

CU*ANSWERS HIGH AVAILABILITY PROGRAM REVIEW

EVENT DATE: 11/27/2013 – 12/3/2013

SUMMARY

As part of a robust business continuity program, CU*Answers actively maintains a high availability (HA) core-processing environment with real-time CU*BASE/GOLD data replication between identical hosts located at the state-of-the-art primary and secondary datacenters. A minimum of twice each year, an event is scheduled to rollover core-processing and operations to the secondary datacenter (located in Muskegon, MI) for a minimum period of three business days. At the completion of each event, core-processing is rolled back and operations resumed at the primary datacenter (located in Kentwood, MI). These events are invaluable in our effort to validate our procedures and ensure our ability of recovering CU*BASE/GOLD core processing in an effective and timely manner.

Case in point: An unplanned emergency high availability (HA) rollover was conducted on the morning of Wednesday, November 27, 2013 to bring CU*BASE/GOLD core-processing online at the secondary datacenter. This was done to replace a failed hardware component on the production host at the primary datacenter (located in Kentwood, MI). Once system repairs had been completed, a rollover event was scheduled on the evening of Tuesday, December 3, to bring CU*BASE/GOLD back online at the primary datacenter, returning to normal daily operations.

Notable characteristics of this event include:

- Demonstration of our confidence in our ability to quickly and effectively bring CU*BASE/GOLD online at the secondary datacenter on short-notice, preventing a potential disruption from escalating. (With redundant components on the systems, a rollover was not mandatory. Our level of preparedness gave us the option to bring up core-processing at the secondary datacenter in a controlled manner, mitigating the risk of a disruption during peak business hours.)
- Operating from systems at the secondary datacenter throughout the extended Thanksgiving holiday weekend (with increased activity during what may be the busiest shopping period of the year)
- Performing month-end processing on systems at the secondary datacenter

This unplanned event demonstrates the value of the investment in high-availability technology resources and rigorous hours of planning and testing to best prepare us for recovering CU*BASE/GOLD core-processing in a timely and effective manner. This report identifies the scenario and circumstances leading up to the decision to perform the emergency rollover as well as any challenges observed, lessons learned, and recommendations for consideration related to this event.

EVENT DETAILS

At approximately 4:55 AM ET on Wednesday, November 27, Operations personnel observed a system message identifying a potential hardware issue. An iSeries Administrator was notified immediately for further investigation. The messages received pointed to a hardware failure for a component on a redundant, hot swappable, pair of RAID array controller boards. Although the system was functioning properly on the remaining controller board, the potential for performance degradation was present.

Recovery teams were contacted and put on alert status to be ready in case further action was required. The system vendor (IBM) was contacted to locate a replacement part and estimated arrival time for installation. A part was located in the Chicago area and placed on a courier for delivery. Given the severe winter weather in that area at the time, the ETA could not be guaranteed.

At 6:10 AM, the decision was made to prepare for failing over core-processing to the HA system at the secondary datacenter. Many factors were considered in this decision:

- **The timing of the event.** Had this occurred a few hours earlier, we may have elected to wait for the replacement part and install it on the production system. Given the proximity to the opening of the business day, the decision was made to move core-processing to the secondary datacenter.
- **The potential for performance degradation.** With only one controller board, the system was not in its optimal state. On the Wednesday before Thanksgiving, the possibility of increased network and system traffic was present. Given the circumstances of this incident, the path of caution seemed like the favorable route to take.
- **Our history of performing regular rollover exercises.** Having skilled recovery personnel and test validated procedures, we knew precisely what needed to be done and how long it would take to bring core-processing online at the secondary datacenter. We were confident this could be accomplished before the start of the business day.

Recovery teams began preparing systems for bringing up production at the secondary datacenter. At 7:00 AM, the rollover procedures were initiated. The event unfolded as follows:

- 7:05 AM - CU*BASE application taken offline along all subsystems and services on CUAPROD.
- 7:30 AM - CU*BASE application and all subsystems and services brought online at the secondary datacenter on CUAHA1.
- 7:35 AM - Recovery teams begin system checks and troubleshooting issues as they are uncovered
- 7:40 AM - Application checks confirmed (GOLD, ItsMe247, CU*Talk, CU*Spy, etc.)
- 7:41 AM - Client logins to CU*BASE/GOLD enabled on CUAHA1
- 7:55 AM - All third party vendor communications confirmed with one exception (FIS)
- 8:00 AM - FIS operations support contacted to refresh communications from the vendor side. This corrected the issue. Communication is confirmed at 8:20 AM.
- 8:15 AM – All ACH files posted (alert sent to clients about the delay)
- 1:00 PM – New hardware component arrives and installed on CUAPROD.

Support calls were attended to as credit unions opened for business. Systems were monitored closely during the extended holiday weekend while Operations performed their regular daily duties from the secondary datacenter. A tentative time for rolling production back to the primary datacenter was set for Tuesday, December 3, at 10:00 PM ET to give the replacement component adequate “burn-in” time and allow administrators to monitor and test

the system before bringing it back online. This extended period also allowed the continuous availability of CU*BASE/GOLD without interruption throughout the busy holiday weekend.

