

CU*ANSWERS HIGH AVAILABILITY PROGRAM REVIEW

EVENT DATE(S): 5/20/2018 – 5/23/2018

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 year, live HA rollover events are scheduled to redirect CU*BASE core processing and operations to the HA data center (located in Yankton, SD) for a period of 72 hours or more. 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 HA rollover events are invaluable in an effort to validate procedures and ensure the ability to recover CU*BASE/GOLD core processing in an effective and timely manner.

CU*Answers has been performing HA rollovers since 2003. By performing these events in a live environment, the risk of disruption is increased. However, the value of knowing in advance what to expect and how to resume operations when under the pressure of a live production environment is irreplaceable.

In this HA rollover exercise turned recovery effort, we are reminded that the controls we count on, day-in and day-out, sometimes fail. When they do, having a test-validated process to recover and restore systems, while protecting the security and integrity of the data, enables us to detect and resolve disruptive incidents more quickly.

Notable characteristics of this event include:

- During the initial rollover to the secondary data center
 - Five credit unions reported problems connecting to the HA host at the secondary data center.
 - One credit union reported problems connecting to the HA host from their Teller Cash Machines using the CFM interface due to file permissions that were not replicated.
 - File permissions on SnapShot libraries were not replicated.
- During the roll-back to the primary data center
 - Discovered that select CU MASTER table entries did not replicate properly during a brief window on Tuesday (May 22) that resulted in:
 - Transaction sequence numbers out of synchronization between the MASTER and associated TRANSx (account) tables
 - Memberships created during time window not present on primary host after the roll-back (transactional data was present but not the account information in the MASTER table)
 - Account maintenance records modified during time window not present on primary host after roll-back
 - Batch ACH process for transactions where sequence number mismatch required multiple attempts to correct (post, reverse, post) causing temporary duplications and/or reversals during the process.

EVENT DETAILS

Sunday, May 20th

The scheduled high-availability rollover began at approximately 10:00 PM ET on Sunday evening. All applications and services were back online by 10:46 PM with the exception of two third-party EFT vendor connections (ELAN and FIS/Metavante). These connections were restored by 11:30 PM after contacting and working with the associated vendor support departments. End-of-Day processing was performed normally on the production host at the secondary data center in Yankton, SD.

Monday, May 21st

On the following morning, as credit unions opened for the business day, five (5) branch locations reported problems connecting to CU*BASE at the secondary data center. Each case was resolved by the CU*Answers Network Services team working with the IT support vendors from the credit unions impacted. Resolutions centered around modifications made to host and network configurations at the credit union (firewall, DNS, routing, etc.). It should be noted that two announcements were published in the weeks prior to the rollover to alert credit unions of the scheduled event and to provide instructions for testing connectivity to the secondary data center in advance.

One credit union reported issues connecting to the secondary data center from newly installed Teller Cash Machine using the CFM integration software. This required a manual configuration change (user authority on the SSL certificate) on the host at the secondary data center to restore connectivity.

Also, on Monday morning, it was discovered that a file used by a data analytics process (Datacap1) was created using the incorrect data (48 hours old instead of previous day). This was the result of postponing specific automated jobs prior to the scheduled rollover to avoid potential delays. In essence, the job to read and parse the previous day's transactions was executed before the job to create the current data file. Once identified, the automated jobs were executed again to generate the correct data files. A programming change was implemented to confirm the date of the input data file to avoid a recurrence during future rollovers.

Lastly, after the rollover, query programs were not able to access the SnapShot libraries created on the host at the secondary data center. This was due to incorrect permissions applied to the libraries. SnapShot libraries are not replicated between primary and secondary hosts due to size and bandwidth requirements. SnapShot libraries are created on each host as needed. Since the libraries are not replicated, permission modifications must be applied manually to both hosts.

