

Trends in Server Virtualization

introduction

Server virtualization has come a long way in a very short time. From its early days in IT test and development pilot projects, with VMware's vSphere being the only game in town, to mass adoption, virtualization-first IT policies, and a range of hypervisors available from companies like Citrix, Microsoft, and Red Hat Linux.

Server virtualization has now found its way into IT infrastructures of just about every size. A recent survey of IT professionals by Spiceworks found 80% of small- and mid-sized businesses have already adopted some form of server virtualization.

Virtualization is now viewed as a mainstream architecture, with many companies committed to deploying all new apps in virtualized environments. But it wasn't always that way. The earliest apps running on virtualized servers tended to be less critical to the business. Web and file servers were common candidates.

Pilot environments have less demanding backup and recovery needs and administrators often relied on existing backup and recovery tools, without considering how virtual environments would affect

recovery. Failure of pilot project apps wouldn't have a financial impact on the business and administrators would have time to figure out how to recover apps and data.

Jump forward a few years and it's now routine to move mission-critical applications into virtual machines. Applications like Microsoft Exchange, Microsoft Sharepoint, Microsoft SQL Server, Oracle, and SAP, invariably have demanding availability and data integrity needs. But, often the tools being used for backup and recovery were not built for the new and unique requirements of virtualized servers.

To complicate matters, applications running in virtual machines now frequently share resources. This builds dependencies into the architecture that are often not given much attention when planning backup tasks. All of this requires rethinking the virtual server data protection strategy to make sure recoverability is aligned with the needs of individual applications and the overall business goals of the organization.

table of contents

CHAPTER 1	<u>Trends in Data Center Virtualization</u>
CHAPTER 2	<u>Why Virtualize?</u>
	<u>2.1 Reduced Capital Expenditure</u>
	<u>2.2 Reduced Operating Expenses</u>
CHAPTER 3	<u>Virtualization's Downside</u>
	<u>3.1 Not Everything Can Be Virtualized</u>
	<u>3.2 Increased Cost</u>
	<u>3.3 Server Sprawl</u>
	<u>3.4 Single Point of Failure</u>
CHAPTER 4	<u>Backup Issues Unique to Virtualization</u>
	<u>4.1 Backup Issues For Small & Mid-Sized Companies</u>
	<u>4.2 Agentless Backup</u>
CHAPTER 5	<u>New Backup & Recovery Capabilities</u>
	<u>5.1 In-Place Recovery</u>
	<u>5.2 VM Migration</u>
	<u>5.3 Changed Block Tracking</u>
	<u>5.4 Advanced Capabilities</u>
CHAPTER 6	<u>The Business of Data Recovery</u>
	<u>6.1 Recovery Point Objective (RPO)</u>
	<u>6.2 Recovery Time Objective (RTO)</u>
CHAPTER 7	<u>Virtualization, Data Protection and Business Continuity</u>
	<u>7.1 It's About The Business</u>
	<u>7.2 Downtime Costs</u>
	<u>7.3 Measuring Risk</u>
	<u>7.4 Workable Plans</u>
	<u>7.5 Fallback</u>
CHAPTER 8	<u>Virtualization and Tape</u>
	<u>8.1 Industry Standards</u>
	<u>8.2 Cost and Durability</u>
	<u>8.3 Portability</u>
	<u>8.4 Good Enough for Google</u>
CHAPTER 9	<u>Data Protection and the Hybrid Data Center</u>
	<u>9.1 When is HA Not HA?</u>
CHAPTER 10	<u>Arcserve® Unified Data Protection</u>
	<u>10.1 Arcserve UDP and Magnetic Tape</u>
	<u>10.2 Arcserve UDP 7000 Appliance Series</u>

Trends in Data Center Virtualization



To read the technology press, you could be forgiven for thinking that virtualization is now the entirety of the IT computing landscape. This simply isn't true.

Sure, a lot of companies have now adopted virtualization-first policies that dictate all new apps must run in virtual environments. Moving legacy applications to virtualized environments, however, is another matter entirely.

Despite an increasing focus on technology and the benefits of IT to an organization's bottom-line, [a recent survey by Spiceworks](#) found that IT budgets are at a standstill. IT hiring is also not keeping track with demand for new technology. As a consequence of these trends, IT staffs are continually expected to do more with less.