On Tuesday, December 3, at 10:00 PM, the roll-back process was initiated bringing CU*BASE/GOLD core processing back to the primary datacenter. This process was completed and systems back online by 10:45 PM. At 11:00, problems were discovered in application log files indicating possible errors in files for three libraries. These errors were corrected by the Application Recovery Team while others looked for the potential cause. Additional testing was performed to confirm integrity of the file system and determine if rolling back to the secondary datacenter was necessary. These tests produced no new problems and at 11:40 the decision was made to remain on the system at the primary production datacenter.

End of Day (EOD) processing was completed successfully without errors or issues, paving the way for Beginning of Day (BOD) operations on the primary production system.

CHALLENGES

Having performed dozens of core-processing rollover exercises in a live, production environment with a veteran recovery staff, the “*element of surprise*” can still catch us slightly off guard. Looking back, there were steps we could have taken to perfect our response to this incident. We are quickly reminded, however, what could have been had we stopped to fine-tune our efforts and the potential system-wide disruption that was avoided as a result. We’ll take that trade-off every time and continue to learn from our experiences.

The following challenges were observed during this incident:

1. There was an initial delay in the response from the hardware support vendor in locating a replacement part and providing an ETA on its arrival and installation. In this incident, the replacement component was part of a redundant configuration (stand-by component was still functioning). There appeared to be a slight disconnect in the sense of urgency applied to this support call based on the observed behavior of the vendor. We had to follow up with the vendor on more than one occasion to retrieve the desired information. In this case the delayed response did not result in an impact to our recovery effort, but under different circumstances, it might have.
 - a. We will follow up with the vendor representative to ensure future support needs are handled promptly.
2. Crisis communications is always a key ingredient in any timely and effective recovery effort. While alerts and announcements were issued to clients and internal staff, there were some (non-client-facing) internal teams who were not aware of the status of the incident nor their specific roles and responsibilities during the recovery effort.
 - a. In response to this challenge, the Crisis Communication Team will meet to discuss and identify appropriate procedures and determine how all internal teams can take part in the rollover exercises (expand the scope of the simulated disruption).
3. On Saturday, November 30, a weekly procedure was performed to “clean-up” inactive and orphaned USER accounts. This procedure inadvertently deleted profiles for some users whose last logon date to CUAHA1 exceeded the 92-day criteria imposed by the procedure. USER login history is maintained separately on each system (not replicated) for auditing purposes. Upon being reported, System Administrators quickly restored the USER profiles on the host.

- a. To prevent this from reoccurring, Operational run sheets have been modified to omit this and similar procedures (relative to USER profile history) from being performed during a rollover event. In addition, we will consider staggering the days of the week that we perform regular rollover events (typically scheduled for Sunday through Wednesday) to help identify potential issues from processes and procedures that are performed on a weekly or monthly basis.
4. During previous rollover events, there have been a limited number of client sites that were unable to connect to systems at the secondary datacenter due to non-standard network configurations at the credit union. During this incident, there were three sites unable to connect to the secondary datacenter. One of them is a pre-conversion site that is currently configured for training.
 - a. Instructions have been developed and published (<http://rollovers.cuanswers.com>) to help clients test their configurations prior to a rollover event to confirm connectivity to the secondary datacenter. To get this message out to the network, a reminder will be sent out on a quarterly basis to all clients.
5. Since the last rollover event (September, 2013) there had been one new third party vendor added to the network (with one credit union client). Procedures for testing connectivity for this new vendor were not added to the HA-Post-Roll documentation. As a result, communications for this one vendor were offline until identified later in the day (Wednesday, November 25, at approximately 1:30 PM ET).
 - a. Procedures have been modified for adding new vendors to the network to confirm that the appropriate documentation is updated.
6. During the rollover back to the primary production system, we experienced file structure errors in three libraries, even though all pre-roll auditing confirmed no errors. Recovery teams were able to recover/rebuild all affected transaction files quickly to restore the libraries. The days following the roll back, a small number of transactions files have been discovered with similar structural errors and repaired.
 - a. A support case has been opened with the data replication vendor to investigate the cause. If a fix is required, a test will be performed to confirm the results. This will be monitored closely during the next scheduled rollover event, tentatively planned for February/March 2014.

CONTINUING EFFORTS AND RECOMMENDATIONS

Whether planned or unexpected, each recovery test and high availability rollover exercise provides us the opportunity to continually seek to improve the process and adjust our procedures. The best way to accomplish this is to “Practice. Learn. Repeat”. The following is a list of action items and projects that we are pursuing to get us closer to that goal:

1. Expand the scope of our scheduled rollover exercises to engage all teams to sharpen our communications and procedures in preparation for a potential large-scale recovery effort.
2. Design an improved communications campaign (in addition to the standard event announcements) to better educate clients on the requirements and importance of complying with standard network configuration settings for connecting to CU*BASE/GOLD.
3. Develop a regular auditing process to ensure all documented recovery (HA/DR) procedures are current and accurate.