

CU*ANSWERS HIGH AVAILABILITY PROGRAM REVIEW

EVENT DATE(S): 10/06/2019, 10/27 – 11/5

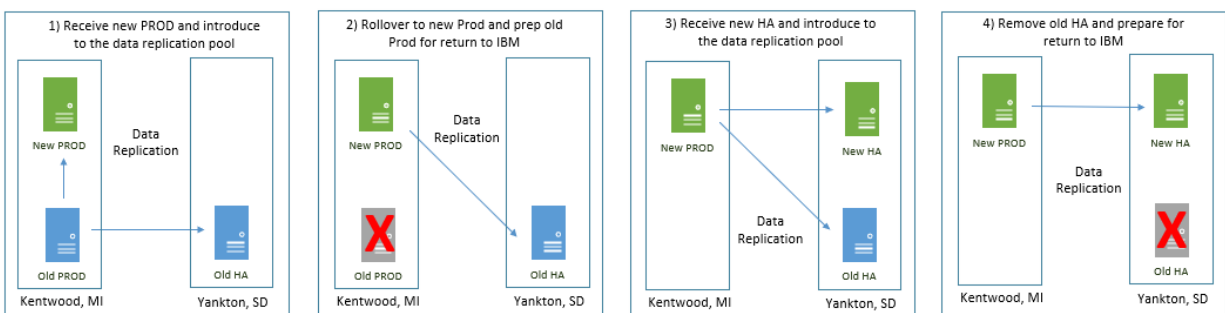
EVENT SUMMARY

All times included in this report are Eastern Time.

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 core processing production and operations to the HA 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.

This report details the scheduled HA rollover performed on October 6th and the unplanned, emergency HA failover conducted on October 27th, putting the capabilities and level of preparedness of the business continuity program to the test. Due to the nature of the emergency failover process, the Event Summary section has been expanded to provide the background between each phase of the event as it unfolded. Additional details are included in the Event Timeline.

Every 36-48 months, the server hardware that hosts the CU*BASE application suite is replaced with new technology to meet the demands of a growing user base and changing application performance requirements. During the month of October, the execution phase of the project was initiated to upgrade both the production (PROD) and high availability (HA) servers to a new hardware platform using the familiar “rollover” process.



This process, shown in the four steps above, involves receiving the new hardware from the manufacturer under a lease agreement, installing and configuring the operating system and applications using a comprehensive system build checklist, introducing the server to the live data replication pool, performing pre-roll tests and audits, then conducting the rollover to bring the new PROD server online as the CU*BASE production host. This process is

repeated for the HA host. Once completed, the old hardware is sanitized per the data security policy and returned to the manufacturer.

As part of this project, the new server hardware was received in early September and prepared for implementation through a series of comprehensive system and application testing from multiple teams. On the morning of **October 6, beginning at 1:00 AM**, the rollover process began to bring CU*BASE core-processing online on the new PROD host. To minimize the impact to end users and other applications that communicate with CU*BASE, the systems were renamed and network addresses renumbered so that the new host mirrored the settings of old host.

The initial rollover for any new system requires a complete set of audits to initialize the data and prepare it to be the source for the replication pool. These additional steps, along with comprehensive post-roll testing to ensure data integrity, required an extended maintenance window (4 ½ hours) compared to most typical HA rollover events (1 hour). By **5:30 AM**, CU*BASE core-processing was back online, running on the new hardware. For the next few hours, teams worked to identify and correct any application issues that surfaced during the initial EOD/BOD daily processing. For the next two weeks, the new PROD host performed well without any issues.

On **Wednesday, October 9th**, a hardware failure occurred on a redundant voltage regulator. Manufacturer support (IBM) was contacted and a replacement part dispatched with a system maintenance window scheduled for **1:30 AM to 3:00 AM** the following morning. With the new replacement part installed, the PROD server was operating once again.

On the morning of **Sunday, October 27th**, a maintenance window was scheduled from **5:00 AM to 6:00 AM** to power down the new PROD host and move it from the temporary staging area to the permanent location in the computer room. The PROD host is composed of multiple processing units and disk arrays, all contained within a six foot tall rack enclosure on casters. The rack enclosure was moved a total of three feet across a raised floor in the data center. Upon powering on the new PROD host, a system error message was received indicating a hardware failure, preventing the server from booting up. Once again, IBM support was engaged. Diagnostic tools determined that one of the CPUs had failed. A new CPU module board (containing eight CPUs) was dispatched, as was on-site technical support for installation. After installing the replacement CPU module, a new error message was received prompting the dispatch of additional parts. (This process was repeated until the host was finally recovered at 4:00 AM on Tuesday morning, a total of 46 hours from the initial CPU failure).

