CU*ANSWERS
A Credit Union Service Organization

iSeries Failover

5/28/2010
Revised 6/10/2010

Table of Contents

Message from the CU*Answers CEO

The experts will tell you "of course your disaster recovery procedures and investments will pay dividends," but in the back of every business person's mind is always the question whether the insurance policies you buy are giving you the return on investment you are counting on. Well, last month we had an event that confirmed that years of investing, training, and practicing do pay dividends when trouble calls the plans forward and into play.

On May 28, 2010, the Friday morning prior to Memorial Day weekend, we had to declare a disaster condition that required an emergency failover to our backup computer resources. In every event like this a CEO worries that you do not really know until you know, and for our network we now have some experience with this procedure that makes me feel confident in the return on investment for our High Availability Strategies.

The accompanying report will outline the event and is designed to be entered into your Board minutes or official record. Hopefully you have been documenting the regular series of HA rollover exercises we do every year with your Board of Directors. (Check out http://www.cuanswers.com/hagapanalysis.php for more details.) Those events are the proof of concept practice events that keep us in a ready state for situations like this. Add to that our annual disaster recovery tests that we complete each year with IBM in Chicago and you can see the layered strategy we employ as a CUSO to protect your member's data and your operations. Look for an upcoming report on the 2010 Disaster Recovery Test (performed May 18-20 at our hot site in Chicago) in the next few weeks.

As disasters go, this was simply a local event in our computer room, not the kind of stuff disaster movies are made of. Still, these localized events are an excellent gauge of the effectiveness of our High Availability Strategy. This incident will serve as an invaluable learning tool with which we'll further develop our response and protection for your members and operations. Although we successfully averted what may have been a real crisis, we won't lose a step in keeping our testing and preparedness procedures up to date. That said, it is nice to have a real example of the disaster recovery strategy when you report on the next HA Rollover Test to your Board, even if we hope we never have to experience another again.

Randy Karnes, CEO

1.0 Incident

On May 28, 2010 at approximately 6:41 a.m. ET the CU*Answers primary computer processing system (PROD) suffered a system shutdown. This PROD box was new. The box underwent 10 days of quality testing during the break-in period and had been in service for approximately 38 hours before the system failure. The Incident Response Team (IRT) attempted recovery of PROD and contacted executive management. At approximately 6:57 a.m. ET the IRT initiated the disaster recovery failover process from PROD to the High Availability system (HA). At approximately 8:00 a.m. ET the failover was completed and all core systems were restored. Credit Union staff regained access to CU*BASE at approximately 8:05 a.m. ET and to third party services (**It's Me 247**, CU*Talk, ATM, Debit and Credit Card services) around 8:15 a.m. ET.

No member data was lost or affected during the outage. Three files that were corrupted during the outage related to non-critical application event messages.

2.0 Affected Services

All CU*BASE core processing services and related services (**It's Me 247**, CU*Talk, ATM, Debit and Credit Card services) were unavailable during the outage. The outage lasted approximately one hour and five minutes.

3.0 Recovery

3.1 Hardware Failure

At approximately 6:30 a.m., OPS received an unusually severe drive failure error message from the PROD system. OPS contacted the iSeries (System i) Administrators, who engaged in troubleshooting efforts. About 10 minutes into the process, the entire PROD system shut down due to a problem with the hot swap process in the RAID controller. The iSeries Administrators notified the CIO, the CIO informed the CEO, and the CEO informed the Chair of the Board.

3.2 Alerts

Alert progression occurred as follows (times are approximate; full Alert detail in section 6.0):

6:30 a.m. ET	First critical error message to Operations (OPS); iSeries Administrators contacted
6:41 a.m. ET	IRT and CIO
6:43 a.m. ET	CEO
6:57 a.m. ET	Chairperson of the Board
7:15 a.m. ET	First CU*BASE Alert to Clients
9:09 a.m. ET	All Clear CU*BASE Alert to Clients
5:04 p.m. ET	Follow-up Alert to Clients summarizing the day's events

3.3 Failover Initiation

The iSeries Administrators and Executive Management agreed this was a disaster situation and engaged in failover from PROD to the HA system. The failover was initiated at 6:57 a.m. ET and completed at 8:00 a.m. ET.

3.4 Service Restoration
The failover process was a success and CU*BASE services were restored by 8:05 a.m. ET to Credit Union staff. Third Party services were restored and available no later than 8:15 a.m. ET.

3.5 Troubleshooting
At approximately 11:00 a.m. ET representatives of IBM arrived to assist with troubleshooting the problem. iSeries Administrators in conjunction with IBM and the CIO determined the hot swap error was the cause of the problem.

