

CU*ANSWERS HIGH AVAILABILITY PROGRAM REVIEW

EVENT DATE(S): 3/08/2020 – 3/15/2020

SUMMARY

As part of an ongoing business continuity program, CU*Answers actively maintains a high-availability (HA) core processing environment with near real-time data replication between identical hosts located at two geographically dispersed, state-of-the-art data centers. A minimum of twice each year, live HA rollover events are scheduled to redirect CU*BASE production and operations to the secondary data center (located in Yankton, SD) for a period of one full business week or longer. At the conclusion of the rollover event, core processing is redirected back, and operations resumed at the primary data center (located in Kentwood, MI). These live production HA rollover events are invaluable to validate procedures and ensure the ability to recover CU*BASE GOLD core processing in an effective and timely manner when incidents occur that threaten to disrupt business operations.

The rollover performed on March 8th through March 15th allowed recovery teams to achieve several significant goals including:

- This was the first live production rollover performed on the new HA host installed as part of the server upgrade project implemented in November 2019.
 - The server upgrade project also included changes to the IP address space at the secondary data center.
- This rollover provided the opportunity to perform system maintenance by upgrading the IBM operating system on both the PROD and HA hosts from V7R2 to V7R3.
- As part of the HA rollover strategy, alternate teams participated in both the rollover and rollback exercises to test the effectiveness of ongoing cross-training and succession planning programs.
 - This would prove to be very timely as the date of the rollover aligned with the beginning of the global COVID-19 outbreak and execution of the corporate pandemic plan to reduce exposure by minimizing the number of staff in the office at any one time.

The rollover event was completed successfully with relatively few issues. Those issues included the restoration of orphaned stored procedures in multiple CUBASEBETA library files on the HA host after the rollover, a condition not uncommon during the initial live production run on a new sever, and missing or incorrect logical record format files for MEMBER, TRANS, HTRANS, TAXADD, ECT, and FILExx libraries on the PROD host prior to the rollback due to a bug found in the replication software. Both issues were detected and corrected with minimal impact to client credit unions and members. The discovered bug in the replication software has been addressed by the vendor and a project initiated to upgrade the application prior to the next scheduled rollover.

The HA rollover event began at 3:00 AM ET on Sunday, March 8th, which coincidentally was at the same time as the switch to daylight-saving time for our time zone. At that time, recovery teams began the role-swap process to bring CU*BASE GOLD online on systems at the Site-Four data center in Yankton, SD. Production of CU*BASE GOLD

continued from the Site-Four data center until the scheduled roll-back process, beginning at 3:00 AM on Sunday, March 15th.

The following sections identify challenges observed, lessons learned, and recommendations for consideration related to this event.

EVENT DETAILS

As identified in the Summary section above, the planned HA rollover began on **Sunday, March 8th** at 3:00 AM ET and completed by 4:25 AM. Two weeks prior to the rollover, the operating system on the HA host was upgraded to V7R3. The duration of a normal HA rollover process is 45-60 minutes. This rollover was extended to 90 minutes to allow adequate time for application testing on the new operating system before enabling live transactions.

Connectivity with all third-party EFT vendor networks was restored by 4:30 AM. During each rollover process, all third-party EFT vendors function in “stand-in” mode adhering to predefined settings on transaction types and amounts (set by the credit union). Once all post-roll application testing was completed, the Operations Team performed normal EOD/BOD processing to prepare for the new business day.

Later that day, an application error was detected that led to the discovery of orphaned stored procedures within the CUBASEBETA libraries. These residual stored procedures were cleared to resolve the error messages.

During the early morning hours on **Monday, March 9th**, a small number of credit unions reported error messages in CU*BASE that led to the detection of missing or incorrect logical files (record format) for multiple libraries. A similar error was experienced during the [HA failover performed in October 28, 2019](#) that resulted in the creation of a script to identify and correct the logical files generating the errors. This same process was applied to resolve the error messages during this rollover event. Working with the vendor, teams were able to identify a bug in the third-party replication software. As a precautionary measure, the script to identify and correct logical files was also applied after the rollback process to validate the integrity of logical file record formats for all libraries.

Support teams received calls from three credit unions (of the 200+ on the network) indicating that select branch locations were not able to connect to the new HA server. CU*Answers Network Services technicians assisted credit unions in correcting network routing and configurations for the modified IP address space on the HA network.

On **Tuesday, March 10th**, with live production for CU*BASE GOLD provided from systems at the Yankton, SD data center, the PROD host in the Kentwood, MI data center was taken offline to receive an operating system upgrade to V7R3.

