arcserve®

WHITEPAPER

High Availability is not a Luxury.

Eliminating Downtime for Small and Mid-Market Organizations

Information technology (IT) provides enormous value for the small and mid-market business, but it also represents a tremendous point of weakness. When markets are global, employees work around the clock and business is effectively always on, any interruption to application availability can quickly lead to lost revenue, lost productivity, lost brand value, and regulatory problems. Taken to the extreme, extended downtime can even threaten the survival of your business.

How then should your organization deal with this type of existential threat? The painful reality is that most organizations do not deal with it well.

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, as the term implies, business continuity is a business problem, and it requires a business approach to fix it.

Here's a quick way to figure out if your existing business continuity plan is leaving you exposed:

- **If your plan requires significant manual intervention, you're exposed;**
- **If your plan accepts loss of data beyond a few seconds for critical systems, you're exposed;**
- **If your plan cannot restore access to critical systems in minutes, you're exposed;**
- **And, if your plan depends on 30-year old backup and recovery technology, you're definitely exposed**.

Backup and recovery has been the go-to technique for protecting IT systems for 30 years, but it was developed in a much simpler time. Backing up data to tape or disk, or performing snapshots — the modern equivalent of a backup — creates a point-in-time image of application data. Restoring from a point-in-time copy is never going to bring your data more current than the most recent backup. Whether your copy is from 15 minutes ago or is two days old, recovering from a backup means you must face the consequences of data loss. That may be OK for some systems. But, for many of your most important business applications, data loss will be catastrophic.

Backup and recovery techniques were developed for relatively unsophisticated computing processes, back when there were 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, be it 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

HA is the dream solution if you're looking to insure your systems against downtime and data loss, but, for many years, HA has had a bad rap. The technology has been viewed as too complex and too expensive for small and mid-market businesses. Common wisdom was that only large enterprises with deep pockets and plentiful IT resources could realistically deploy HA solutions. This criticism has rung true up until recently.

HA 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 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 a universal approach to business continuity that lowers the cost of data protection, simplifies disaster recovery, and eliminates data loss and downtime.

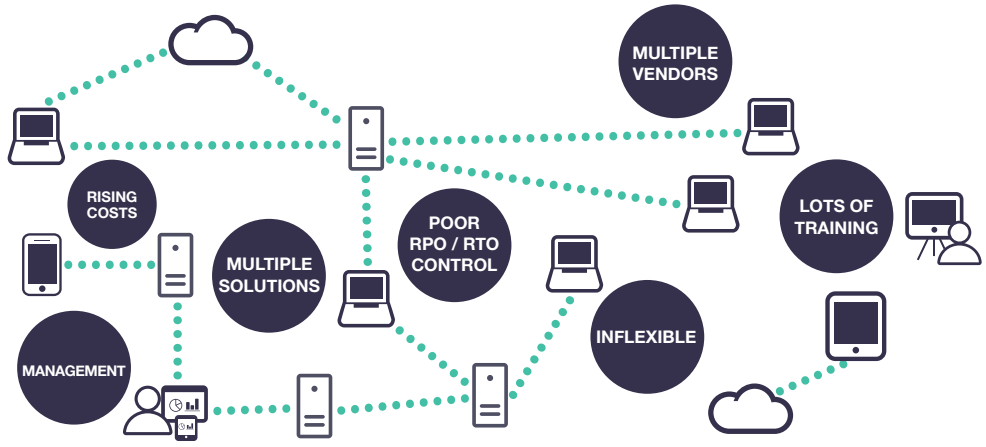

Figure 1 – The hidden risks of complexity in using siloed backup solutions

HA 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. What is certain is that the constraints limiting HA deployment have been lifted, and you're now free to move on to other problems that can dog your business continuity plans.

Top 10 Business Continuity Gotchas

Everyone talks about how to do disaster recovery the right way, but it's just as helpful to take a look at the pitfalls that await you if you do it wrong. Here is our break down of the top 10 gotchas of disaster recovery and business continuity planning.

1 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. 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.

2 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 has to take into account all eventualities, from the ordinary to the cataclysmic.

3 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.

4 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.

5 Do You Have A Plan?

If your disaster recovery plan is a Post-It note on the backup tapes under your sysadmin's bed, 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.

6 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.

7 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.

8 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 have to work to get the application back online before the organization starts to suffer significant losses.

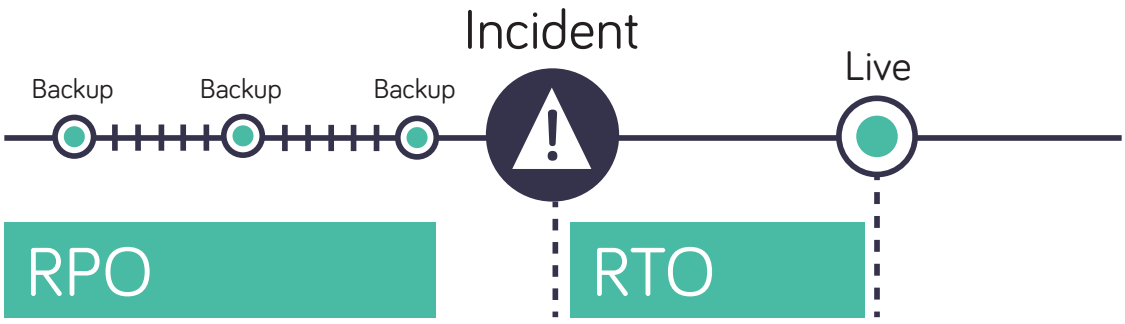

Figure 2 – Understanding how frequently your different applications and data sources should be backed up (Recovery Point Objective – RPO) and how quickly you need it back (Recovery Time Objective – RTO) is a crucial component of your Business Continuity plan.

