the Availability Digest

Tackling Switchover Times

October 2006

The biggest advantage of an active/active system is that failover is almost instantaneous – at least, from the perspective of the user. However, the move to active/active can be a daunting task. As a consequence, many companies are first trying to improve their switchover times from hours to minutes and perhaps even seconds. This gives them time make the active/active move sometime in the future.

Two large enterprises which are currently undertaking this task are Sabre Holdings, a travel products company, and Mayo Clinic. Their stories are told below.

Sabre Holdings

The Company

Sabre Holdings Corporation, headquartered in Southlake, Texas, USA, is a multibillion dollar publicly listed company (NYSE TSG) that provides online travel-related products such as Travelocity to the public. Its applications include reservations, ticketing, and the storage of profiles for its customers. Sabre also provides distribution and technology solutions for their travel industry clients.

Sabre operates a large heterogeneous data center utilizing NonStop servers for some of its critical data processing needs, The NonStop servers are surrounded by many industry-standard servers performing communications and other tasks. Sabre takes on the responsibility for the customer interface and all application development and outsources the data center operations to EDS.

Sabre's Data Center

Within Sabre's data center are four NonStop fourteen-processor S88 servers. These are organized as two two-node systems, with one pair of systems being the active pair and the other pair being the backup pair.

The database is partitioned across the two databases of a pair. The backup database is kept in synchronism

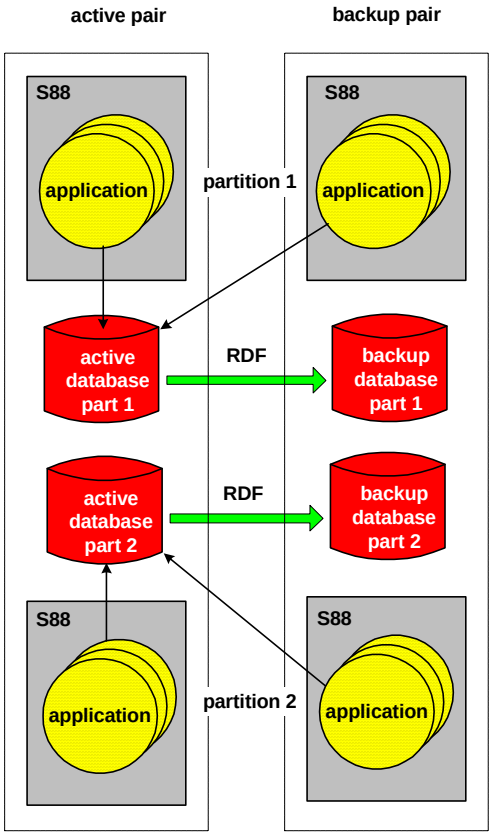

The Sabre Production System

with the active database via HP's RDF (Remote Data Facility) data replication product.

All four systems are used in a load-sharing configuration, with applications running on all systems. However, since RDF target databases cannot be used for update purposes, applications running on the backup pair use the active databases. During busy hours, the systems are handling 250 to 300 transactions per second, many of them being quite complex.

Sabre also runs a two-system test facility configured in the same way except that the databases are not partitioned.

Its Early Switchover Experience

Sabre's initial efforts were aimed at reducing switchover times for planned maintenance. Its goal was to achieve planned switchover times in the order of minutes. The early switchover practice was as follows:

- Hold (or reject) all inbound traffic.
- Stop the application processes on the active system.
- Allow TMF (HP's Transaction Monitoring Facility) to complete on the active system.
- Reset the database access control lists (ACLs) to switch the roles of the active and backup databases.
- Start the application processes on the backup system.
- Release the inbound messages that had accumulated.

At this point, the roles of the active and backup systems had been reversed; and the system that was previously acting as backup took over active processing. Service was restored.

Much to Sabre's chagrin, switchover times turned out to be in the order of four hours. Investigation showed that in addition to software and script bugs and other operational errors which could be corrected, there were some fundamental flaws in the switchover procedure:

- An excess of manual validation of procedure results was required.
- Some batch jobs were scheduled as TMF was trying to complete, thus keeping it from completing.
- TMF would not complete until all transactions were complete. Therefore, it was held up by the longest running transaction.
- Application startup times were excessive.
- Key individuals were not available when problems came up.

The Current Procedures

Using this experience, Sabre made several changes to minimize its switchover problems:

- It rewrote applications to eliminate long transactions.
- It identified processes that could run for a while in read-only mode (such as the process that allowed a user to change his profile).
- It scripted as much of the switchover procedure as was feasible.
- Change management was improved to ensure that the application versions on the active and backup systems were the same.
- Complete documentation was generated to allow all participants the opportunity to learn what to do and to know the current state of the switchover process.

