One Bank's March Towards Active/Active

A Strong Reminder Why Successfully Testing Your Business Continuity Plan Actually Matters

Paul J. Holenstein
Executive Vice President
Gravic, Inc.

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, uni-directional and bi-directional, 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.

Not many banks avoided exposure to the recent subprime crisis and speculative real-estate mortgage meltdown. One bank that did, due to its rational credit policies, remained the number one lender in its area while other financial institutions severely restricted credit to their customers.

However, the bank was not as well-prepared for Mother Nature. The bank is located on the Pacific Rim Ring of Fire (Figure 1) and lost its data center for several hours during an earthquake. When the failover to the backup system did not go as planned, the bank said, "Never again!" and began its march towards active/active.

The Bank's Disaster Recovery Plan – Faith and Hope

The bank operates two HP NonStop data centers that are located about 50 miles apart as an active/backup pair. Given the geographic challenges facing the bank, this short distance put the data centers in a reasonable "geographic redundancy" probability zone to minimize the risk of a single disaster consuming both data centers.

Prior to the earthquake, the backup database was kept synchronized with the production database via a data replication product and network topology that supported only active/passive architectures in which the target environment had to be an exact mirror of the source.

Ring-of-Fire
Figure 1

Therefore, the bank followed what is unfortunately a common practice. Although it periodically performed failover testing, its tests often ran into many small problems and ended up incomplete. In the end, the bank relied on faith and hope that its backup system would come up in a reasonable amount of time should the primary system fail.

Many companies take similar shortcuts to make it through their tests, but such expedients can lead to catastrophic results. This bank realized that by leveraging its existing technology in more powerful ways, it could build backup environments that are always in a known-working state, and improve its availability profile without adding significant cost or complexity.

Mother Nature Strikes

ATM/POS Downed by Earthquake
Figure 2

One fateful day, an earthquake struck and caused the production systems to fail. The bank initiated its disaster recovery plan. As might have been predicted, the bank suffered a *failover fault.* It could not bring its backup systems into operation. The most critical outages were those of its online banking services and its ATM/POS network (Figure 2). This was a terribly critical time since people needed to buy supplies and take other actions to recover from the damage caused by the earthquake.

The problem was further aggravated by the fact that the IT staff had been evacuated from the primary data center due to concerns about structural damage. Hours passed before the bank's staff could reenter the production data center and bring up the production system in order to restore ATM and POS services to the community.

The Search for Continuous Availability

This disastrous experience led the bank to realize that its approach to disaster recovery was unacceptable. It concluded that it needed a backup system that was known to be working and that could be tested frequently without affecting its customers. Additionally, it wanted to leverage the new approach to avoid application outages for common procedures such as O/S upgrades and deployments of new application versions.

The Limitations of the Bank's Architecture

The bank came to understand that its backup approach was constrained by the data replication technology and network topology that it had adopted. It had chosen a replication product and network architecture that did not support the functionality needed to ensure rapid and reliable failover:

- The replication product required the production and backup systems to be configured exactly the same.
- The replication product prevented backup applications from opening the database in update mode. Therefore, applications on the backup system could not be running with the database mounted for fast failover.
- The replication product provided only uni-directional replication. The bank could never move to a configuration in which both systems were actively processing transactions, informing each other as to the database changes they were making.

The fact that applications could not be running on the backup system limited the bank to an active/passive configuration, in which the backup system was idle except for being a replication target. Compounding this challenge was the fact that configuration errors could not be detected until a failover occurred. Moreover, there was no way to easily test the backup system's applications without taking a production outage. Between the complexity of the failover process and the problem of configuration errors, failover not only was difficult and time-consuming, but it also was unreliable.

The Road to Availability Improvement

The bank decided that it had to move from an unreliable disaster recovery architecture to a known, working disaster tolerant architecture. Disaster recovery means that the IT systems can recover from a disastrous event and continue operating, even if that means hours or days of downtime. Disaster tolerance means that recovery is so fast that no one notices the outage or at least is not inconvenienced by it.

Implementing a disaster tolerant architecture can be a daunting task. However, it can be accomplished via a controlled process that achieves incremental improvements.

The Availability Improvement Process

The bank's availability improvement process proceeded as follows:

Step 1: Define Requirements

The bank began by reviewing its options for a new replication product. In order to support fast and reliable failover, the replication engine had to have the following characteristics:

- The replication engine had to have the capability to allow applications to be up and running on the backup system with the database mounted, ready to take over processing in an instant's notice (we refer to this step as a sizzling-hot-takeover architecture).[1]
- The backup database always had to be in a consistent state during replication so that it could be used immediately following a failover.
- The production and backup systems had to be decoupled so that they did not have to be configured identically, thus eliminating failover faults due to configuration errors.
- The delay in replicating database changes (the replication latency) had to be small to minimize data loss following a production system failure.
- The replication product had to support bi-directional replication so that the failed system's database could be easily resynchronized upon its return to service.

