HIGH AVAILABILITY IS NOT A LUXURY

Eliminating Downtime for Small and Mid-Market Organizations

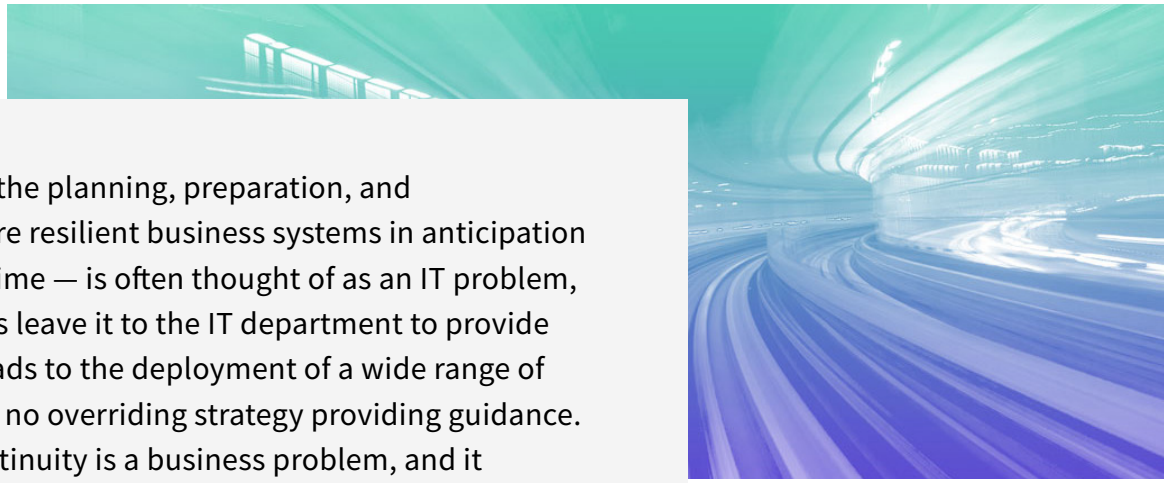

Business continuity — the planning, preparation, and implementation of more resilient business systems in anticipation of unscheduled downtime — is often thought of as an IT problem, and most organizations leave it to the IT department to provide a fix. This invariably leads to the deployment of a wide range of tactical solutions, with no overriding strategy providing guidance. In reality, business continuity is a business problem, and it requires a business approach to fix it.

Your existing business continuity plan is leaving you exposed if it:

- Requires significant manual intervention
- Accepts loss of data beyond a few seconds for critical systems
- Cannot restore access to critical systems in minutes
- Depends on legacy backup and recovery technology

Backup and recovery techniques were developed for relatively unsophisticated computing processes, back when there where regularly scheduled periods of time when no one would be using the system. The always-on business applications that you now rely on for day-to-day operations need a technology that guarantees continuous system availability and eliminates the threat of data loss, without relying on a backup window.

Modern high-availability (HA) technology continuously streams application and data changes to a remote location. When disaster strikes, such as an earthquake, a power outage, or a bungled software install, failover to an up-to-date copy of your system is automatic and instant. HA eliminates downtime and eliminates data loss.

HIGH AVAILABILITY FOR THE MASSES

High availability is the dream solution if you're looking to insure your systems against downtime and data loss. However, the technology has been viewed as too complex and too expensive for small and mid-market businesses. Many thought that only large enterprises with deep pockets and plentiful IT resources could realistically deploy HA solutions. This is no longer the case.

High availability typically uses a combination of replication and server heartbeat technology to keep IT systems at a remote location synchronized with applications in the primary data center. In the past, this meant dedicated high-bandwidth networks between two physical locations and redundant copies of server, storage, and networking hardware, with specialized applications and operating software. The cost of this redundancy has always put HA out of reach for smaller organizations.

Today, low-cost, high-bandwidth networks are ubiquitous, to the point of being a business necessity. In addition, a wide variety of service providers make it simple to spin up virtual servers on demand at very low cost. These infrastructure advances now mean that HA technology is available to more organizations at a much more modest price tag.

The dramatic drop in HA infrastructure costs has put the business continuity plans of many organizations at an inflection point. Uncoordinated, often overlapping backup solutions are abound in the data center. If you have relied on backup and recovery for business continuity you probably find that these siloed solutions are a maintenance nightmare, sapping productivity, and, more importantly, dramatically complicating disaster recovery.

MODERN HA SOLUTIONS OFFER AN APPROACH TO BUSINESS CONTINUITY THAT LOWERS THE COST OF DATA PROTECTION, SIMPLIFIES DISASTER RECOVERY, AND ELIMINATES DATA LOSS AND DOWNTIME.