Tuesday, May 22nd

Scheduled system maintenance was performed on the secondary host at the Kentwood data center to apply operating system and firmware updates. The same updates were applied to the host at the Yankton data center several days prior to the rollover event. Later that afternoon, personnel monitoring the data replication between hosts noticed some intermittent errors being reported in the iTERA application logs. These errors are not uncommon due to the volume of transactions that occur at any given time within CU*BASE. The replication software is designed for "self-healing" and will reattempt the replication at an optimal time for performance. Shortly before 5:00 PM, the frequency of error log entries increased and teams engaged the replication software vendor to assist in diagnosing and troubleshooting the

application. The error condition was resolved by clearing and resynching two logical files that are associated with library MASTER tables. This is standard practice and was not determined to be a concern at the time. Teams continued to monitor the logs with no additional errors found.

Wednesday, May 23rd

Daily operations on Wednesday went as expected with no significant issues reported. The scheduled roll-back to the host at the Kentwood data center was started at 10:00 PM. All systems were back online by 11:00 PM at the primary production data center. All post-roll application tests were completed successfully.

Once third-party EFT vendor connections were back online, stand-in transactions started to post and random errors began to appear in the ISOPOST job log. The randomness of errors and lack of a clear scope of the cause made troubleshooting problematic (see the Challenges Observed section below for details).

Thursday, May 24th

During the early morning hours following the roll-back, recovery teams had collected information to determine that there was a mismatch in select transaction sequence numbers, preventing transactions from posting for some credit unions, and generating the ISOPOST errors. Although the scope of the condition was not clear at the time, teams utilized scripted tools to check for and correct the sequence numbers for the detected types of transactions experiencing the errors. It became apparent that trying to correct the issue while new transactions were coming in was problematic, and that normal operational processing windows offered little margin.

Teams were divided between finding and fixing the actual cause and correcting the transaction sequence numbers so that batch jobs (i.e. ACH postings) could be processed. Initial attempts at correcting the sequence numbers proved to be limited in scope (see the Challenges Observed section below for details), requiring the reversal and reposting of transactions over the course of Thursday and Friday until all transactions were completed.

Later that morning, as members began to process transactions in online/mobile banking applications, teams began to see similar transaction sequence errors. The decision was made to bring online/mobile banking offline for a brief period (starting at 8:00 AM and ending at 8:50 AM) while recovery teams worked to resolve the issue. In hindsight, had teams known the actual scope of the issue, the interruption to online/mobile banking might have been limited to only those credit unions experiencing the errors (see Challenges Observed section for details). During that time, the tools used for correcting the sequence errors was enhanced/expanded to make corrections for all transaction types.

By mid-day, reports of missing accounts for new members was reported as a new issue. It was determined that these memberships were created during a window on Tuesday (approximately 11:00 AM – 5:00 PM). All of the transaction data was stored on the system but the reference to the account in the MASTER table was missing. This expanded the scope of the issue to more than transaction sequence numbers. Equipped with the understanding of the timing when the errors occurred, teams were able to compare MASTER tables between the live production host, secondary host, and archived libraries from Tuesday (the day that the replication errors were occurring). This gave teams the information needed to determine a clear scope of the issue, the number of credit unions affected, and the multiple types of data fields within the MASTER tables that had not replicated properly.

A program was written to target the suspect/missing data field, compare the sources of data (primary, secondary, and archive), determine the accurate data value, and report the findings. Once these reports were generated, a game plan was designed to address all of the corrections to be performed that evening and into the day on Friday.

Throughout the event, alerts were posted relative to the impact to credit unions as information was available. http://alerts.cubase.org/alerts_for/cuanswers-online/. A recap of the event (attached) was sent to all credit unions on Thursday, May 31st, as a precursor to this report.

CHALLENGES

As our Information Technology systems and networks continue to grow in both size and complexity, the ability to orchestrate a perfectly-tuned environment becomes more challenging. Rather than scale back and settle for less, innovation requires pushing forward, learning along the way, while preparing for the unexpected. The controls we have implemented and strategies designed that ensure our ability to adapt and recover from adverse circumstances must be tested periodically, not just for regulatory purposes, but also for our own confidence and due diligence. The purpose of putting the controls and strategies to the test is to uncover any gaps or weaknesses. It is in this process that we gain experience and knowledge and strengthen our level of preparedness moving forward.