On the morning of **Sunday, March 15th**, beginning at 3:00 AM, recovery teams started the process to roll-back CU*BASE production to the primary data center. This was completed by 4:05 AM with all services back online.

CHALLENGES AND CONTINUING EFFORTS

Every rollover event, planned or unplanned, provides an opportunity for a valuable learning experience. Even those that appear relatively smooth on the surface, often require decisions to be made and resolutions to apply behind the scenes. Every recovery team member gives their all to minimize the impact to clients and members, while performing their job with an intense focus. Challenges observed during this rollover event include the following:

1. Planned system maintenance on the PROD and HA POWER9 hosts included an IBM operating system upgrade from V7R2 to V7R3.
 - This rollover was preceded by the IBM operating system upgrade for the QC and DEV hosts, as well as multiple self-processor hosts running CU*BASE GOLD. This allowed software developers to thoroughly test core processing applications prior to upgrading the host for CU*Answers online credit unions.
2. First live production run for the new HA host installed at the secondary data center as part of a server upgrade project.
 - By design, the server hardware for both PROD and HA are identical in build and resources and are scheduled to be replaced every 36-48 months to leverage new technology and scale to meet the demands of a growing CU*BASE credit union network.
 - Even with near-real-time data replication and high availability strategies in place, the initial live run on any new server can present unexpected issues that require a deeper level of planning, testing, and troubleshooting to detect and remediate problems, while minimizing impact to the end users.
 - This HA rollover on a new host was no exception, as orphaned stored procedures were discovered, as were default permissions on data warehouse files that are replicated manually.
3. Missing or incorrect logical (record format) files discovered in multiple CU*BASE libraries after the rollover to the HA host.
 - This is a repeat of an issue that surfaced during the [October 2019 emergency HA failover](#) where replication of logical files used by CU*BASE GOLD applications were missing or not updated as part of the data replication process. At the time, the cause of the problem was considered to have been the result of a hardware failure on the new PROD host requiring a failover to the HA host.
 - The primary difference between a “rollover” and “failover” is the ability to communicate between the two hosts. In a normal rollover, the two hosts communicate with each other and gracefully swap roles between primary and target (aka production and standby). In an unplanned failover, the two hosts are not able to communicate, requiring a forced promotion of the HA host to primary without the ability of performing system audits to ensure data integrity.
 - Working with the vendor for the replication software during this rollover event, a bug was detected in the current version of the application.
 - A new version has been developed and a project initiated to upgrade both PROD and HA prior to the next scheduled rollover. In this instance, the same script created to correct the logical files for the October 2019 failover was applied to this rollover.

4. Incorrect permissions applied to select Data Warehouse files following the HA rollover.
 - Due to the size and short retention lifecycle of daily Data Warehouse (DW) files, they are not included in the data replication strategy. Instead, they are regenerated on the HA host following a rollover.
 - These DW files are created with default permissions that are modified based on the content and access requirements of the files. Security auditing is used to detect and report permission anomalies so that the necessary corrections can be applied.
 - The correct permissions were applied to match those of the production server.

5. Slow transmission speeds for large file transfers were observed again during this HA rollover event between the GoAnywhere host and HA host. Files that had taken 5-6 minutes in the production environment were now taking 45-60 minutes or longer from the secondary data center.
 - Additional testing was performed to learn more about a potential cause. While the time required to complete the file transfers is longer, the actual impact to operations is minimal. One can expect two servers in the same data center over a high-speed core backbone to communicate more efficiently than to two servers separated by 750+ miles over a gigabit Internet VPN.
 - Modifying the buffer size on the host for these file transfers was considered. However, additional research will need to be performed to determine the global impact to system resource utilization before any changes are made in a production environment.

6. Data communications between the Site-Four PROD host and the CU*Answers HA host was dropping traffic as a result of the IP address changes implemented on the HA network.
 - This was detected Sunday morning and resolved by modifying the Extended Business Network (EBN) VPN configuration.

CLOSING REMARKS

Whether planned or unexpected, each recovery test and high-availability rollover exercise provides the opportunity to continually improve the process and adjust procedures to accommodate the ongoing changes in technology and applications that make up the core-processing environment. The value and significance of these exercises are multiplied when we consider the ever-changing threat landscape from hardware component failures, dependency on third-party vendors and supply chains, and the frequency and scope of today's natural disasters including global pandemics.

The investment made over the past two decades in building and testing its Business Continuity Management Program has positioned the CUSO to navigate the storms on the horizon and enable it to reach for new opportunities and serve its owners and client credit unions in innovative ways.

Report submitted by Jim Lawrence, CBCP | CU*Answers | Vice President of Business Continuity