While the hardware recovery team was working with IBM support to restore the new PROD host, the rest of the team was engaged to perform an emergency HA failover to bring CU*BASE core processing online on systems at the Yankton, SD data center.

A “**failover**” is required when communication is lost with the PROD host. Under normal conditions, a “**rollover**” is performed to gracefully swap roles between the PROD and HA hosts, ensuring that all processing queues are transferred and replication reversed. Under adverse conditions where communications to the PROD host are lost, the role swap process is forced and manual steps are necessary to clean up any unfinished processing and troubleshoot any issues that are generated.

The emergency HA failover process began on **Sunday, October 27 at 8:30 AM** and concluded at **9:30 AM**. For the remainder of the day, teams worked to resolve issues that surfaced as expected during a failover process. Two notable issues included intermittent application errors in CU*BASE and network latency between web servers for the online/mobile banking applications and the HA host in Yankton.

These issues peaked as credit unions opened for the business day on **Monday, October 28**. By the time the root cause of the intermittent CU*BASE errors was identified and a script created to correct the configuration files for each credit union impacted, user activity volumes on the system prevented recovery teams from making the necessary repairs. After multiple attempts, including contacting each credit union individually to ask users to log off while the script was executed, the decision was made to force all CU*BASE users to log off at the same time at **1:15 PM** for a period of 15 minutes to allow the script to complete for all credit unions. At **1:30 PM**, all CU*BASE users were allowed to log back in, and teams worked to bring each third-party EFT vendor back online, as well as online/mobile banking applications, including features such as MAP/MOP, RDC, Indirect Lending, Instant Card Issue, etc. By **2:30 PM**, all systems were functioning normally. The network latency issue was resolved the next morning, at approximately **6:30 AM**, after working with the firewall manufacturer support team.

During the early hours of **Tuesday, October 29**, IBM and the hardware recovery teams announced that the new PROD host had been restored and had passed initial system diagnostics. At 1:00 PM, the host was placed in the replication pool to begin the data synchronization process, in preparation for bringing CU*BASE back online at the production data center in Kentwood, MI. That afternoon, teams gathered to identify roll-back options and required system testing necessary to regain the level of confidence in the restored host. Rather than rush the new PROD host testing and risk impacting month-end processing, the decision was made to schedule the roll-back for Tuesday, November 5th.

Demand for system resources on the CU*BASE environment vary on any given day and typically peak on the first day of the month. The number of transactions processed each week tend to peak on Friday. When the two are combined (when the first of the month falls on a Friday), the level of monitoring is heightened and proactive measures are considered to properly allocate system resources to the most critical processes.

With this in mind, recovery teams planned in advance and positioned support teams for month-end processing on **Friday, November 1**, operating from the HA environment. By 9:30 AM, reports were received indicating system performance issues with online/mobile banking applications and third-party EFT transactions. During the troubleshooting process, teams discovered that both the replication and CU*BASE production traffic was being transmitted over the same network interface on the HA host (a configuration setting that occurred during the emergency failover performed on October 27th). Once replication traffic was routed over a separate interface, network latency began to fade and system performance returned to normal.

During the period from **Friday, November 1** through **Sunday, November 3**, multiple tests were performed on the repaired PROD host including repeated power cycles to simulate the conditions that spawned the hardware failure resulting in the recovery effort. Confidence in the reliability in the host was required prior to running live production of CU*BASE core processing.

On **Tuesday, November 5**, from 2:00 to 3:30 AM, the rollback was initiated bringing CU*BASE production back to the primary data center. The old PROD and HA hosts were sanitized and prepared for delivery back to the manufacturer, ending the host upgrade project.

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

EVENT TIMELINE

All times reported are Eastern Time.

Week of September 2nd

- New server hardware for both PROD and HA are received, giving teams two weeks to install and prepare the hosts and three weeks to replicate data and perform initial offline testing.