The inevitable fallout from flat IT spending and reduced IT staffing is that projects involving legacy applications get lower priority. Yes, newly developed apps get to live on virtual servers, courtesy of the dictates of virtualization-first policies. But, the mantra of "if-it-isn't-broke-don't-fix-it" will ensure that, short of natural end-of-life (EOL) driven spending, there is often no budget to move legacy applications on physical servers into the virtual world.

Most organizations, regardless of size, are now dealing with hybrid IT environments containing both virtual and physical servers. This is not an ideal situation, by any means. Hybrid environments complicate just about every aspect of server administration. They reduce administrator productivity, increase the cost of management tools, spread the knowledge of management tools too thinly among administrators, and can lead to compromised application availability and data integrity.

Why Virtualize?

So, what is all the fuss about virtualization? Why are IT organizations so enamored with this technology? The simplest answer is that it saves money.

When business applications are deployed on physical servers, administrators and capacity planners put their heads together to figure out how big a server is needed for the application. They make estimates for how much processing power, how much memory, how much storage, and how much network bandwidth the application needs.



The capacity estimate has to anticipate future growth in demand for the application and peaks in demand. The need to provide adequate capacity for growth means that physical servers inevitably run at, on average, 20 percent of their rated capacity. This means that 80 percent of the capital invested in server processor power, memory, storage, and network capacity sits idle.

2.1 Reduced Capital Expenditure

Virtualization enables a physical server to host more than one virtual server. It also provides the capability to easily move virtual servers between different physical servers to balance demand for resources. The physical servers running virtualization software frequently run at above 80 percent of their rated capacity. The consolidation of business applications on a single

physical server, each with their own discreet operating environment, can dramatically reduce the number of servers in the data center. With fewer physical servers, IT organizations are able to reduce their capital expenditure, freeing these funds for use elsewhere in the organization to increase revenue growth. Of course, there are other benefits, too.

2.2 Reduced Operating Expenses

Reducing the number of physical servers in the data center also saves energy, an important consideration when carbon-footprint is a metric tracked by investors and shareholders. And, it enables a data center to host more applications, a critical factor when data center real estate is becoming more valuable.

From an administrative perspective, virtual machines are much easier to set up, and break down. If an application needs a new server, an administrator can provision the resources much faster than they could

do if setting up a physical server. This often reduces provisioning from weeks to hours, or less, benefiting rapid application development.

Virtual machines are also simpler to administer than their physical counterparts. A single virtualization administrator is able to manage many more machines than if they were looking after physical devices. This benefits administrator productivity and can help alleviate staffing issues.

Virtualization's Downside

There are, of course, downsides to virtualization, too. Not all business applications are appropriate for running on a virtual server. While there are cost savings to virtualization, the technology can also lead to increased expenditure. As with many technologies, uninformed use can exasperate the very problems that it is intended to solve.



3.1 Not Everything Can Be Virtualized

Not all applications are great candidates for virtualization. Applications that are very sensitive to performance may not be a good fit. These apps are unlikely to tolerate sharing physical resources with others and the overhead of running a hypervisor on the same hardware may be unwelcome.

There is a wide variety of applications that require physical appendages to their servers, often with unique driver software. Because hypervisor software has to appeal

to the majority of application use cases, these unusual applications are often not supported.

Not all application software is capable of being virtualized. For some, it can be licensing agreements limitations that prevent virtualization. For others, it could be complexity. Many organizations have older legacy applications that are mission-critical, but are so complex due to many years of upgrades and changes that moving them to a virtual platform would be too risky.

3.2 Increased Cost

There is a cost component that can impact the adoption of virtualization. While virtualization can reduce operating costs in the long-term, there are up-front expenses associated with implementing the technology.

The host servers used to run each virtualization hypervisor must be capable of supporting the

performance needs of all virtual servers. These servers are likely to be more costly than the physical servers they replace.

Server and network administrators must be trained in the art of virtualization. A wide variety of tools are available; many provided by the hypervisor vendor.