The following challenges were observed during this rollover event (listed in chronological order):

Sunday, May 20th

- Immediately following the rollover process, two (2) EFT vendors required a contact support call from the CU*Answers Operations Team. Elan was contacted first to recycle application services on their host to restore connectivity at 10:49 PM. The next vendor contacted was FIS/Metavante. FIS has recently consolidated support departments for the different platforms provided. This included a modification to the phone menu system, making it difficult to track down the appropriate support department. Once located, the vendor recycled services on their host to restore connectivity at 11:25 PM. Contact documentation was updated to reflect the change at the vendor.
 - Throughout the HA rollover event, all third-party EFT applications operate in “stand-in” mode at the vendor until connectivity is restored.
 - There are typically one or two EFT vendors that require contact following any rollover to recycle services on their host and restore application connectivity. Efforts to remedy this with the vendor (including communications prior to the rollover event) have been unsuccessful.

Monday, May 21st

- Five (5) credit unions reported a problem connecting to CU*BASE after rolling to the secondary data center in Yankton, SD. These reports ranged from specific PCs that were configured incorrectly to entire branches that required modifications to firewall, host, and DNS configuration changes at the credit union. CU*Answers Network Services staff worked with the credit union IT support departments and vendors to correct the issues.

- Two announcements were sent to all online credit unions (on May 2nd and again on May 15th) to alert them of the planned HA rollover event and to encourage them to test connectivity from each branch location (instructions provided in the announcement).
- One (1) credit union reported an issue communicating with the host at the secondary data center from newly installed Teller Cash Machines using the CFM interface.
 - Working with support teams at CU*Answers, the problem was resolved by modifying the user authority permissions for the SSL certificate with the CFM interface on the secondary host. This is the first HA rollover since the integration with the CFM interface. All future integrations will require manual permissions modification to both the primary and secondary hosts.
 - As more credit unions integrate with the new CFM interface, it will be imperative that host-based configurations are performed on systems at both the primary and secondary data centers.
- During normal nightly processing, after the SnapShot library files are generated for each credit union, an automated job is performed (Datacap1) to build a file from the newly created libraries containing information for data analytics queries and dashboards. Due to the planned HA rollover, the creation of the SnapShot library files was delayed, resulting in the use of stale data for the data analytics file.
 - The job was re-run after the creation of the SnapShot libraries and the process to update the data analytics file completed. A programming change has been made to confirm the date of the SnapShot library prior to file creation to avoid a reoccurrence during future rollover events.
 - In addition, pre and post-rollover run sheets have been modified to inform operators of the proper sequence in which suspended jobs are to be executed.
- Query programs were not able to access the SnapShot libraries created on the host at the secondary data center after the rollover.
 - This was due to permissions applied to the libraries. SnapShot libraries are not replicated between primary and secondary hosts. SnapShot libraries are created on the active primary production host as needed. Since the libraries are not replicated, permissions changes must be applied manually to both hosts.
 - A post-rollover check has been added to the Operations run sheet to check file permissions after the Snapshot libraries have been generated.

Tuesday, May 22nd

- There were no known issues reported related to the HA rollover on Tuesday. However, intermittent replication errors began to appear on the secondary/TARGET host during the afternoon hours. As noted in the Event Details above, these errors are not unusual and typically are resolved by the self-healing characteristics of the data replication software. Shortly before 5:00 PM, the vendor was contacted to diagnose and resolve the error conditions. This was accomplished by clearing and resynching two logical files on the secondary host. This step was successful in clearing the errors that were occurring at the time.
 - Logical files do not contain data. They are used by applications on the IBM host to provide descriptions of the contents within one or more physical files (record definition, data mapping, etc.). Logical files are loaded at run time in the operating system.

Wednesday, May 23rd

- There were no known issues reported related to the HA rollover on Wednesday during normal production hours. Following the HA roll-back, Operations Team members noted errors in the log file for the ISOPOST application after third-party EFT vendor subsystems were brought online at the primary data center.
 - The error messages were the first indication that a problem was present on the primary host following a successful rollover. During the next 12-24 hours, recovery teams would uncover an intermittent and random condition that had occurred back on Tuesday afternoon that resulted in data replication errors while operating on the secondary host. The process to identify, track, isolate, and correct those errors proved to be a lengthy process; one where the integrity of the data took precedence over the timing of the recovery effort (see next paragraph for details).