Week of September 16th

- The new PROD host is introduced into the data replication pool and access provided to developers to confirm application configuration settings.
- Announcements and alerts are issued to credit unions to inform and prepare them for the scheduled rollover maintenance window.

Week of September 23rd

- Offline application testing and system checks are performed as the scheduled rollover date approaches.
- Recovery teams meet to go over the playbook for conducting the rollover with a hard deadline set for failback if required.

October 6, 2019 (Rollover to new PROD)

- At **1:00 AM**, recovery teams stop the services for CU*BASE core-processing and begin the rollover process.
- By **2:30 AM**, the system name-change and address renumbering process is complete. A system reboot (IPL) is required.
- At **3:00 AM**, the data replication software is reinitialized for the new system name.
- At **4:45 AM**, the initialization process is complete. Recovery teams experience intermittent network connectivity issues relative to the address renumbering process. With a complex network architecture consisting of multiple devices and security layers, this was anticipated, and the necessary troubleshooting and resolution steps are taken to correct the issues.
- At **5:25 AM**, the Operations Team begins to start the services and subsystems for the CU*BASE application. By 5:30 AM, transactions are being received by third-party EFT vendors and online/mobile banking environments are online.
- At **6:00 AM**, the Operations Team begins automated daily processing tasks for EOD/BOD. A significant performance increase is noticed already on the new hardware. Several automated jobs are completing in less than half the time required on the previous hardware. Jobs that require read/writes to tape or file transmissions do not experience the same performance increases.
- At **6:20 AM**, an error message when logging on to the GOLD application was observed by a small number of test participants. A global CUBASE Sign On profile was recompiled on the new PROD host to resolve the issue.
- At **7:00 AM**, applications that access the GoAnywhere server for secure file transmissions were experiencing intermittent network connectivity errors. The CU*BASE host utilizes multiple IP addresses for specific functions. When the addresses were renumbered, the VLAN for the NIC bound to the

appropriate IP address for connectivity to the GoAnywhere server was different than that of the old PROD host. Once this was discovered, the correct VLAN configuration was applied.

- By **7:45 AM**, EOD/BOD processing had completed for all time zones.
- At **8:00 AM**, errors were observed in the log files for the MemberReach application queries. A temporary workaround was applied while developers researched the root cause. It was determined that a database tool (SMP – Symmetric Multi-Processing) was partially responsible for causing the jobs to generate the errors. SMP was not installed on the previous host.
- At **10:00 AM**, it was reported that the CTE (Custom Training Edition) objects were out of sync with the previous host. These objects were resynced to correct the issue.

October 9, 2019

- A system error is observed on the new PROD host indicating a hardware failure (one of a redundant pair of voltage regulators). Replacement part and service is scheduled for later that evening. A maintenance window is scheduled for **1:30 AM to 3:00 AM** to power down the host, replace the faulty part, and bring the system back online.

October 10, 2019

- At **1:30 AM**, the faulty voltage regulator is replaced, and systems are restored by **3:00 AM**.

October 15, 2019

- The old PROD host is removed while the new HA host is added to the data replication pool at the HA data center in Yankton, SD.
- Old PROD host is sanitized and prepared for return to the manufacturer.

October 20, 2019

- CU*BASE 19.10 software release update is performed for CU*Answers credit unions on PROD and HA.

October 23, 2019

- Access to the new HA host is made available for application development and QC teams to begin offline testing.
- A maintenance window is scheduled for **Sunday, October 27**, from **5:00 to 6:00 AM** to power down new PROD and move it to the permanent location in the data center (previously occupied by old PROD).

Sunday, October 27 (Emergency Failover to HA)

- At **5:00 AM**, the new PROD is powered down and moved to the permanent location.
- At **5:30 AM**, a system error message is received when attempting to power on the new PROD host. The support vendor (IBM) is contacted to troubleshoot the error condition.
- Recovery teams are put on alert.
- At approximately **7:30 AM**, IBM support determined that a CPU on the host has failed. A new CPU module (8 CPUs) and service is dispatched to replace the CPU. Estimated time of arrival is 3-4 hours.
- Recovery teams are engaged to invoke an emergency failover.
- Initial preparation is performed on the HA host at the Yankton, SD, data center to prepare for the failover.
- At **8:30 AM**, the failover process begins.