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.

Technologies like Continuous Data Protection are very critical in ensuring that these objectives can be met.

9 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.

Understanding Risk

With the exception of email, it's almost impossible to know which applications pose the biggest downtime risk to your business without getting input from line of business users. RPO and RTO offer metrics to measure this risk. They also indicate which applications are a priority for your disaster recovery efforts.

Both RPO and RTO operate on a continuum. Think of a timeline with the outage event at the center. The RPO point lies behind the outage event and indicates the amount of data loss an application can tolerate. As the point moves further back from the outage event, the amount of data loss grows and the potential cost to the organization increases.

The RTO point lies on the opposite side of the outage event timeline to the RPO. The RTO shows the amount of downtime an application can tolerate before business losses start to mount up. In other words, how fast you need to get the application up and running after an outage.

If information in the system can be recreated from other sources, losing some data during a disaster might be a headache, but may not cause too much of a problem. For example, missing invoices in the accounts payable system can be recreated by asking vendors to resubmit requests for payment. However, if data cannot easily be regenerated — online customer orders, for example — the loss of this information may directly affect revenue, user productivity, your company's reputation and brand, and regulatory compliance.

Similarly, non-critical business systems — monthly reports from a business analytics application, for example — don't have the same downtime impact on the organization as systems involved in day-to-day operations, like a point-of-sale (POS) application, for example. RTO measures the business impact of application downtime and it can help you determine which disaster recovery tools to apply to the application. Periodic backups may be fine for a business analytics application, but because a POS system is likely to be critical to the business, it will demand a high availability solution.

The difference between the RPO and RTO metrics and actual results from regular disaster recovery testing indicates whether you have an application availability gap. It's worth bearing in mind that an availability gap does not always indicate the use of an incorrect approach to business continuity. Organizations frequently have a wide range of disaster recovery technologies from different vendors, many of which overlap, duplicate, and complicate recovery. Testing can flush out problems and inconsistencies in existing business continuity technologies and highlight areas where consolidating on a single approach or single solution vendor can improve RTO.

What Does Successful HA 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?

HA 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, measureable, and automated. Technologies such as Continuous Data Protection, replication and automated failover and failback are critical.

The maturing of HA products has 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 assured recovery in the event of an outage.

Any business continuity approach that can't make the claim of no downtime and no data loss is not high availability. A wide variety of solutions are available that promise to improve disaster recovery, but if they don't eliminate your exposure, they're not HA.

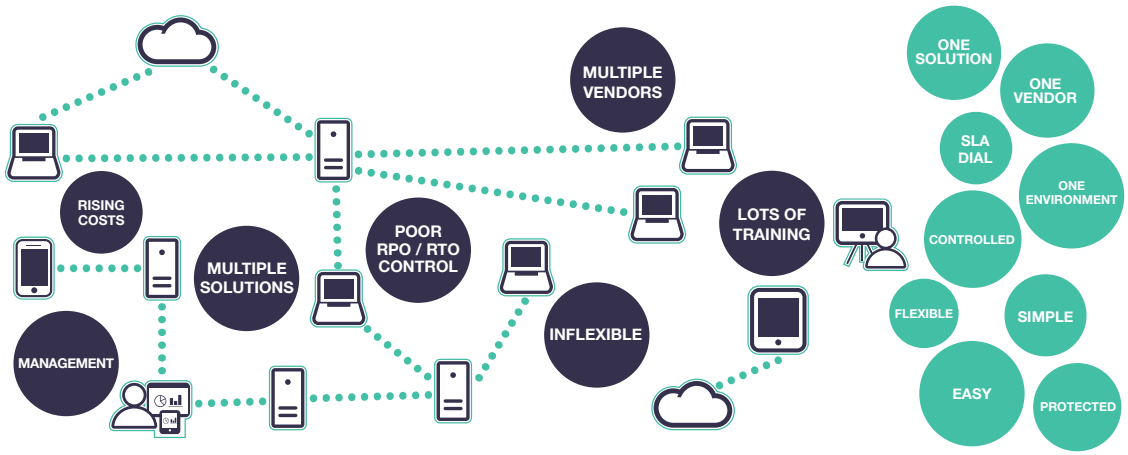

Figure 3 – A unified business continuity solution

arcserve®

WHITEPAPER

About Arcserve® Unified Data Protection

Arcserve has been providing zero downtime protection to organizations around the globe for over 20 years. Now, Arcserve® Unified Data Protection (UDP) delivers a one-stop-shop for all of your data protection and high availability needs. With centralized control, Arcserve® UDP unifies backup, snapshot, replication, and deduplication protection for your virtual, physical, on-premises, and cloud-resident application assets. Arcserve® UDP Assured Recovery™ provides a comprehensive real-time disaster preparedness test process for non-disruptive validation of business continuity plans. To find out more about Arcserve® Unified Data Protection (UDP) and our free 30-day trial, visit: http://arcserve.com/availability

For more information on Arcserve UDP, **please visit** arcserve.com

Copyright © 2015 Arcserve (USA), LLC and its affiliates and subsidiaries. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective owners. This document is for your informational purposes only. Arcserve assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, Arcserve provides this document "as is" without warranty of any kind, including, without limitation, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement. In no event will Arcserve be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, business interruption, goodwill or lost data, even if Arcserve is expressly advised in advance of the possibility of such damage.