3.3 Server Sprawl


Ironically, server sprawl — a condition that virtualization holds out the promise of solving — can, in fact, be exasperated by the ease of spinning up virtual machines. Server sprawl became a significant issue in the data center when servers were being deployed without a sufficient understanding of their impact. This often resulted in data centers full of under-utilized server hardware that consumed precious energy and floor-space.

Server virtualization addressed this problem. Consolidating many physical servers on a single virtual server alleviates energy and floor-space constraints. However, the ease of provisioning virtual machines can lead to renewed server sprawl.

3.4 Single Point of Failure

Finally, a glaringly obvious downside to server virtualization is the fact that hosting multiple virtual servers on one piece of hardware introduces the potential for a single point of failure. If the physical server running the hypervisor fails, all applications running on virtual machines hosted by the hypervisor will become unavailable.

Ensuring data availability and data integrity in a virtual server environment demands a new approach to data protection. Although many virtualization deployments rely on existing data protection techniques, assuming what worked for physical servers will work for virtual servers, virtualized infrastructures pose challenges. Not least of these is that the environment is unlikely to be all virtual.



Backup Issues Unique to Virtualization



As applications like ERP, CRM, and email have moved to virtual machines, data protection in the virtualized environment has undoubtedly become

more important. Unlike less critical applications, these new workloads often have no tolerance for data loss, and leave very little room for downtime. Unfortunately, data protection software vendors have frequently found themselves playing catch-up to rapid changes in the virtualization operating environments.

In the early days of virtualization, there were no officially sanctioned methods for backup tools to interface with the hypervisors controlling virtual machines. Without formal APIs, this led to a lot of ad-hoc approaches to backup and recovery. These approaches were not sanctioned by the vendors of the virtualization software and upgrades to the hypervisor inevitably broke the backup tools, putting data protection at risk.

One of the goals of server virtualization has been to make more efficient use of physical server hardware. Previous architectures would use perhaps ten to thirty percent of a server's CPU, on average, leaving plenty of capacity for periodic workloads like backup. Virtualized server environments now typically see utilization greater than 80 percent. This leaves very little excess capacity for other workloads.

With utilization on virtualized servers running so high, there is less free capacity to accommodate backup software running on the host OS in a virtual machine. This software takes vital resources away from business applications running on the same machine, and can impact the performance of applications running on other virtual machines on the same virtual server.


In addition to backup issues, early implementations of virtual server data protection software had significant recovery issues. Due to lack of granularity, restoring data to a virtual machine was frequently an all-or-nothing proposition. This is a big problem if, for example, you only want to restore a single corrupt file.



4.1 Backup Issues For Small & Mid-Sized Companies

If you're a small- to mid-sized enterprises (SMEs) chances are you are at a crossroads when it comes to choosing the right data protection solution for your data. Moving to virtualized servers is not going to happen overnight. It is highly likely that your environment has a mix of physical and virtual servers running a variety of business applications, and will do for some time. This brings up several issues.

Conventional backup tools are primarily designed to protect physical server environments and are often not virtualization-aware. Your system administrators are intimately familiar with these tools and will have a variety of trusted standards and procedures in place that smooth management of backups and enable swift and accurate recovery.



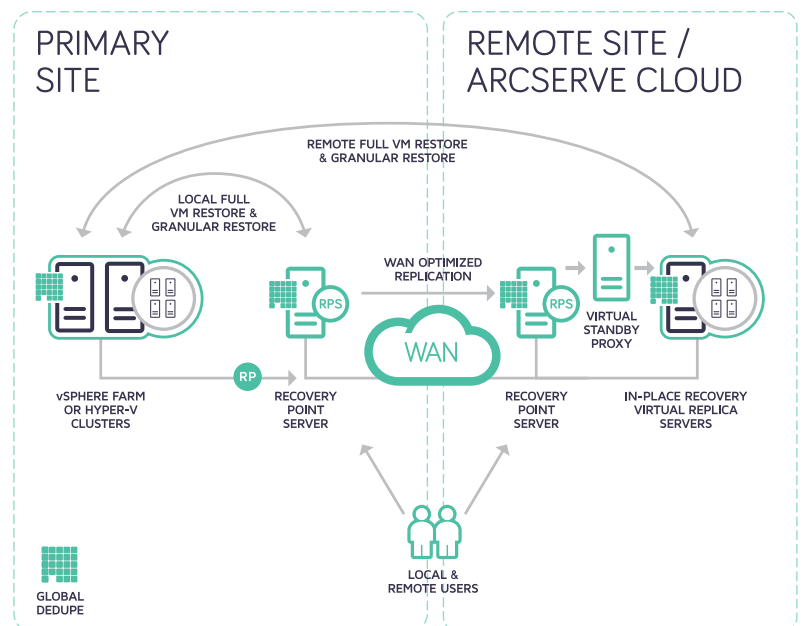
New backup solutions are often designed to only protect virtual servers. Many vendors of backup products were slow to respond to virtualization requirements due to the ad-hoc nature of the environment. This left an opening for smaller, more nimble start-up companies. The new tools these companies developed solve the specific problem of backing up virtual servers but they generally do not provide protection for existing physical servers or for more conventional backup scenarios involving tape.