Thursday, May 24th

- Shortly after midnight, recovery teams were able to determine that the cause of the ISOPOST error messages was due to a mismatch in transaction sequence numbers for a growing number of applications. The root cause of the mismatch was unknown at the time.
 - A script (tool) was developed to resynch the transaction sequence numbers based on the information at hand. It was later determined that the range of the types of transactions included in the original version of the correction tool (specifically TRANS1) was insufficient and that other types of transactions (TRANS2 and TRANS3) were also experiencing similar sequence errors.
 - With the issue still ongoing, at 8:00 AM, a decision was made to suspend the online/mobile banking application (as a precautionary measure) to prevent additional transactional errors and to give recovery teams the opportunity to uncover the full scope of the problem at hand. The online/mobile banking applications were restored at approximately 8:50 AM.
 - Had teams known the actual scope (specific credit unions impacted) at the time, the option to bring online/mobile banking offline for only those credit unions could have been executed, instead of impacting all credit unions.
- During the initial ACH batch-file postings that morning, several errors were generated. The decision was made to reverse the ACH postings and troubleshoot the errors for resolution. The posting and reversal was performed multiple times until the process was free of errors.
 - This created some confusion and concern for the credit unions affected (14) as transactions seemed to appear and disappear without explanation as files were posted then reversed.
 - Alerts were issued informing the credit union not to attempt manual corrections, and that the problem was being addressed.
- During this period, it was learned that transaction sequence numbers were not the only records out of synch between hosts following the rollover. It was reported that a small number of memberships created were missing, as were some membership information maintenance records (address changes, etc.). Recovery teams sought to determine if these incidents were related, and if so, determine the common characteristics.
 - With this new information in hand, it became clear that the missing information occurred during a specific window of time on Tuesday, while operating from the secondary data center.
 - This was confirmed by running a data query to compare data in production, on the secondary host, and archived data from the end of day on Tuesday, May 22nd. At that point, teams were able to identify all of the records that were not synchronized correctly between the primary and secondary hosts.

- A plan was created to divide the records by type and implement the solution to correct them (using an automated restore process where possible, or manual methods where required).

Friday, May 25th

Final edits and corrections were made by teams working with those credit unions affected during the rollover and recovery event to verify G/L balances and process 29 DFU's for member account issues.

CONTINUING EFFORTS AND RECOMMENDATIONS

Whether planned or unexpected, each high availability rollover and recovery test provides the opportunity to continually improve the process and adjust procedures. It allows us to test the technologies in place, to identify where they are fallible, and design the necessary tools and skills to resume operations quickly.

The following is a list of action items and projects relative to this rollover/recovery event that will be pursued in an attempt to draw closer to that goal:

1. Create a more distinct announcement for communicating to credit unions about scheduled HA rollover events and the implications for not testing connectivity to the secondary data center.
 - On the Monday following the initial rollover, five credit unions (out of 160+) reported not being able to connect to the secondary data center. While a small percentage overall, the impact to those affected can be significant, and in most cases, avoided. For future rollover events, announcements will be created that look less like standard alerts or marketing materials and emphasize the impact of not performing the pre-event connectivity testing.
2. Improve communications during a rollover/recovery effort between teams by assigning the role and responsibility for each scheduled event.
 - During previous rollover exercises, once the post-roll application tests were completed, daily processing was turned over to the Operations Team with instructions to notify the appropriate recovery team if any issues should surface. While this has proven successful for most events, on occasion, there are those rare instances that require a broader escalation of teams and skills to see the big picture.
 - This rollover event was one of those rare instances. While the experts were focused on troubleshooting this issue at hand, having a team member dedicated to keeping others informed and knowing when to engage additional support may have accelerated the time it took to accurately identify the scope of the issue, rather than chasing symptoms as they emerged.
 - Future rollover events will have this communications role defined for participation until the "all clear" sign is given.
3. Improve the communications to client credit unions regarding issues and disruptions that surface during a rollover/recovery effort.
 - Throughout the duration of the rollover/recovery event, alerts were posted for client credit unions as information became available. These alerts provided minimal information and were intended to communicate that teams were aware of the multiple issues and were working to resolve them. While it is always our goal to be transparent and communicate promptly, it is important to remember that incomplete information or even too much information can sometimes be worse than none at all. This process was complicated by the fact that the cause

and true scope of the issues were not verified until mid-day. There were additional alerts posted for disruptions that were not related to the rollover, adding to the confusion by client credit unions.

