Survival skills for a digital disaster
Why preparation is key
By Dave Bell, President, Cyber Solutions, LLC

How many times have you heard this story? Company ABC suffers a technological disaster, loses all data, attempts to recreate data from scratch, then goes out of business. It's a bit like a Hollywood horror movie — we know the premise, but we've become immune to the disturbing results. Some of us might even utter the words, "That could never happen to me."

This article assumes that it will indeed happen to you. And, when it does, will you be ready?

There are five keys to a successful backup and disaster recovery (BDR) solution. These keys differ slightly depending on the size of your business, but the concept is the same.

Key No. 1 — Centralize data. Your company data is much easier to manage when it is stored in a centralized location. Consider a small office that has three workstations. The user at each station stores their work on their local hard drive; Excel files, Word files, QuickBooks files are all scattered across three machines. The more workstations, the more complex this problem becomes.

Instead, designate one machine as the central storage location. Allow other users to access this machine to access files. As the firm grows, you should consider implementing a true server that can act as a central location for data management.

Key No. 2 — Find the right backup software. You have two choices at a high level: file-based backups or image-based backups. As the name implies, file-based backups create a copy of each file that is configured within the backup set. On the other hand, an image-based backup will take a snapshot of the *entire* machine. Think of this snapshot as a bubble containing everything you need to quickly recreate the entire machine from scratch (operating system, programs, data, etc.). In the event of a disaster, this bubble can be easily moved to another machine or up to the cloud. It also provides the exact benefits of a file-based backup where individual files or folders can be restored. For these reasons, image-based backups are the preferred method.

Key No. 3 — Replicate to a remote location. You should have your snapshots located in at least two locations. First, your backup software should store each snapshot on a device in your office (USB drive, network-attached storage, warm-spare server, etc.). This local copy provides for fast recovery of files or folders in the event of a non-catastrophic disaster, like deleted files or Ransomware infection. If you have the resources to implement a warm-spare server, you can quickly recover from a complete hardware failure by loading a snapshot on the warm-spare server and be fully functional within minutes.

Second, your backup software should replicate this snapshot to a remote location either in the public cloud or a private data center. This provides protection against catastrophic events like fires, floods and theft.

Also, both of these tasks should be automated. Gone are the days of a staff member "swapping tapes."

Key No. 4 — Simulate a disaster. Hopefully you are doing some form of the first three keys already. Unfortunately, this is where the vast majority of firms and businesses stop. Disaster simulation takes significant planning and hard work. However, it is imperative to ensuring disaster recovery *and* maintaining business continuity.

To put it simply, you should simulate “pulling the plug” on the device that contains your centralized data. For a smaller business, you may simply need to restore your image-based backup to a different machine and test your applications on the new machine. However, for larger businesses, a full-blown disaster recovery cutover is necessary.

Run a controlled-disaster simulation

As an example, we run a disaster simulation for a client on a regular basis. Their environment consists of the following components:

1. Five virtual servers balanced across two physical servers
2. On-premise email
3. Remote access via Citrix for field staff
4. Dozens of applications for staff
5. Two offsite warm-spare servers for disaster recovery
6. Image-based backup software that creates local snapshots every hour
7. Snapshots are replicated offsite

Since we’re in control of this simulation, we can schedule the cutover during a slower business period and during a time when a smaller set of users will be affected. This allows us to test complete functionality without the risk of significant downtime for the entire staff if something goes wrong.

Here’s the typical schedule for the simulation:

Thursday — cutover day

1. We suspend email delivery to the production servers. This firm uses a third-party offsite spam filter, allowing email to be queued while delivery is suspended. If you are not using this type of solution, be aware that email may not be queued (or only queued for a brief period) before a non-delivery report will be sent to the sender. If you are using hosted email through Microsoft Office 365 or other provider, you can skip this step.
2. We run a final snapshot and then pull the plug on the production servers. (In reality, we power them down gracefully in order to prevent corruption. This provides us a reliable means to fallback if necessary.)
3. We turn on the replicated critical servers within the disaster recovery (DR) site. Once these servers are running, we run some quick diagnostic checks to ensure the servers are functioning properly and communicating with each other.
4. A test is conducted from an external workstation to ensure a connection can be made to the DR site and applications can be launched. In this example, Citrix is used to create a remote connection. However, you can use a VPN or other remote access software to create the connection.
5. If we have any doubts, we revert back to our production servers within the office (it is much easier to do so now before the DR site potentially contains production data). If we move on, our DR site will start to receive email and Citrix logins. Typically, this is done within DNS or within the third-party spam filter.
6. We configure backups once we confirm email is flowing to the DR site. Even while testing within the DR site, we *must* have backups configured in case we suffer a disaster within our DR site. Yes, we’re paranoid.

Friday — user-testing day

1. We monitor the servers at the DR site for performance.

2. We also make note of any issues that users may experience.

Saturday/Sunday – Implement Key No. 5

Key No. 5 — Know how to get back. So you might think it's time to high-five and go home. Unfortunately, once the testing is complete, you must know how to get back to the office servers. In this test, we simply copy the virtual machines from the DR site and transport them back to the office over the weekend. Depending on the amount of data, the time to copy can be significant. Be sure to account for this when scheduling.

At the conclusion of our testing, we schedule a debrief meeting in order to implement any changes to our documentation. This also gives us a chance to address any user concerns that may have come up.

Survive disaster

The simulation above is a very specific example, but it can serve as a basic framework for you to plan your own disaster recovery simulation. Overall, no matter the size of your organization, a disaster recovery simulation is imperative to surviving a digital disaster.

Don't let your firm be the star of the next technology-related horror flick. Take steps today to prepare yourself. When it comes to digital disasters, you don't want to be famous.

David Bell has served in the IT industry for over 20 years. He currently serves as President for Cyber Solutions, LLC. You may reach him at 651.379.5741 or dbell@cybersolutions-web.com