homogeneous and heterogeneous data replication, and its business continuity and data and application integration features.

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.