- Credit unions with open branches Sunday are alerted (open at 10:00 AM).
- An alert is posted for all CU*Answers online credit unions.
- At **9:30 AM**, the failover process is completed, and CU*BASE is brought online at the Yankton data center. All services are online except for WorldPay ISO connectivity. Teams engage the vendor to troubleshoot the issue.
- At **9:45 AM**, recovery teams observe slow performance conditions for online/mobile banking after the failover.
- At **10:55 AM**, recovery teams attempt to route WorldPay traffic over the backup circuit to the vendor. This does not resolve the issue.
- At **11:00 AM**, one credit union reports having trouble printing receipts. A script is created and executed to repair a file used for the print queue. (The scope of this issue is determined to include more than receipt printing as credit unions opened for business Monday morning.)
- At **11:10 AM**, connectivity is restored to WorldPay after a firewall configuration change is made at the Yankton data center.
- At **11:15 AM**, the replacement CPU module arrives at the Kentwood data center. Upon installation, a new system error message is received, requiring more diagnostics.
- At **2:00 PM**, IBM support determines that an additional CPU has failed (on a different module) and dispatches the appropriate parts.
- At **6:15 PM**, second replacement CPU module arrives from IBM and is installed. Upon boot up, a new system error message is received. IBM escalates support.
- Teams continue to diagnose performance issues with online/mobile banking. Indicators do not point to the HA host, but to a device somewhere in the path from the production web servers to the host. Network packet captures are obtained at each hop to determine the source of the latency.
- At **9:00 PM**, recovery teams attempt to failover to the backup ISP to determine if the latency issue follows. This does not correct the latency issue.
- At **10:30 PM**, the decision is made to perform EOD/BOD processing on the HA host, acknowledging that a roll-back before Monday morning is unlikely.
- At **11:10 PM**, recovery teams attempt to failover to the backup firewall and monitor the latency issue. The firewall failover does not resolve the issue. A support ticket is opened with the firewall support vendor.

Monday, October 28

- At **6:00 AM**, an issue is discovered involving logical files on the HA server with configurations that were damaged or incomplete during the failover (not synchronized with the production server). This is resulting in intermittent issues with users in CU*BASE for multiple credit unions.
- Online/mobile banking applications are taken offline and third-party EFT vendor networks are placed in stand-in mode. A script is created to inspect each logical file and correct those that are not configured properly. With users at credit unions logging in to CU*BASE, file locks are generated causing the script to abort.
- List of services taken offline include:
 - Online/mobile banking
 - MAP/MOP online membership applications
 - Instant Card Issue
 - RDC
 - NADA
 - Indirect Lending Services

- Text banking (for all cuasterisk.com credit unions and self-processors, since this is a shared service across the entire network)
- Third-party EFT vendor networks (stand-in at CU*Answers)
- At **7:00 AM**, working with the firewall manufacturer support, teams identify the source of the network latency and correct the issue.
- At **7:30 AM**, the script to correct the logical files is modified to run on a single credit union at a time. As more users log into CU*BASE, this process becomes more difficult as additional file locks are generated.
- At **9:30 AM**, teams begin contacting each credit union to have them log off all users so that the script can be executed. This proves to be very time-intensive with a low success rate.
- At **12:30 PM**, after considering available options, the decision is made to boot all CU*BASE users off the system for an estimated 10-15 minutes so that the script can be executed to resolve the logical file issue. An announcement and alert are sent out to inform all credit unions that this will occur at 1:15 PM.
- At **1:10 PM**, a system message is sent to all CU*BASE users via the CU*BASE login screen informing them that all logins will be booted off at **1:15 PM**.
- At **1:15 PM**, all CU*BASE/GOLD users are booted off and the script executed successfully.
- At **1:30 PM**, CU*BASE/GOLD is back online.
- From **1:30 to 2:00 PM**, all third-party EFT vendor networks (except for WorldPay) are brought back online, one at a time, allowing transactions in the queue to process before activating the next vendor network.
- At **2:00 PM**, online/mobile banking is brought back online for all credit unions.
- At **2:10 PM**, MAP/MOP, indirect lending services, NADA, and instant card issue are all back online.
- At **2:15 PM**, RDC is brought back online to allow documents in the queue to process.
- At **2:50 PM**, WorldPay EFT network is brought back online. At this point, all services are back online with no further issues reported.
- Later that evening, IBM dispatches and installs a new VPD (Virtual Product Data) module that contains data vital to system operation including licenses, activation codes, component serial and identification numbers, etc.
 - Root cause analysis by the manufacturer (IBM) of the parts replaced during the recovery effort determined that the VPD module was not in fact defective, but that one of the CPU modules installed contained an empty configuration code (p/n) preventing it from communicating with the VPD. This contributed in the delay by support teams in repairing the host and bringing it back to an operational state.

