CU*ANSWERS HIGH AVAILABILITY PROGRAM REVIEW

EVENT DATE: 4/28/2016 – 5/18/2016

SUMMARY

As part of an ongoing 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 two geographically dispersed, state-of-the-art datacenters. A minimum of twice each year, HA rollover events are scheduled to redirect core-processing and operations to the secondary datacenter (located in Muskegon, MI) for a minimum period of 72 hours. At the completion of each event, core-processing is redirected back and operations resumed at the primary datacenter (located in Kentwood, MI). These HA rollover 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.

Procedures followed during this rollover exercise were similar to the previous event, performed on September 13-23, 2015. Each of these rollover exercises brings with it a unique set of circumstances and challenges, but common among them are the goals and objectives of a successful continuity and recovery program.

Notable characteristics of this event include:

- This rollover was planned to perform necessary system diagnostics and maintenance as a result of three brief intermittent disruptions to CU*BASE that started back in February.
- This rollover was originally planned for April 17-20 but postponed due to symptoms of a similar condition (believed to be responsible for the CU*BASE disruptions) detected on the HA host.
- This rollover was performed just hours following a CU*BASE interruption that occurred near the close of the business day on April 18.
    - Typical rollovers are scheduled to begin at 10:00 PM ET. This rollover was scheduled for 9:00 PM ET to allow recovery teams an additional hour for troubleshooting any residual effects of the disruption earlier in the day on PROD.
    - Due to the timing of the disruption toward the close of the business day, the announcement of the HA rollover was not issued nor alert posted until 5:44 PM ET. As a result, some credit unions may not have been aware of the event until the open of the next business day, at which time CU*BASE core-processing was already being provided from our secondary datacenter.
- Typical rollover periods are scheduled for three business days. This rollover event was scheduled without a known rollback date (to be determined by the results of the diagnostics testing and the speed at which replacement components could be installed). The total duration of this rollover was 20 days, including a CU*BASE/GOLD 16.05 software release and End-of-Month processing from the secondary datacenter.

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

EVENT DETAILS

On Thursday, April 28, at 9:00 PM ET, the production host was taken offline and rollover procedures initiated. By 10:40 PM, recovery teams began testing core-processing applications on systems at the secondary datacenter. At 11:00 PM, all applications were confirmed and CU*BASE/GOLD brought online.

Once production was directed to systems at the secondary datacenter, teams began to diagnose and repair PROD at the primary datacenter. This included the replacement of 85 Toshiba hard drives that were determined to have faulty firmware responsible for generating a higher than normal amount of disk errors. Utilizing hot-swappable technologies allowed us to perform this exhaustive maintenance task while the system remained online without interruption to data replication. Details about this process are included elsewhere in this report.

On Wednesday, May 18, at 10:00 PM ET, the rollback process was initiated bringing CU*BASE/GOLD core-processing back to the primary datacenter. This process was completed and systems back online by 11:05 PM.

CHALLENGES

As noted above, the circumstances leading to the decision for this rollover event can be traced back to a CU*BASE disruption that occurred on the morning of Saturday, April 9, where GOLD sessions were dropped and several users reported having difficulty logging back in for approximately one hour. This event, similar to an earlier disruption back on February 3, escalated the effort for deeper diagnostics and testing to determine the root cause.

To accommodate this level of system maintenance, a rollover event was scheduled for Sunday, April 17. On the evening of Friday, April 15, system logs on the HA host indicated that the problem may not be isolated to the PROD host. Rather than roll production to the HA host, the decision was made to postpone the rollover event and perform additional testing to confirm system reliability. Before the rollover event could be rescheduled, a third CU*BASE disruption occurred on Thursday, April 28. This prompted the decision for an emergency rollover later that evening.

This third disruption provided the clues needed for the system manufacturer (IBM) to determine cause and resolution. This disruption resulted in damaged objects (files) on the disk drives requiring replacement of the affected drives and restoration of the damaged objects on the PROD system.

Assessing the scenario, recovery teams presented two options for performing the required repairs. The first option (preferred) included replacing the disk drives one at a time, allowing the system to rebuild each one as part of the RAID technology (Redundant Array of Independent Disks). This option provides a much slower restoration process but would allow us to maintain full data replication throughout the process. The second option included bringing the stand-by system down, pulling all of the faulty disk drives (Toshiba), replacing them with new disk drives (Seagate), and restoring the system through backup media and data replication. This option could provide a more timely recovery process, however, it would require a significant period of time (multiple days) without full data replication for CU*BASE.