Step 2: Choose a Data Replication Engine

The bank next evaluated the various replication alternatives that were available for NonStop systems. It chose the Shadowbase replication engine from Gravic, Inc. (www.gravic.com/shadowbase) as the one that best satisfied its requirements.

Shadowbase replication supports bi-directional data replication and data integration. Applications can be actively running on both systems and simultaneously updating the application database. Replication is process-to-process, leading to small replication latency times. Shadowbase technology can replicate between heterogeneous systems, so maintaining identical system configurations is not a requirement.

Step 3: Switch Replication Engines

Before taking any further steps, the bank wanted to make sure that it was comfortable with its new data replication engine. To ensure this, it replaced its original replication engine with Shadowbase software doing the same active/passive job.

Once it was satisfied with the performance and functionality of Shadowbase replication, the bank took advantage of this step to upgrade its HP NonStop S-series servers to HP NonStop NS servers. It installed NS servers

[1] Ideally, to achieve fully continuous, load-balanced application availability that leverages all of the capacity of the available systems, the applications on the backup system should also be able to process transactions simultaneously with the production system. We call this lofty goal a fully active/active system. Whether or not this goal is attainable in your environment should not impact your efforts to leverage the other advantages mentioned in this article for improving application availability.

in its production and backup data centers and used its new replication engine to synchronize the two NS servers and the backup S-series server with the production S-series server. When this step was complete, it switched its transaction load to the new production NS server replicating to the backup NS server and retired its older S-series servers. The entire upgrade was accomplished with little if any downtime.

Step 4: Configure Bi-directional Replication

The bank's next step was to extend to bi-directional replication, making failover testing simpler. If backup applications are not already running, then the bank starts them, switches the network, and tests the backup system. The production database is maintained in a current state by bi-directional replication. Therefore, fallback is simply a matter of rerouting the transaction stream back to the production system.

Equally important, applications on the production system can continue running. If the testing on the backup system is against test or verification accounts, the production system can continue processing production requests.

Step 5: Configure the Fast Failover System

Once it has become comfortable with bi-directional replication, the bank will be in a position to reconfigure for fast failover. It can put both systems into operation with all applications up and running. Transactions could be sent to either system for proper processing. However, the bank will direct all transaction activity to only one system; the other system will serve as a sizzling-hot-standby.

Since the application is already running on the standby node, the standby system can be tested frequently by simply sending test or verification transactions to it. These transactions will verify the full end-to-end processing capability of the standby application. Therefore, in the event that the standby system should be needed, the bank will know it is operating properly and the system will take over with no failover faults.

An additional advantage of configuring in this mode is that it simplifies switch-over processing to the point that management can direct the staff to switch over frequently, making sure that both nodes are always ready to assume the processing load. Frequent testing leads to ensuring that the staff is comfortable and well-versed in the failover process.

With only the acquisition of a proper replication engine and some system reconfiguration, the bank will have moved from multi-hour unreliable failover to multi-second reliable failover. It will have achieved its goal of continuous availability with no change in its hardware configuration.

Step 6: Move to Active/Active (Future Option)

At this point, the bank will be in a position to take this process one step further if it wishes. It could put both of its