Tuesday, October 29

- At **4:00 AM**, IBM support has successfully recovered the new PROD host and begin running additional diagnostics to validate system reliability.
- At **6:00 AM**, the new PROD host had passed all diagnostics. Recovery teams performed the final steps to complete the emergency failover initiated on Sunday morning, establishing the proper roles for both SOURCE/HA and TARGET/PROD.
- At **8:00 AM**, recovery teams meet to discuss roll-back options now that new PROD is back online.
- At **1:00 PM**, new PROD is ready to be placed back in the replication pool and begins to receive data from HA. Initial estimation for replication completion is 51 hours.

Wednesday, October 30

- At **8:00 AM**, options for the roll-back are evaluated and the decision is made to schedule the event for Tuesday, November 5, from 2:00 to 4:00 AM.
- This schedule will allow teams to complete the necessary testing and comprehensive system audits before bringing CU*BASE production back on new PROD. Trying to rush the roll-back with minimal testing increases the risk of data integrity issues, potentially impacting the month-end processing.
- At **6:00 PM**, it is reported that the data replication process has completed and that systems are synchronized. Testing is scheduled for the following days, including hardware reboots to simulate the conditions that resulted in the failover. All tests were completed successfully over the weekend.
 - The estimated completion time for replication to “catch-up” is based on the current conditions (i.e. available bandwidth, system resources, number of active transactions, etc.). This value fluctuates as conditions change. In this case, the initial estimate of 51 hours was completed in 29 hours, taking advantage of low activity during the overnight period.

Thursday, October 31

- Recovery and Support Teams meet to discuss plans for monitoring and responding to any issues that may occur during peak transaction volumes and month-end processing on Friday, November 1st.

Friday, November 1

- At **9:00 AM**, reports are received that users are experiencing slow performance for online/mobile banking applications. Performance indicators on the HA host are normal. Latency is detected on the HA network in Yankton, SD. Web servers in Grand Rapids, MI, are waiting on a response from the HA host. Requests are queueing up, causing timeouts in web-based applications.
- At **9:30 AM**, transaction queues for third-party EFT vendors is exceeding the normal volumes, causing a backlog.
- At **10:20 AM**, data replication between hosts is temporarily suspended to isolate the latency issue. Network packet captures detect that replication and CU*BASE production traffic is traveling over the same network interface card on the host system.
- At **11:00 AM**, the route configuration for replication traffic is modified to redirect data over a separate interface. This does not immediately resolve the latency issues; however, performance improvements are observed over time as replication jobs begin using the alternate path. Teams monitor conditions as the queues for both web application requests and ISO transactions continue to shrink back to normal volumes.
- An alert is issued to report the cause of the network congestion as an unintentional configuration that occurred during the emergency failover performed on October 27.

Tuesday, November 5 (Rollback to PROD)

- At **2:00 AM**, the HA rollback is performed to bring CU*BASE online back on the new PROD host at the Kentwood data center. Core processing is back online at 3:00 AM.

Friday, November 8

- The old PROD host is sanitized and prepared for return to the manufacturer as per the lease agreement.
- The hardware upgrade project for both PROD and HA has completed.

CHALLENGES AND CONTINUING EFFORTS

Performing a “fork-lift” upgrade of hardware for the CU*BASE core-processing environment in a relatively short maintenance window is both a recovery and high-availability effort requiring careful planning and orchestration with strict attention to details. Each step of the process introduces a level of risk requiring a response strategy to execute when things do not go as planned. This report details the effort of recovery teams at CU*Answers during significant hardware failures in a live production environment.