The switchover procedure was enhanced to take advantage of these improvements. The procedure that is now followed separates the switchover process into two stages – a preliminary stage, during which as much of the switchover work as possible is done without taking down the system, and the actual switchover stage, during which the actual switchover occurs. This process proceeds as follows:

- Preliminary Stage:
    - Carefully validate as much of the application environment as possible.
    - Stop all of the processes that can operate in a read-only mode for a short period of time and start them on the backup system.
    - Start frozen copies of update applications on the backup system.
- Switchover Stage:
    - Hold all inbound traffic.
    - Stop all of the processes on the active system.
    - Allow TMF to complete.
    - Reset the ACLs to switch the roles of the databases.
    - Thaw the frozen processes on the new active system.
    - Release the inbound traffic.
    - Stop the read-only processes, and restart them as update processes.

At this time, the new active system is ready to take over. Sabre now sees switchover times in the order of three to twelve minutes as compared to four hours before. The time that read-only processes are prevented from updating ranges from nine to ninety minutes.

The Future

All of Sabre's efforts so far have been focused on switching over for planned outages. The next step is to achieve unplanned switchovers in as little time.

Sabre's goal is to achieve five 9s of availability. This translates into about five minutes of downtime per year, not five minutes per switchover. Therefore, it must get its switchover times down into the seconds rather than minutes.

To do this, Sabre must refine its current procedures further. But the company realizes that an active/active architecture may be the ultimate solution and is beginning to build a business case for that. Already, in order to realize a degree of disaster tolerance, Sabre has separated its data center in Tulsa, Oklahoma, into two geographically separate sites.

Mayo Clinic

The Clinic

The Mayo Clinic is probably the most recognized health facility in the U.S., if not in the world. Its primary facility is in Rochester, Minnesota, with major clinics in Jacksonville, Florida, and Phoenix, Arizona.

In Rochester alone, there are over 1,700 physicians and scientists; a like number of residents, fellows, and students; over 24,000 allied health staff; and over 100 operating rooms and 1900 patient beds. Annual clinic activity averages 70,000 admissions and 48,000 surgeries which are

supported by over 33,000 lab tests per day. All of this activity generates over four million transactions per day.

Paper Medical Records

Mayo began its medical record program in 1907. The volume of archived medical records had increased to over five million by the early 1990s, each hand-indexed for research data mining.

This volume of paper records was approaching the unmanageable, so in 1992 Mayo began its move to electronic medical records (EMR). By the end of 2005, the generation of paper records ended.

Electronic Medical Records

The primary consideration in going to a paperless system is that doctors must have access to a patient's records 24 hours a day, 365 days a year. Otherwise, patient care and safety may suffer.

These documents have to be up-to-date with all of the latest lab results. They are complex documents containing text, images, and lab results. The data must be available as the patient moves around the hospital and must be accessible to doctors at the patient's bedside (doctors now drag a wireless laptop behind them on an IV-like mobile pole as they move from patient to patient).

Mayo chose NonStop servers to run the electronic medical record application, Carecast, supplied by GE Healthcare (formerly IDX, and before that, Phamis). Mayo started with a NonStop VLX system in 1992 and now is running a pair of 24-processor S88000s (configured from a sixteen-processor system and an eight-processor system) as an active/backup pair which they call their production system and their business continuity system. Each system has one terabyte of data storage.

The two systems are connected via a high-speed ServerNet backbone, and the backup database is kept synchronized with the production database via a GoldenGate data replication engine. The production system supports 20,000 Carecast terminals and communicates with over eighty heterogeneous interfaces to laboratory systems and other devices.

Mayo's Availability Goals

A major challenge that Mayo faces is that it must perform maintenance upgrades on the systems once per quarter. Mayo was doing this by switching over to the backup system, upgrading and testing the production system, switching back, and then upgrading the backup system. Each switchover took one to two hours. Consequently, the system was down for two to four hours during the upgrade process. This added up to more than eight hours of planned downtime per year, equivalent to not even three 9s of availability (which is eight hours of downtime per year); and this did not include any downtime due to unplanned outages.

Manual procedures were put in place so that records could be accessed manually and treatments and medications ordered during an outage. These procedures are cumbersome and are used only as a last resort.

Mayo has set a current goal of four 9s (less than one hour per year of downtime), with a longer term goal of five 9s (less than five minutes of downtime per year). Ultimately, it wants to achieve zero downtime from either planned or unplanned outages.