Many SMEs will find it necessary to support multiple backup tools in the short-term, but this will increase software and support costs and reduce administrator productivity.

4.2 Agentless Backup

When thinking of backup, we need to distinguish between the application and the machine. Traditional backup software requires an agent be installed on the host OS to communicate with the backup server that catalogs and stores backup data. Agent software helps the backup become application-aware.

Virtualization environments are increasingly relying on agentless backup. This approach backs up the entire virtual machine but has less understanding of the applications running on the host.



New Backup & Recovery Capabilities



The unique nature of virtual servers has generated several opportunities for extending the use of data protection techniques. VMware, for example, hosts each virtual machine in a core vmdk file, with several smaller supporting files for logs, configuration, and the like. Being able to capture an entire virtual machine by performing an image-based backup of a handful of files has its advantages.

5.1 In-Place Recovery

The vmdk file contains an encapsulated host operating system and the applications running within it. Restoring this file to a suitable device can allow the virtual machine to be restarted straight from the backup. This is known as an in-place recovery.

In-place recovery can get a failed server up and running quickly. The virtual machine is restored to another server or a backup device and restarted. This gives users almost immediate access to their applications and data, albeit at a recovery point equivalent to the most recent backup. The straightforward nature of in-place recovery means that it is amenable to automation,

further simplifying the recovery process. There are limitations to this technique, however.

If recovering a virtual machine to a backup device, the device is unlikely to have the same capacity and resources as the virtual server being backed up. Applications running on the device will likely have degraded performance and suffer from other limitations. Also, in-place recovery invariably means users will have to experience a second period of downtime at some future time to transitioned the virtual machine back to its original server.

5.2 VM Migration

One of the benefits of server virtualization is that virtual machines are easy to migrate between servers. There are any number of reasons why you might want to do this: moving workloads between physical and virtual servers, or vice versa; host server maintenance and upgrades; and balancing workloads.



A backup of a virtual machine contains everything needed to set the virtual machine up and running for an in-place recovery.


This also means the backup can be used as a vehicle for migrating the virtual machine, eliminating the need to schedule ad-hoc migration tasks.

5.3 Changed Block Tracking

Hypervisor vendors have added APIs into their operating systems specifically for backup software. For VMware the vStorage APIs for Data Protection are specifically for backup. VMware provide their own integrated backup facility, VMware Data Protection (VDP), but third-party vendors have also incorporate the backup APIs into their solutions, benefiting overall stability of the backup tool and giving virtual server admins greater selection of backup solutions to choose from.

APIs also provide greater insight into the workings of the hypervisor, allowing for a much more granular level of data protection. Changed Block Tracking, for example, allows backup software to understand what has changed since the last backup. Like traditional incremental backups,

Changed Block Tracking enables backup software to perform a snapshot of a virtual machine and then periodically back up only changed data. This allows for more frequent backups with much less data being transmitted to the backup server. Recovery is shortened because, after a failure, only changed blocks need to be restored.



5.4 Advanced Capabilities

Virtual machine data protection also lends itself to a number of other capabilities.

Synthetic Full Backups

Synthetic full backups have been a feature of data centers for several years, but have only found a small number of use cases. With the ability to perform periodic full backups of virtual machines followed by frequent changed block backups, it is again possible to create synthetic fulls, which can be used for point-in-time recovery or archived for compliance.