Challenges observed during these events are listed below, divided between the scheduled rollover to new PROD and the emergency failover to HA that followed:

Scheduled Rollover to new PROD

- After performing the system renaming and network address renumbering step in the rollover process, the new PROD host (now with the IP addresses of the former PROD host) experienced network connectivity issues.
 - While teams initially suspected a problem with MAC addresses and cached ARP entries on the associated network appliances, it was discovered that the VLAN configuration for the temporary switch ports (used for staging the new PROD server) did not match those of the production server.
 - Multiple network interfaces are utilized for the CU*BASE application host, with specific access requirements for each within the TCP/IP stack. Modifying the VLAN configuration on the switch resolved the network connectivity problems.
 - Additional details regarding the proper VLAN configurations will be included in the documentation for future hardware upgrade projects.

- During initial testing, some team members received errors when attempting to log in to GOLD.
 - It was learned that some logins were using a default GOLD screen configuration (27x132 instead of the production 24x80). This was resolved by recompiling the CUBASE profile configuration files.
 - This is an example of an issue that can surface during the initial rollover on new hardware. By utilizing an all-hands-on-deck approach, the issue was detected and resolved promptly.

- During daily processing after the completion of the rollover, errors were detected in the system logs for automated tasks that included the use of database queries. At first glance, it appeared that the query description files were not synchronized with the former host (level check). Updating each query and running them manually seemed to have resolved the issue. This was performed on each query file as a temporary fix.
 - The next day, it was discovered that a database tool (SMP – Symmetric Multi-Processing) was installed on the new server that did not exist on the previous server. While SMP can be used to optimize certain repetitive database tasks, it can also present problems for sequence-sensitive tasks (a condition experienced after the rollover to the new PROD host).
 - As part of the resolution for these errors, the SMP feature was disabled at the OS level and enabled at the application layer, where control is very specific and targeted for those functions that can benefit from the optimization.

- During testing after the completion of the rollover, it was discovered that the CTE (Custom Training Edition) library files were not synchronized with the prior server.
 - These libraries were resynced with old PROD to resolve the issue.
 - It is not uncommon during rollover events for developers to recompile an executable file or resync a data file to ensure that all attributes are carried over to the new host. This is where the comprehensive post-roll testing process is put to the test.

- With the move to 24x7 processing across six time zones, the duration of the rollover to new PROD was itself a challenge to schedule around. Where a typical HA rollover can be performed in approximately one hour, the process to rollover to new hardware requires a longer maintenance window. The initial schedule was based from the previous event (performed in 2016) which required 4 ½ hours. The 2019 new PROD rollover, even with the increase in performance, required the same amount of time.
 - We can anticipate that maintenance windows for performing system upgrades will continue to be squeezed and reduced in the future. In preparation for the next new PROD rollover (est. 2022), teams will evaluate measures to comply with the expectations while working with vendors to optimize the process without risking the integrity of the data.

Emergency Failover to HA

The original project timeline discussed and approved included the step of powering down the new PROD host and moving it to the permanent spot in the data center (vacated by the outgoing PROD host). In hindsight, especially following the initial hardware failure (voltage regulator) on the new PROD host, additional discussions around the sequence of events and the option of performing an HA rollover as a precautionary step before moving the new PROD host, might have minimized (or postponed) the risk of what actually occurred.

- In reviewing the decision process, this is a step (moving the host) that has been executed multiple times without incident. In the 2 ½ weeks that had passed since the voltage regulator failure, there were no indications of any pending hardware issues on the new PROD host. The IBM hardware system components include firmware that is highly sensitive to predictive failures. There were no signs that a second hardware failure was around the corner.
- The maintenance windows for performing the multiple steps required to upgrade both the PROD and HA servers parallel to production are extremely tight in the 24x7 processing environment we operate in across six time zones. The dates selected competed with other significant events such as end-of-month processing and the CU*BASE 19.10 software release.
- Future server upgrade projects will require an even more proactive approach to mitigating the risk of potential downtime during each step of the process, while at the same time, the expectations are that maintenance windows will continue to shrink, requiring creative solutions and strategies.

One of the factors that contributed to the prolonged restoration (46 hours from initial hardware failure) of the IBM host involved the dispatching of multiple parts from remote locations, some as far as four hours away from the production data center. The new servers installed in this project included an upgrade in hardware platform, one in which the manufacture has yet to establish a parts repository near the Grand Rapids area. This was part of a commitment by the vendor to improve response times based on prior incidents while operating on former platform.

- We will continue to put some pressure on the support vendor to improve local inventory of replacement parts for this new platform.