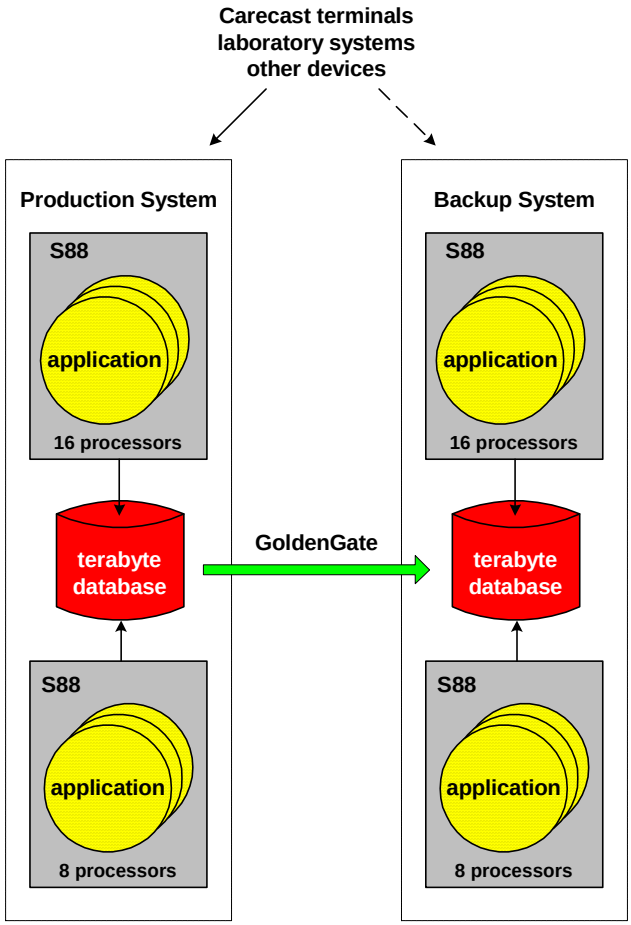

Mayo's Electronic Medical Record System

To achieve this goal, Mayo organized the SWitch Improvement Group (SWIG). This group set about understanding the reasons for long switchover times and instituted changes to correct many of the problems that were uncovered. By mid-2005, SWIG had reduced the two- to four-hour quarterly outages to 15 minutes. How was this accomplished?

Reducing Switchover Time

The SWIG group first analyzed the reasons for the excessive switchover times. As a result, it made the following changes:

- Mayo upgraded its backup system to be a full production system. In this way, Mayo did not have to switch back to complete a maintenance procedure. It now upgrades the

backup system and then switches over to it and runs it as the production system until the next quarter. By doing so, Mayo immediately cut its switchover time in half.

- SWIG automated the switchover with scripts which not only speeded the switchover but also shortened the post-switchover testing interval.

- SWIG now schedules the copying of certain files to occur before or after the switchover so that they do not have to be copied during the switchover.

- SWIG initiated the use of HP NonStop AutoSYNC from Carr Scott Software to copy certain files dynamically during the day so that these files would not have to be specially copied as part of the switchover process.

The fifteen-minute switchover time is made up of about nine minutes to bring down the current production system and about six minutes to bring online the new production system. The nine minutes required to shut down the current production system is mainly taken up with shutting down the Pathway environments. This is due to several factors:

- There are over 1,000 TCPs (terminal control processes) running.

- Several servers run in timeout loops and can't be stopped until their timeouts expire.

- There are more servers to be stopped than necessary since Mayo spawns many static servers to avoid the delays of having to spawn dynamic servers as needed. This improves response time and reduces system load but causes a longer shutdown time.

The six minutes required to bring up the new production system is largely due to:

- having to wait to start certain servers on the backup node until after the production node is down. These servers cannot otherwise be running because they make database updates on startup and would cause collisions with the primary system.

- reestablishing sessions for up to 4,000 laboratory devices and 35,000 users.

- copying over and validating nonaudited files that cannot be copied before or after the switchover or that cannot be kept in synchronism with AutoSYNC.

The Future

Reaching four nines as SWIG has done is necessary but not sufficient. SWIG's next task is to reduce downtime to five 9s, or less than five minutes per year.

Ultimately, Mayo would like to achieve 0% downtime. For this, it is beginning to consider what needs to be done to go active/active.

Summary

The quest for high availability is the quest for fast switchover. If going active/active is deemed to be too large a step to take, one must focus on minimizing switchover time.

Certainly, the extensive use of scripts, frequent training in and testing of the switchover procedures, and ensuring the availability of key people are a must. Beyond that, the factors that

are causing switchover latencies in an installation must be understood and corrected, as has been done by Sabre and Mayo.