- This rollover/recovery effort will be used as a case study to find the gaps and weaknesses in communications with key internal and external parties and enhance the process for future events. All alerts and correspondence will be examined during a post-event debriefing and recommendations made to the Executive Counsel no later than June 30, 2018. Procedure improvements will be included in the July revision of the Business Continuity Plan.
4. Expand the decision-making group when errors or problematic conditions occur during a rollover event.
- For example, during this rollover, replication errors were reported and the vendor engaged on Tuesday afternoon. This information was only communicated to a small group who were most familiar with the application. Especially during rollover events, it's important to report and communicate to a broad group and challenge any assumptions that might lead to issues prolonging the roll-back or (in this case) recovery process.
5. Expand the range of the pre-roll audits performed prior to the rollover and roll-back process to mitigate the risk of reoccurrence.
- The standard schedule for a HA rollover event has been Sunday through Wednesday. This gives teams the experience of performing the rollover event on both a weekend and weekday. With lower traffic volumes observed on a typical Sunday, additional audits are performed to ensure the integrity of data between hosts at the primary and secondary data center. With higher levels of traffic and sheer number of transactions observed on a business day, a truncated list of audits is typically performed prior to the roll-back. This has been the standard process for the past ten years of performing HA rollovers.
 - In response to this rollover event and the issues observed with the replication process, time-consuming audits will be modified (broken into smaller processes) to focus on key data areas in an attempt to more accurately detect a mismatch between hosts. In addition, additional system tools will be developed within CU*BASE to detect a similar mismatch in key data fields.

CLIENT COMMUNICATIONS

The following announcement was issues to all credit unions on May 31, 2018 as a recap of the event.

ATTENTION ONLINE CREDIT UNIONS

As you know, last week was a hectic one with various issues, some of which were related to the HA rollover. As in the past, our teams are already pulling together all the details to be compiled for the official HA Rollover Gap Analysis, which includes a complete timeline showing when things happened and how they played out over the week. That document will be reviewed by our Board later this month then sent directly to all CUs. In the meantime, below is a quick recap of the key issues related to the rollover.

CU*Answers HA Rollover Test Timeline

- The rollover from the Production to the High Availability server was processed between approx. 10:00p ET and 11:30p ET Sunday night, May 20 (most services were restored by 10:46p ET).
- Rollback from HA to Prod was processed between approx. 10:00p ET and 11:00p ET on Wednesday night, May 23.

Data Replication and the HA Rollover Test

Now that the dust has settled, the symptoms noticed by our teams and our CUs on Thursday and Friday eventually pointed back to an error earlier in the week during the normal replication process where data from the MASTER table was not replicated properly between the production and backup server. During the brief period while the error was occurring, data for any CUs for whom activity was being logged would not have been replicated properly to the backup server (meaning it was either logged incorrectly or was missing altogether).

The MASTER table is where **general membership data** is stored, including brand-new memberships, addresses or other general info for existing members, and even passwords for audio response (online banking passwords are stored in a different table). If changes made to this data were not replicated to the backup server, then those changes also weren't there after rolling back to that server.

The Heart of the Matter: The Trans Sequence

Even more importantly, MASTER also stores a **trans sequence #** that is used by all posting programs to assign sequential numbers to all transactions on that member's sub-accounts. If the MASTER trans seq # is incorrect or missing, posting programs can't figure out how to post the transaction and everything comes to an abrupt halt. With a batch program such as ACH, we might need to reverse what got posted before the halt then process again. Other programs such as teller or account adjustments require a different response, but the net effect is the same: strange behavior where it looks like transactions didn't post, or they did post but then seemed to disappear or were reversed, etc.