Multi-Hypervisor support

Although VMware was the first modern server virtualization solution, it is by no means the only one. Microsoft Hyper-V is a powerful contender in the data center and there are a number of lesser virtualization solutions from vendors such as Oracle, Red Hat, and Citrix. It is important that any data protection solution support multiple hypervisors. Although there are products that lend themselves to one environment or another, the administrative, maintenance, and productivity costs of maintaining multiple data protection solutions should be kept in mind.

The Business of Data Recovery

It is critically important to understand how sensitive each area of your organization is to data loss when evaluating data protection options.

This information informs technology selection, provides the foundation for your backup and recovery and business continuity planning, and lets IT know the consequences of a failure to recover each business application. This is even more important in a virtualized setting where an outage to a physical server can affect many different applications.

There are two industry-standard metrics used to record a business application's tolerance of downtime and data loss: **recovery point objective (RPO)** and **recovery time objective (RTO)**. These metrics are units of time and indicate how much data application users can tolerate losing (RPO) and how quickly an application has to be back online before the organization begins to suffer significant losses (RTO). RPO extends back from the time of an outage and RTO extends forward.



6.1 Recovery Point Objective (RPO)

RPO is a measurement 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 will be lost.

6.2 Recovery Time Objective (RTO)

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 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. Knowing your RPO and RTO allows you to define levels of service that you can deliver against.

Virtualization, Data Protection and Business Continuity



The unique capabilities of backup and recovery in a virtualized environment can significantly strengthen your ability to protect application data. But virtualization can also create problems if you don't already have a good business continuity plan in place.

When employees work around the clock and business is effectively always on, business continuity planning is critical. Any disruption to normal operations can quickly lead to lost revenue, lost productivity, lost brand value, and potential compliance issues.

Virtualization can enhance your ability to get mission-critical systems back online after an outage.

In place recovery, for example, lets you quickly mount a backup copy of a virtual environment on a different server and restart the virtual machine. This can dramatically reduce the time it takes to recover.

Business continuity is rarely a straightforward matter of restoring a single virtual machine. Interdependencies between applications on different virtual machines and between virtual and physical servers can introduce complications. Avoiding the following five common missteps will help you maintain perspective in planning for business continuity.

7.1 It's About The Business

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. It 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.


7.2 Downtime Costs

Too often, organizations assign a dollar value for disaster recovery planning before analyzing the financial risk of downtime and data loss to the business. If you can't quantify how much you can lose from an outage to critical systems, it will be difficult to know how much you can spend to avoid these losses.

Recovery point objectives (RPOs) and recovery time objectives (RTOs) will help you understand how sensitive each area of your business is to downtime and data loss. Knowing the RPO and RTO for each business application will help you calculate the cost of downtime and enable you to define levels of service that you can deliver against.

7.3 Measuring Risk

Exactly what events classify as a disaster will vary 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 more common—such as faulty 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 can mount up quickly.



7.4 Workable Plans

If your disaster recovery plan is a Post-It note on the backup tapes under your system admin'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.

Maintaining a disaster recovery plan is only helpful if it works, and the only way to ensure that your plan works is to test it. Testing the plan under simulated disaster conditions can be challenging, and time-consuming. Fortunately, today's leading data protection solutions provide the ability to automate testing of your disaster recovery preparedness. This can now be done without downtime and without taking production applications offline. Look for data protection solutions that help you create environments for non-disruptive testing of your disaster recovery plan.

7.5 Fallback

Moving systems back to the production environment 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 fallback 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 fallback 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.

Virtualization and Tape

Tape has been a reliable medium for storing production backups and archives for decades and it continues to play an integral role in the backup strategies of many companies. Disk-based data replication and the cloud have reduced demand for tape as a tier-one backup target, but the benefits of tape are such that it continues to have a role to play. It's critical to include tape backup capability in the evaluation of virtual server backup and recovery solutions, especially if existing physical servers rely on tape backup.



8.1 Industry Standards

One of the benefits of a technology with a long track record is industry standardization. Linear Tape-Open (LTO) has proven to be an enduring format, with the most recent iteration of the standard supporting up to 2.5 TB of raw data per cartridge and transfer speeds of 160 MB per second. These standards are still evolving. Future LTO cartridges are anticipated to store up to 48TB (120TB compressed) and transfer data at terabyte-per-second speeds. For large datasets, the capacity and speed benefits of tape dwarf alternatives.