3.6 Replication
At approximately 5:00 p.m. ET replication services were restarted. Plans were engaged to rollback from HA to PROD within the next few days. Successful rollover to PROD was completed Wednesday, June 9, 2010 at approximately 10:00 p.m. ET.

4.0 Post-Mortem

4.1 Proximate Cause – Hot Swap RAID Error
The cause of the event was a hot swap error in the RAID controller. In a RAID array of drives, when a drive fails the next drive is supposed to immediately engage. The error prevented the swap from taking place and shut PROD down. In this case, the entire hot swap process was unnecessary because PROD is mirrored, but unfortunately the OS on the iSeries did not allow the Administrators to override the hot swap. Administrators are working with IBM to determine if there is a way to override the hot swap process.

4.2 Disaster Recovery Procedures
CU*Answers regularly performs HA rollovers which are planned events to test recovery readiness. A failover is a disaster recovery event in which we perform such a rollover in response to a system failure. Disaster Recovery procedures and tests prepared the IRT for this anomaly and were critical in reducing recovery time.

4.3 General Employee Response
Employees acted calmly and provided updates to clients as warranted.

5.0 Recommendations

5.1 PROD Replacement Timing
The PROD system is replaced every three years. CU*Answers will no longer attempt to replace PROD during the last week of a month. New guidelines for PROD replacement will include the following:

- PROD will be replaced at least 5 days ***after*** end of month processing. PROD will ***not*** be replaced around a holiday.
- No critical event will be scheduled at the same time (new releases, conversions, etc.).
- Critical employees will endeavor to be available the 3 days prior and post replacement.

6.0 Alert Detail

6.1 CU*BASE System Issues

Published May 28, 2010 at 7:15 a.m. ET by Operations Manager
http://alerts.cubase.org/cubase-system-issues/

The online CU*BASE system and all related services are currently unavailable due to hardware issues. We are diligently working on restoring processes as quickly as possible. Please check back for frequent updates as we have additional information.

6.2 System Update

Published May 28, 2010 at 9:09 a.m. ET by Writing Team
http://alerts.cubase.org/system-update/

CU*BASE core processing is back online, and other related services are coming back online now. Here are the details of this morning's occurrence.

At approximately 6:41 a.m. ET on Friday, May 28th, CU*Answers' primary computer system (iSeries) experienced a server hardware failure. The Emergency Response team notified the CIO immediately, the CEO was notified at 6:53 a.m., and the Chair of the Board was notified at 6:57 a.m. It was then decided that we would roll from the Production computer to the High Availability system via a failover process.

As part of our normal periodic procedures, we perform HA rollovers which are planned events to test our recovery readiness. A failover is a disaster recovery event in which we perform such a rollover in response to a system failure. Disaster is a strong word here, but it is the proper one: the system went down and the team had to react.

As part of our high availability and business continuity planning, we began replicating transactions and data to backup systems. At approximately 8:00 a.m. ET, we restored all core systems to the backup system. Credit union staff was able to access the system starting at approximately 8:05 a.m. and third party services (**It's Me 247**, CU*Talk, ATM, Debit and Credit Card services) began coming back online starting at approximately 8:15am ET.

Our emergency response team is engaged in reviewing any potential data loss, although there are no indications of member data loss at this time. We have three files with potentially corrupt data that relate to system messages; they are not associated with member data. All credit unions should be on high alert for any abnormalities or hints of issues caused by this morning's failover. Please contact a CSR immediately should you encounter any issues.

System teams are currently engaged with IBM support and they are assisting us with diagnosing the cause of this morning's event. Later today we will communicate how long we expect to be running on the HA system and when we anticipate rolling back. In the meantime, we are also evaluating whether to resume replication to our old Production system (remember we just replaced this box on Wednesday), or whether we'll immediately start to replicate to the new system as it comes back on line. More information will be posted later today.

We appreciate your patience and will send more information as quickly as we can. Your teams should be documenting this event as a disaster recovery incident and be noting the event with your Board in your Board minutes as part of your disaster planning due diligence.

6.3 System Update (cont'd)

Published May 28, 2010 at 5:04 p.m. ET by Writing Team
http://alerts.cubase.org/system-update-contd/

As a follow-up to the system failure this morning, IBM support diagnosed the problem as being due to a specific hardware issue. The Systems team was able to replace the failed hardware. Our team has been monitoring the box all afternoon and is feeling very confident that it is ready to be put back into the mix for normal replication, which will resume this evening.

We are very pleased with the HA system, which has been in place for roughly a month and has been handling the processing very well. We will continue processing on the HA system and we will be ready for a normal month end on Monday night. An ongoing evaluation of the new box is being performed so that we can determine the date on which we will roll back.

Final results of the diagnostics performed will be published next week for your Board packets.