WHITEPAPER

Building higher IT business continuity in the face of disaster

Most disasters strike when you least expect them, Some – like hurricanes – give you more warning. But in either case, it's a real problem if disaster brings your business to a grinding halt. Forces of nature, malicious acts, or even a simple mistake can have a long-lasting negative effect on your business. This idea paper offers some first steps and follow-up steps you can take to upgrade your disaster preparedness.

How will data loss affect your organization?

Losing your data or critical systems can mean:

- Revenue loss from the inability to conduct business
- Lost customer trust or confidence
- Financial penalties for violated SLAs
- Legal or financial penalties for compliance lapses
- Excessive recovery and repair costs for lost systems and data

Organizations that would have absorbed one- or two-day outages now find that hours-long outages are a matter of crippling concern. No longer is it just your company's employees who are affected by a failure to recover quickly. Partners and customers also link to your IT operation through an increasing number of outward-facing applications including email and database applications. Operational continuity is critical.

Hours to recover? Better if you can do it in a half hour or less.

True business continuity is made up of three pillars:

- **High availability.** Today's internal and external users have much higher expectations when it comes to server availability: they want to see applications available all the time, 24 hours a day in many cases. This raises the bar for recovery time objectives (RTOs), where critical virtual and physical machines must be recoverable in a matter of minutes, rather than hours.
- **Onsite** recovery. It's always faster and more convenient to recover locally if you can, but if you lose your server room or data center, local recovery is not an option. That leads us to:
- **Offsite or cloud recovery.** The first level of defense against a data center or server room is an offsite DR location. But when damage is widespread, in the case of tornados or hurricanes, your DR facility may also be area that is pummeled by the same forces that took your data center offline. When geographically widespread disasters occur, the value of a cloud-based disaster recovery becomes even more apparent as you work to make everyone's access to applications as robust as possible. For many companies, cloud is fast becoming a primary defense against data loss, especially if it can be managed from any laptop, netbook or tablet with an internet connection and a Web browser.

First steps to move to a stronger disaster recovery position

What's your best path to a stronger disaster recovery plan? Here are first steps for small and medium-size businesses as well as small enterprises.

- *Take a quick inventory* of all your IT-related business processes. This includes everything from financial applications, logistical functions, email, outward-facing functions and more.
- *Rank them for recovery priority.* Think about which applications are necessary for your company to generate revenue or are critical to business continuity? What data can't your customers do without? What's critical to running your internal accounting and finances? And what is required for compliance? With those variables in mind. Now, create a list in in descending order to establish your DR recovery sequence.
- *Establish a Recovery Time Objective* (RTO) for each function. Ask yourself, "How fast do I need to recover this application?" Email and transaction-based applications that people inside and outside the company depend on at all times will probably be near the top of this list, whereas applications that are less frequently accessed, such as a human resources applications, may be low on the list because they're ancillary to your immediate business continuity requirements.
- *Establish a Recovery Point Objective* (RPO). How much data can you afford to lose for business process? How important is the data that you could lose? Applications related directly to business continuity, where data changes significantly every day, will top this list. Back office processes may be lower on the list.

Create a "Break Glass in Case of Emergency" plan. Define where you want to keep your DR data and systems. If you are located in an area that could be hit is regional weather events like hurricanes, floods, or wild fires, then select a secondary location outside of your region that you can fail over to when disaster strike at your primary location. Your choice could include cloud-based recovery.

Next Steps to increase your DR responsiveness

- *Determine which of your RTOs and RPOs can be supported by your existing backup and recovery scheme.* This will allow you to figure out pretty quickly which of your processes are going to "fall through the cracks." Certain applications, like very heavily used SQL or Exchange applications may need to be backed up even more frequently, and if your current backup scheme can't support anything more frequent than once daily backup, you may wish to consider investing in a newer, more aggressive disaster recovery solution.

- *Consider your DR options.* If your current system is not up for the job, select a Disaster recovery solution that best meets the business and recovery objectives you have developed in the previous steps of this plan (see below in the next section for more information). Once it's installed and in production, make sure your staff is trained how to use it.

- *Assign responsibilities so everyone knows what to do* when a disaster strikes. Assign everybody involved in the DR plan a specific task. Don't expect the relevant personnel to always be at the disaster site or to be in control immediately. Implement necessary duplication and redundancy for people, just like you would do with computers.

- *Test, test, test!* One of the worst feelings an IT administrator can have is discovering a backup is corrupted in a disaster recovery scenario. Why wait to find out when it's too late to do anything about it? Test your backup data for corruption when you back it up. Newly developed software allows you to test for recoverability automatically. Use these available tools.

- *Practice, practice, practice!* The more experience your team has successfully carrying out a simulated disaster recovery, the more comfortable and quick they will be to succeed when the real thing happens.

Reviewing your backup product for DR suitability

Planning for DR is critical, but prioritizing and planning alone won't help you recover quickly if you don't have a disaster recovery solution capable of executing your recovery objectives. If you have had a backup and recovery solution for four or five years, you may be missing out on essential performance capabilities that can mean the difference between a recovery too slow to save your company and an alternative scenario where the disruption is merely inconvenient.