8.2 Cost and Durability

Magnetic tape is an exceptionally cost-effective and durable medium. Bit Error Rate (BER), mean time between failure (MTBF) rate, and bit rot — the gradual decay of data stored on magnetic media—are all lower for tape, compared to disk, and manufacturers routinely quote a 30-year lifespan for tape cartridges. It is undoubtedly the safest medium for long-term data storage.

8.3 Portability

Cloud storage can offer definite friction-free benefits when storing production data offsite, but it's not appropriate for all applications, or all organizations. For example, it's not unusual for regulators to demand corporate archives be stored in the country of origin. This can be difficult to guarantee when backup data is hosted in a third-party service provider's data center somewhere in the cloud. Backup data volumes can also prohibit network transfers to cloud storage providers. And, as Andrew S. Tanenbaum observed, "Never underestimate the bandwidth of a station wagon full of tapes hurtling down the highway."

8.4 Good Enough for Google

In 2011, Google's Gmail service experienced a very public outage. A bug introduced in a new software release managed to propagate through multiple data centers before it was spotted. Gmail users woke up to empty mailboxes and Google did not have the option to failover, because backup data centers had the same problem. Google eventually resolved the problem by restoring user mailboxes from tape backup. Many industry observers were surprised by Google's use of tape, but it highlights the versatility of the medium. Organizations large and small can learn from Google's experience of recovering from the Gmail outage and apply those lessons to their own data protection and business continuity plans.

Data Protection and the Hybrid Data Center



Today's data centers are more often than not a hybrid, with IT architects selecting the most appropriate infrastructure for business applications based on a list of proven best-in-class alternatives. The cloud now offers an abundance of off-the-shelf business processes available as Software-as-a-Service (SaaS) offerings, for fast deployment and pay-as-you-go economy. Virtualization-first policies frequently dictate the host for new home-spun applications and packaged software solutions that must be run from the safety of the corporate data center. And, physical servers continue to play a role for legacy applications and those processes that demand the highest performance or specific hardware tie-ins.

Ensuring the availability and integrity of data in this new hybrid data center is a challenge. Legacy backup and recovery tools may be great for physical servers and tape, but not so good with virtual servers. New software for backing up virtual servers may not work well on physical servers, and almost certainly won't be robust enough for tape. With IT budgets and staffing at a standstill, it's more important than ever to avoid the trap of having silos of data protection software for each server architecture. This can quickly raise costs and sap administrator productivity.

9.1 When is HA Not HA?


High availability (HA) technology is now a very real option for the mission-critical apps of mid-market companies. The technology is no longer the complex, esoteric approach to business continuity that it once was.

The goal of HA is simple: zero downtime and zero data loss. Tools do exist that lay claim to being HA, but if they are not able to eliminate your exposure to downtime and data loss, they're not HA.

As Google discovered during their Gmail outage, there is a role for conventional backup even when you have an HA configuration up and running. However, there should be no confusing conventional backup techniques for HA.

Arcserve® Unified Data Protection

Hybrid data center architectures are complex enough without adding the additional risk of incompatible backup and recovery solutions. However, this is often the path organizations take when deploying a combination of legacy tools and newer single-purpose data protection technologies. The lack of integration inevitably creates silos that add unwanted management, increase overhead costs in the infrastructure, and complicate recovery – effectively working against the objective of reducing risks to operational data.



Arcserve® Unified Data Protection (UDP) is a single, unified data protection solution that provides the flexibility to support a wide variety physical, virtual, and cloud IT platforms, a diverse range of application RPO and RTO requirements, and an array of backup media, including tape.

Arcserve UDP integrates disk-to-disk backup, tape backup, replication, high availability, and global deduplication in a highly scalable, single architecture. With agentless deep integration to support a variety of hypervisors and a modern

task-based approach to administration, Arcserve UDP is able to automate complex repetitive tasks and provide all data protection and recovery from a single pane of glass management console. Assured Recovery™, a unique capability of Arcserve UDP, offers automated, risk-free testing of disaster recovery scenarios, without requiring end-user downtime.