High availability may be appropriate for your business, but without performing a detailed analysis of business systems to determine their recovery needs, you won't know which applications will benefit.

TOP 10 BUSINESS CONTINUITY GOTCHAS

Here is our break down of the top 10 gotchas of disaster recovery and business continuity planning.

01 IT'S ABOUT THE BUSINESS, NOT TECHNOLOGY!

Disaster recovery, high availability, backup and recovery, business continuity, call it what you will, the aim is the same: keep the business up and running no matter what the circumstances.

Too often, organizations let technology take the lead and dominate the conversation. What is often forgotten, and is essential to remember, is that disaster recovery is about satisfying a business need and must be driven by business requirements.

Before trying to work out how to implement disaster recovery, you need to spend time thinking about why? Talk to business leaders to understand their priorities. For some it will be email, for others the online order entry system, for others Microsoft SharePoint or other applications. The point is, you won't know what systems are the most important unless you ask business users. Understanding the needs of the organization will let you set priorities that dictate your disaster recovery technology choices.

02 IT'S A CATASTROPHE, OR MAYBE NOT

When you think about disaster recovery, you probably picture hurricanes, floods, terrorist attacks, and the like, not a software upgrade gone wrong with an inadequately thought out rollback procedure or a hardware error on a critical piece of networking equipment. Planning for the worst-case scenario and being tripped up by trivial day-to-day errors is very common. Your disaster recovery planning must take into account all eventualities, from the ordinary to the cataclysmic.

03 HOW CAN YOU ASSIGN BUDGET WITHOUT KNOWING THE COST OF DOWNTIME?

Too often, organizations assign a dollar value for disaster recovery planning before determining the financial risk of downtime and data loss to the business. Unless you can quantify how much you can lose from an outage to critical systems, it will be difficult to state how much you can spend to avoid these losses. Your approach to disaster recovery must be aligned with the needs of the business. This means assessing the financial cost of downtime before allocating a budget. Don't forget to include regulatory compliance in your cost of downtime calculations. There are often financial penalties for unmet legal obligations.

04 IT'S ABOUT MEASURING RISK

Exactly what events classify as a disaster can change from organization to organization, and even from department to department. Some events — earthquakes, for example — are potentially so catastrophic that it is obvious the organization must protect itself against their occurrence. Other events may be common — such as failed network hardware — yet have an outsized financial impact. When thinking about disaster recovery, it is essential to ask: What are we trying to protect ourselves from? Don't overlook the commonplace. Small losses from common problems mount up quickly.

05 DO YOU HAVE A PLAN?

If your disaster recovery plan is a Post-It note on the backup tapes, you're in trouble. As crazy as it sounds, a surprising number of organizations don't have a disaster recovery plan. It is essential that you develop a formal document detailing all applications, hardware, facilities, service providers, personnel, and priorities, and you must obtain buy-in to the document from all stakeholders in the organization. The plan must represent all functional areas and offer clear guidance on what happens before, during, and after a disaster.

06 WE'VE GOT A PLAN, BUT WE DIDN'T TEST IT

Maintaining a disaster recovery plan is only helpful if it works. The only way to ensure your plan works is to test it. Testing the plan under simulated disaster conditions is essential, but it can also be challenging. Performing disaster recovery testing is expensive and takes time and resources away from day-to-day operations. However, unless recovery is fully tested at the application level, you will inevitably encounter difficulties during a real-world disaster. Look for data protection solutions that help you create environments for non-disruptive testing of your disaster recovery plan.

07 WHO IS RESPONSIBLE, AND FOR WHAT?

A real-life disaster event will be chaotic and confusing. If key staff do not understand their disaster recovery responsibilities, the recovery process will be long and fraught with problems. Your disaster recovery plan must clearly state the roles and responsibilities of everyone involved, including what to do if key personnel are not available. These people must also be involved in testing your recovery plan.

08 RECOVERY POINT WHAT? RECOVERY TIME WHO?

It is critically important to understand how sensitive each area of your business is to downtime and data loss. This information informs your disaster recovery technology selections, provides the foundation for your disaster recovery planning, and lets you know the consequences of a failure to recover each business application.

Two metrics are used to record an application's tolerance of downtime and data loss: recovery point objective (RPO) and recovery time objective (RTO). Both metrics are measured as units of time. RPO extends back from the time of the disaster and RTO extends forward.

RPO is a measure of data loss. The larger the RPO, the more data loss an application can tolerate before it becomes a problem for the business. Think of it as the point in time that you can successfully recover data up to. All data between that point and the time of the disaster is gone.