Legacy backup schemes are hard pressed to keep up

Because speed is of the essence in a disaster situation, even with disk-based legacy file-and-folder backup and disaster recovery solutions may be too slow to meet the needs of today's always-on computing infrastructure.

They're only designed to recover half of your computing infrastructure – the files – leaving you to carry out hours long bare metal system-level restores, requiring you to purchase identical spares that can be turned on in the event of a server failure, or forcing you to buy additional system-level recovery software. This can unnecessarily complicate your ability to recover quickly if disaster strikes and create productivity losses that are simply unrecoverable. This is especially true if you depend on tape for recoveries of business-critical machines. The risk is too great to take when minutes count: each server recovery can take hours and the potential for a failed recovery is high, either from mechanical failures or corrupted data that foils a recovery.

Image-level backups a key beginning point to faster recoveries

Image-level backups are a key beginning point for a robust business continuity plan, and there are great solutions to choose from that meet even the most stringent business continuity requirements. Ideally, you'll want a continuity plan that deploys across all of your Windows server resources regardless of the server's location. Check to see if both your virtual and physical machines will be manageable from a single user interface. Also important: control the amount of data being produced particularly in virtual environments. Address this issue by employing data deduplication and file compression. When it comes time to recover, smaller data sets can be restored faster than large ones. Preferably these two data reduction utilities should be fully integrated into your backup and recovery software; they'll be easier to use.

Second, look for the flexibility to recover at any level (file, folder or system levels) and in any direction, including from physical to virtual and virtual to virtual, without having to preconfigure hardware in advance. With such a system in place, you can establish a certainty about your ability to recover, and that's what a successful business continuity plan is all about.

Third, check your backups for data corruption. Any experienced backup administrator already knows this is a great idea, but the fact is, it's very time-consuming to do. In fact it's easier said than done, as the traditional process of testing a backup requires about the same time and effort associated with a disaster recovery. It can take hours even with disk-based backups. That's why nobody really has the time to check every single backup.

Recovering hybrid environments

There are very few organizations that haven't embraced the benefits of virtualization, both to lower costs and increase resilience to system-level failures. But there are also very few organizations that don't also maintain physical servers as well, and some of them are running business-critical applications. Given the reality of a hybrid computing environment, how can you protect everything well enough that you can be certain to recover quickly in a disaster? Let's begin by looking at two approaches that don't work:

- Older legacy solutions featuring file and folder backups scale poorly to hybrid environments. And they don't take advantage of key capabilities built into the fabric of virtual platforms.
- Virtual-only backup and recovery solutions force you to maintain one backup scheme for virtual and another for physical. This can be a real problem when disaster strikes if you and your staff are using two unrelated solutions to recover your physical and virtual machines.

Instead, look for a solution that works seamlessly across virtual and physical environments. That means you'll only have to learn DR once, not twice, increasing the likelihood for a quicker recovery.

Recovering from a distance

Too much is riding on your data to depend exclusively on local backups, which is why offsite storage has been a default solution for decades. However offsite file and folder storage really only applies to data archiving, too slow to support a rapid disaster recovery effort. Those who can afford a disaster recovery facility have a great deal more to work with to ensure a quick recovery, with machines on standby and data replicated and ready to run after a failure. However, if you're not a big company, you're less likely to have your own site, and it's not just the cost of an offsite facility that drives it. It's the cost of outfitting that site and staffing it. And that leads us to the cloud.

The rise of cost-effective private or public cloud environments is transforming everything about disaster recovery, but it works best when you move to an image-based backup model. Choose a DR solution that up a complete image of a system and its data, so everything you need for a quick recovery is there when you need it. Look also for incremental forever, block-level backups, combined with global deduplication. Combined, they'll allow you to transfer the smallest possible amount of data across communications lines and dramatically maximize your disk storage usage.

Introducing AppAssure disaster recovery software

Unified Backup, Replication and Recovery

Dell AppAssure 5 Backup, Replication and Recovery software delivers local and offsite backup, replication, disaster recovery with cutting edge snapshot backups and ultra-fast recovery technologies. AppAssure delivers:

- **Push-button failover** to standby virtual machines
- Easy **bare-metal restore** to similar and dissimilar hardware that stands up servers in minutes, with full access to data
- **Integrated data replication, deduplication and encryption**
- **Certain recoveries.** AppAssure Recovery Assure™ technology automatically tests for corruption so you can be certain the backup will be completely recoverable.
- **Protection for both cloud-based and local data via one interface.** AppAssure gives you the freedom to manage your AppAssure-protected machines from wherever you are with a Web browser and an Internet link.
- **'Government-grade' encryption for data security:** AppAssure keeps your cloud-resident data safe from prying eyes with 256-bit AES encryption and password protection.

To learn more about keeping your data safe and about backup and recovery on virtual machines, visit AppAssure at www.appassure.com