

Tips For A Successful Recovery: Coming Back From The Brink

By Sixto Ortiz Jr.

8 October 2010

Processor

Whether they're due to natural causes, sabotage, or equipment failures, data center disasters can have serious consequences for a business. Besides the immediate expenses incurred, there is a potential for loss of confidential data and perhaps irreparable damage to a company's reputation.

In the midst of a crisis, it's difficult for admins to focus on anything other than getting the data center back online and recovering any critical data. Implementing a recovery from a data center disaster is vital; whether that recovery is successful hinges on the readiness of a business to deal with an unthinkable disaster.

Assess What's In The Data Center

In the aftermath of a disaster, it is important for admins to know what data resides in their data centers so it can be quickly determined what needs to be recovered. Without that knowledge, recovery efforts will flounder and precious time will be wasted.

Dr. Mickey S. Zandi, managing principal, consulting services, SunGard Availability Services (www.availability.sungard.com), says SMEs that need to recover from a disaster must understand the level of system interaction and interdependency. Admins must know what is in that data center and the applications, systems, components, and elements operating in the data center. Then, personnel must understand how those components are linked to each other.

Involving the business can improve the assessment process by helping administrators properly prioritize recovery efforts. Karen L. Cole, CBCP, SPCI, and CEO at Assura Consulting, says admins should get input from the business on expectations for recovery times of systems and data. IT sometimes assumes that all data in the data center must be immediately available, which is not the case. Having the business establish recovery requirements reduces the cost of IT recovery strategies and the stress on IT teams involved in setting recovery priorities by determining what needs to be immediately recovered and what can wait.

Have Backup, Will Travel

For recovery efforts to flow smoothly, admins should be prepared to deal with data center disasters well before they occur. But, what specific items should admins focus on when preparing for recovery?

A required area of preparation is planning and deploying a secondary data center that runs the same applications concurrently and delivers the same services under the original SLA, says Nir Ilani, director of product marketing for application delivery at Radware (www.radware.com). Rolling out such a project, Ilani says, involves covering all data center operational aspects, including storage, networking, servers, electricity, and cooling.

Administrators must also take into consideration capacity demands and the projected load at the time of a disaster. Once in place, a global server load balancing solution should be implemented in both data centers so all traffic and user transactions are transparently redirected to the secondary data center, he adds.

Cloud computing can also have real-world utility in terms of disaster recovery. The easiest way to utilize a cloud computing approach is to simply back up data to the cloud, says Andi Mann, vice president of virtualization product marketing at CA Technologies (www.ca.com). But, administrators can go beyond simple backups to the cloud by choosing to back up applications, systems, and even entire services to the cloud. This, he adds, will require the ability to run the entire system and application stack in an offsite location.

Easily Forgotten Considerations

Many data centers have generators, but without fuel, those generators are useless. According to Cole, many data center managers do not have an agreement with a fuel provider that can provide emergency service. This leads to slow service and price gouging as administrators scramble for fuel supplies, so service levels for processing should be negotiated ahead of time. This will ensure that the required power is available during disaster recovery operations.

Another important consideration, Ilani says, is the potential for massive traffic surges for Internet-based services during a disaster, which could amount to 50 times more than routine traffic. Thus, admins must ensure that network infrastructures are able to process such an increase in capacity and must scale capacity on demand.

The ability to perform actions remotely is also critical to implementing a successful recovery. Jeffrey Godlewski, technology specialist at CDW (www.cdw.com), says organizations should determine which functions and specific positions are compatible with remote work, even if they are not typically done remotely. Administrators should also consider how well telephone and messaging systems will support personnel redeployment and should analyze telecommunications bandwidth and whether it is sufficient to support the redeployment of all remote-capable positions, he adds.

Other considerations along the area of remote work capabilities include the remote access technology in place and its ability to scale, the deployment of sufficient remote or mobile computing devices, and an assessment of available remote access security tools and their scaling capabilities to support redeployment of personnel, he adds.

Backup & Restore Considerations

Having data backed up provides organizations with peace of mind, but a number of items must be carefully considered prior to restoring data. Jeff Pederson, manager of data recovery operations for Kroll Ontrack (www.krollontrack.com), says backups should be restored to a different volume to ensure that all important files are correct and secure on the backup before potentially overwriting data on the active volume. Pederson also recommends that if there is a RAID problem, personnel should test the backup by restoring it to a different location or image on each drive from the RAID before attempting a rebuild. Sometimes, he adds, a RAID rebuild does not work correctly and can make a problem even worse.

Also, Pederson says, personnel should not create any new files on the disk needing recovery or continue to run virtual machines until the important data is recovered. This is because new files can overwrite the files needing recovery if restoring the backup fails. Virtual machines using snapshots and thin provisioned virtual disks still in use after the data loss can overwrite files that need recovery. Personnel should also avoid running file system repair tools on a virtual disk unless a good backup has been validated by restoring it to a different volume. These repair tools assume there is a good backup of the data and can overwrite file pointers to make a system consistent.

Finally, in a flood situation, Pedersen says, hard disk drives should be kept as moist as possible to prevent disk corrosion and allow recovery experts to clean and dry the platters correctly with minimal damage to the platter surfaces. If the drive dries out, he adds, there is potential for damage or destruction to the platters, making it more difficult to retrieve data.