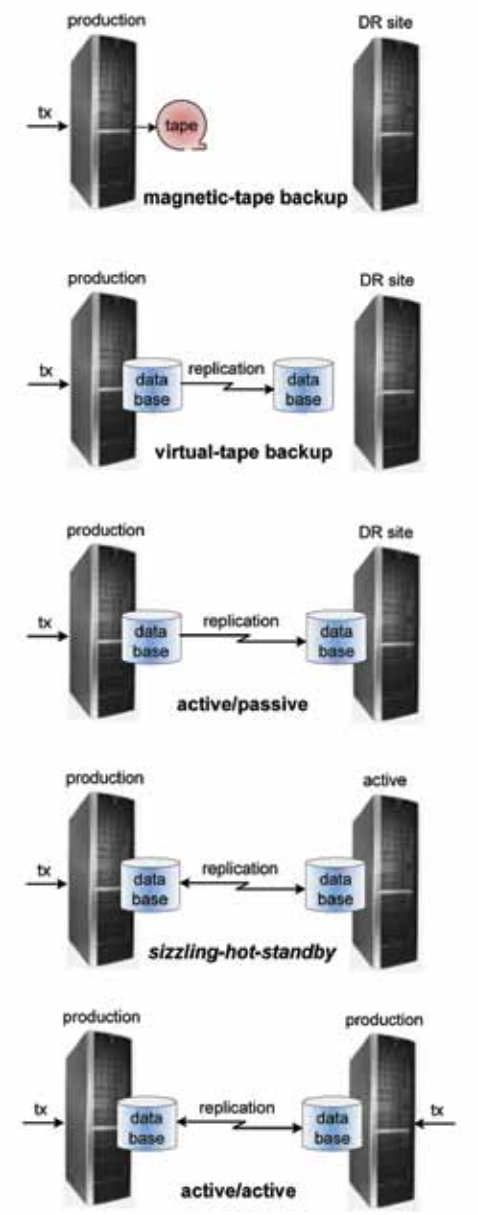

Moving to Continuous Availability
Figure 3

systems into active production. This is called a fully *active/active system*. Since both systems can process transactions, the transaction workload can be split between the two systems. Should one system fail, it will only be necessary to reroute all further transactions to the surviving system.

Fully active/active technology is the approach that has been implemented by many other banks and financial institutions, including the Royal Bank of Canada, Fifth Third Bank, and FDC.

Summary

The success that this bank has achieved in moving in a controlled fashion towards continuous availability teaches some important lessons:

- Make sure your failover procedures actually work. Do not settle for testing shortcuts that

One Bank's March Towards Active/Active

continued from page 22

could lead to failover faults when a real disaster strikes.

- Replace technology and products that stand in the way of improving your availability.
- Do not give up if you think that implementing higher availability is too complex. As we have discussed, it can be done in an incremental fashion, one step at a time.
- Do not give up if you think moving to higher availability architectures is too difficult or is unattainable. If a fully active/active architecture would not work for your application environment, by all means strive for its slightly lesser sizzling-hot brethren.

As shown in Figure 3, each architectural step, from magnetic-tape backup to virtual-tape backup to active/passive to sizzling-hot-standby to fully active/active, moves you closer to continuous availability. The migration is a process that is managed and controlled to ensure success on your schedule and at your comfort level.

After all, if you already are running a backup site, you already have accepted the cost of redundancy, which is the first requirement for improved availability. Why, then, accept outages measured in hours and the possibility of catastrophic failover faults when you could have continuous availability for the cost of a data replication engine and some reconfiguration? That is the lesson this bank is happy to share.

NonStop Going Green

continued from page 18

doesn't use any more than what it needs.

- My NonStop systems have low usage or are even idle once every day or once a week. Is there any way I can save on power during times of low usage?
  Yes! You can enable Dynamic Power Savings power regulation mode to automatically save power when appropriate. You can also shutdown unneeded systems using Data Center Power Control policies.
- Can I see my power and cooling usage in any form?
  ICpwr in HP SIM will provide power and temperature monitoring.
- What is the history of my power and cooling usage so I can predict the future?
  ICpwr will tell you exactly how much you are using over time and help predict future needs.
- How can I set the power regulation policies for my data center?
  By using the power regulation features, NonStop in J06.14 onwards.

Why isn't NonStop going 'green'? It already has.

Vinay Gupta joined Tandem in 1994 after graduating from Indian Institute of Technology. He has architected and designed many NonStop manageability and serviceability applications (e.g., OSM and NonStop Essentials) and has been the manageability architect of NonStop NS-Series servers, NonStop BladeSystems and CLIMs. He is the driving force of integrating NonStop with HP Systems Insight Manager (SIM) and HP IT Performance Suite products. Vinay leads pan-HP Support Automation Architecture, self-healing Initiatives and availability data collection, storage and analytics. He also contributes heavily to many other pan-HP workstreams related to Converged Infrastructure (CI) architecture, security, common event management, common diagnostics, etc. Vinay is the primary inventor of 2 granted patents and 3 pending ones.

Tom Kondo has been at HP NonStop for 15 years. Currently he is the telecommunications architect and technical liaison between the NED and OpenCall business units. He also is a key architect for the NonStop networking, storage, and system interconnect and the corresponding Architecture Review Board. In addition he serves on the HP Converged Infrastructure (CI) architecture team as well as the HP CI Security architecture team and CATA reviews. Tom has contributed to several technologies over the years including NSAA, Power Management, Quality, Stock Exchange products, and many more. Prior to HP, he was a Principal engineer at Amdahl. Tom is the primary inventor on seven patents and is co-inventor of one that is pending.