Our First Response

Missing data is in some respects easier because you can detect what's *not* there more easily than data that is *there* but is just slightly *different*. Our logs showed the error in the MASTER tables, and we did take the recommended steps to adjust and resynch the data. However, as we later discovered, that process was imperfect and trans sequence numbers were still out of synch on some MASTER records when we rolled back. This started causing ripple effects as soon as a posting process tried to post transactions to the affected accounts.

Summary of Reported Symptoms

This list summarizes specific anomalies reported either by our teams or by credit unions on Thursday and Friday:

Symptom: Memberships created during the rollover period were not visible

- *Impact:* Less than 20 accounts across all CU*Answers online CUs.
- *Status:* All data was resynched and recovered; issue resolved as of Thursday afternoon, 5/24. The affected CUs have been contacted.

Symptom: Changes made to MASTER membership info (address, phone numbers, etc.) were not visible

- *Impact:* The total varied by CU but less than 120 CU*Answers online CUs saw at least a few memberships affected.
- *Status:* All data was resynched and recovered; issue resolved as of Friday afternoon, 5/25. The affected CUs have been contacted.
Note to Internal Auditors: You may notice changes in your CUFMAINT log performed by CU*Answers data center employees as a result of this repair work.

Symptom: Transactions not appearing, or being reversed/duplicated

- *Impact:* Affected some accounts at 14 CU*Answers online CUs, for transactions done via posting processes such as ACH, share draft processing, account adjustments, and teller, along with some related GL out-of-balance conditions.
- *Status:* All issues affecting members directly were repaired by Friday afternoon, 5/25. Client Services is still working on a few behind-the-scenes running balance issues and G/L balancing issues. These will be completed soon, no later than the end of this month so that everything is correct for month-end. **The affected CUs have been contacted.**

*Symptom: Changes to CU*Talk Audio Response passwords did not "take"*

- *Impact:* Less than 259 memberships across 62 CUs.
- *Status:* Decision was made NOT to attempt to repair this programmatically. New passwords will not work; members will either need to change their password again or contact the CU for a reset. **Any concerns or questions, don't hesitate to contact a CSR.**

*Service Interruption for **It's Me 247** and CU*Talk on Thursday*

- *What:* **It's Me 247** online/mobile web banking and CU*Talk audio response were offline between 4:30 AM and 8:50 AM ET on Thursday, May 24. This was a precautionary measure while we investigated reports of missing transaction data, which were later found to be related to the transaction sequencing issue described above.
- *Impact:* All CU*Answers online CUs.

- *Status:* Both systems were brought back online as of 8:50 AM ET on Thursday morning, 5/24.

Sidebar: SnapShot and the Recovery Process

As a sidebar, one of the tools we used in the recovery and re-synching process was the new daily SnapShot library (FILExxS). As you've seen in recent communications from Asterisk Intelligence, the [SnapShot library](#) gave us an additional picture of production data from EOD prior to the rollback. This proved to be very useful as we took care of the various transaction posting issues, since it helped us avoid the more time-consuming process of restoring from tape backup.

Keeping You in the Loop

Communicating during events like this is more art than science. Usually it's only in hindsight that things become clear enough to explain them, and it's important that we not compound the confusion by reporting incorrect or misleading details. Some people want us just to "tell us when it's fixed" while others want to feel like they are getting the inside scoop on all the nitty gritty. Remember also that at the time symptoms first started appearing, it was not immediately obvious what the root cause was, or even that the symptoms were all necessarily related to the rollover. It is always our goal to be transparent, but incomplete information or even too much information can sometimes be worse than none at all. We appreciate your patience as we continue refining that art.

A Business Continuity Process You Can Believe In

The high availability process requires that we continually test our ability to keep the heartbeats of data between two boxes in sync to be ready for emergencies and eventualities of *forced* recoveries. In a real disaster, are we ready to deal with the inevitable chaos? In the end, this event – what we learned and what we will change about our processes – is the foundation for our continued success in the future. Be confident that our ongoing work toward this state of readiness is the correct course of action for our network.