After the emergency HA failover on Sunday, October 27, network latency was experienced on the HA environment, resulting in slow performance for online/mobile banking applications throughout the day. Recovery teams worked through the process of capturing network packets to trace the hops across the multiple devices, beginning with the web servers at the production data center through to the HA host located 750 miles away in Yankton, SD. Later that evening, the symptoms pointed to the redundant firewall appliances at the HA data center. The support vendor was engaged, and resolution applied early Monday morning.

- While network monitoring tools are in place at both the production and HA data centers to trend performance and uptime metrics as well as detect and alert when outages occur, there was nothing in place to detect the latency that was experienced in advance (by comparing response times to that of the production environment).
 - Recovery teams will research options for implementing additional ongoing monitoring to detect any performance issues. One factor that adds a degree of difficulty to this goal is the significant difference in amount of traffic between the production and HA network. Operating on the HA network for 2-3 weeks each year during rollovers may not provide an accurate baseline to determine when deviations occur.

On Monday morning, October 28, following the emergency HA rollover, multiple attempts were made to correct CU*BASE configuration files that were resulting in errors in CU*BASE. A script was developed to inspect and correct the errors in the configuration files. The script, however, required isolated access to each file. As users logged into CU*BASE and online/mobile banking, or transactions posted by ISO/EFT vendors, file locks were generated, causing the script to abort. Multiple attempts were made, including sending email alerts and calling each credit union individually to instruct users to log off. This process resulted in delays, lengthening the amount of time services were offline for members. The decision was made to force all users off the system at the same time (1:15 PM), execute the script to correct the errors, and bring CU*BASE back online (1:30 PM).

- Fortunately, these types of scenarios are rare and unprecedented. As a result, there was very little history or experience to look back on in developing a strategy to correct the issue. Often, we tend to prioritize our recovery attempts based on the amount of impact to the end user (least to greatest). In this case, given that the script to correct the errors only required a few minutes to execute, the decision to yank the bandage off (boot users out of CU*BASE during the morning hours) and correct the issue would have minimized the cascading effect that was experienced.
- In addition, more effective communication methods to users within CU*BASE are being developed so that future incidents requiring action by the user can be implanted in a timely manner with less confusion.

During both scenarios detailed above where system/network latency impacted performance of CU*BASE and online/mobile banking applications, recovery efforts involved bringing entire applications offline or booting all CU*BASE users offline at the same time. In some cases, having the ability to turn off only select features of an application or a specific group of users or vendors might have provided more options for recovery teams and limited the impact.

- These objectives may be too complex and expensive to apply to existing applications but will be considered for future revisions.

One of the most critical components in any recovery effort is that of communications. During the emergency HA failover on Sunday, October 27, an alert was posted, followed by an email announcement on Monday morning. Most credit union staff arrived in their office with little knowledge of what had occurred until they read the email message and alert.

- The existing channels for communicating with credit unions was originally developed with a business-day mindset, Monday through Friday, 8-5. A new or enhanced process for communicating 24x7x365 is required to alert and keep credit unions informed during emergency incidents.
 - The recovery team has defined this as a high priority objective for 2020.

CLOSING REMARKS

Earlier this year, in a [report](#) detailing the emergency HA rollover performed as a result of a February extreme winter storm that included widespread utility power outages and a hardware failure on the generator at the production data center, a statement was made in reference to the so-called rarity of these unscheduled, emergency recovery efforts. Over the years of developing and testing the business continuity program, we have learned that these scenarios are not rare at all. In fact, as we near a close of 2019 and of the decade, the trend has been toward an increase in the number and scope of disruptive events for organizations across the country.

As applications and networks become more complex and integrated, while at the same time, tolerance for downtime becomes practically non-existent, businesses are rethinking their strategies for achieving the level of resilience required to compete and succeed. As a core data process and CUSO for a network approaching 300 credit unions, we will continue to be aggressive in investing in and testing our capabilities through live rollover events.

While the details of this report highlight the events surrounding an October emergency HA failover, the identified Challenges and Continuing Efforts provide a precursor to the strategies that are required to take our network and operations into the next decade, preparing for the next disruptive event. In conclusion, we as a cooperative, learn and improve more during one unscheduled, emergency recovery effort than in a dozen planned rollover events.