RTO is a measure of an application's importance to ongoing business operations. The smaller the RTO, the faster you must work to get the application back online before the organization starts to suffer significant losses.

If you don't know the RPO and RTO of each application, you're in the dark when it comes to disaster recovery. Whatever you do to ensure recovery after a disaster will be guesswork. RPO and RTO allow you to define levels of service that you can deliver against.

09 RECOVERY WILL TAKE LONGER THAN YOU THINK

For many organizations, thinking about disaster recovery stops when backup tapes leave the data center. But understanding how long it will take to recover key business systems, and how much critical business data will be lost after a disaster, is essential. Even if you can access offsite backup copies, there is no guarantee that you can recover applications in a timely manner. Do you have access to equipment that can read the data? Can you restore the data and rebuild application systems fast enough to satisfy business users? Do you have the bandwidth to recover data from a cloud service provider? Understanding how long it takes to recover applications, and the effect of downtime on the business, may prompt you to make different technology choices.

10 GOING HOME

Going back home after failing over to a disaster site is an often-overlooked component of disaster recovery planning. It's easy to see why. When we think of disaster, our minds focus solely on protecting valuable assets. Little thought is given to what happens to those assets after the disaster event has passed. The ability to failback to production systems is every bit as important as the ability to failover. Unless carefully planned, a backup data center is unlikely to have the same capacity or performance as the production site.

Without a failback plan, you may perform a successful initial failover and then see losses mount as your business limps along for weeks operating from an inadequately provisioned backup site.

WHAT DOES SUCCESSFUL HIGH AVAILABILITY LOOK LIKE?

It's no secret what successful high availability looks like: no application downtime and no application data loss. But is this realistic for small and mid-market organizations?

High availability technology is no longer the complex, esoteric approach to business continuity that it once was. Large corporations have been using high availability techniques to protect their most critical business applications for years. The technology has been tried and tested and is widely accepted as a standard disaster avoidance tool. It is simple, repeatable, measurable, and automated. Technologies such as continuous data protection, replication and automated failover and failback are critical.

Modern HA products have brought the price within reach of small and mid-market companies. This combined with lowered infrastructure costs — broadband, server virtualization, multiple service providers — and dramatically improved usability are making HA a very real business continuity alternative for organizations of all sizes.

Downtime and data loss are a fact of life for businesses that rely on IT. Offsetting this risk with the right technology must be a consideration at the earliest stages of the software development and product deployment lifecycles. Understanding the protection level demanded by each application lets you allocate the appropriate resources. By the time an application is in production use by business users, its RPO and RTO must be clearly identified, and the appropriate business continuity solutions implemented to provide recovery in the event of an outage.

A wide variety of solutions are available that promise to improve disaster recovery, but if they don't eliminate your exposure, they're not high-availability.

Arcserve Replication and High Availability

With Arcserve Replication and High Availability, reducing system downtime and data loss has never been easier. You maintain complete availability across virtual and physical server systems with powerful capabilities that have one common purpose: to keep your business up and running. Confidently deliver on the most stringent service-level agreements (SLAs) with real-time LAN and WAN replication for data onsite, offsite, and in the cloud. Use data rewind to reverse changes made to an application or the OS files, and preserve business operations with automatic and push-button failover, automated end-user redirection, and push-button failback.

How it Works

Arcserve Replication and High Availability synchronizes the data on your servers and a second physical or virtual replica server that you provision locally, at any remote location, or in the cloud. Full System Protection for Windows automates the provisioning for your replica server, using a virtual server, to speed up deployment.

Once synchronized, byte-level changes are continuously replicated from your production to replica server. Application-aware replication, real-time server and application monitoring, automatic and push-button failover, automated end-user redirection, and push-button failback functionality are all designed to provide maximum system uptime. Unlike complex distributed cluster and storage area network (SAN) replication solutions, you get availability for systems and data in one single solution.

About Arcserve

Arcserve provides exceptional solutions to protect the priceless digital assets of organizations in need of full scale, comprehensive data protection. Established in 1983, Arcserve is the world's most experienced provider of business continuity solutions that safeguard multi-generational IT infrastructures with applications and systems in any location, on premises and in the cloud. Organizations in over 150 countries around the world rely on Arcserve's highly efficient, integrated technologies and expertise to eliminate the risk of data loss and extended downtime while the reducing the cost and complexity of backing up and restoring data by up to 50 percent. Arcserve is headquartered in Minneapolis, Minnesota with locations around the world.

Explore more at arcserve.com