Complicating the decision was the potential for uncovering additional damaged objects that could occur as a result of the rebuilding process for the new disks while reading data from the existing disks. The decision was made to proceed with the first option with every effort made to ensure the protection of data. Teams worked multiple shifts starting on May 5 and ending on May 11, replacing a total of 85 disk drives without interruption to CU*BASE core-processing, all while maintaining full data replication.

With the scheduled GOLD 16.05 software release on May 15, the decision was made to postpone rolling back to PROD (with 85 new disk drives) until May 18. The process to replace 96 Toshiba disk drives on the HA host began on May 23 and concluded on May 28.

In addition to the challenges and issues listed above, the following incidents occurred during the rollover event (not related to faulty disk drives):

- Daily security reports
    - On the morning of April 29, following the rollover, it was noted that daily CU*BASE security reports were not delivered to our internal compliance teams.
        - Since the last rollover event, new email relay servers had been installed and not configured to receive messages from the host at the secondary datacenter. Once this configuration was corrected, reports were delivered once again.
- Membership application disruption
    - On May 2, it was discovered that online membership applications were not being received for processing.
        - The servers that host the membership application service (MAP/MOP) were installed since the last HA rollover event. It was determined that firewalls installed at the secondary datacenter were not configured to allow network traffic from these new servers. This has been corrected.
- CU*BASE user configuration
    - On the morning of April 29, immediately following the rollover, system security controls detected a small group of user accounts that had been disabled on the PROD host but enabled on the HA host.
        - One of the processes normally performed during a rollover event involves auditing these select user accounts. Due to the circumstances and urgency of this particular rollover event, this process was not performed. Documented procedures have been modified to ensure compliance for future rollovers.
        - This incident represents a great example of the value of and need for layered security controls. Where one control failed (missed procedure), another control detected the potential problem and alerted the appropriate teams for prompt response.

CONTINUING EFFORTS AND RECOMMENDATIONS

The first two of the three challenges itemized above are related to new application servers installed at the primary datacenter that were not configured for network access to the secondary datacenter. This needs to become part of the implementation process; not only planning for normal daily operations but anticipating and testing access that may be required during disruptive scenarios. This is one of the many reasons we perform regular rollover events to help uncover these misconfigured environments and correct them.

The third challenge listed is related to host-specific data that is not part of the normal replication process. This data is part of a manual (sometimes automated) synchronization between hosts. This is a relatively rare omission as a result of the urgency of the circumstances. The controls in place helped to mitigate any risk (as intended).

Whether planned or unexpected, each recovery test and high-availability rollover exercise provides us the opportunity to continually 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 relative to this rollover event that we are pursuing to get us closer to that goal:

1. Review and revise the documented process for installing and configuring new application servers to ensure access to all required hosts at both the primary and secondary datacenters.
    a. This will become even more paramount later in 2016 when the HA host is relocated to the Site-Four datacenter in South Dakota.
2. Review and revise the documented procedures for synchronizing host specific data between rollover events.
3. Celebrate our success, then press forward
    a. When we step back and evaluate all of the planning and effort invested over the years to achieve the level of preparedness the CU*Answers network has obtained through rigorous HA rollover exercises every six months on 'live' production systems, we gain the confidence needed to continue building the networks and systems that will meet the demands of tomorrow.

As we assess this HA rollover event, we expected early on that a hardware component was the major contributor to the disruptions we were experiencing. The manufacturer (IBM) performed extensive testing to exclude the more common culprits (controller boards, drivers, etc.) and narrow our focus. What we did not expect was having to replace a total of 181 disk drives (85 on PROD and 96 on HA) in a live production environment without incurring additional downtime. We are convinced that all of our investment and preparation over the years were necessary to put us in a position to pull off this recovery.

In closing, the two systems involved (PROD and HA) are scheduled to be replaced this August and September as part of a three-year lease agreement.

Report submitted by: Jim Lawrence, CBCP | CU*Answers | Manager of Business Continuity and Recovery Services