

Cloud Backup and Recovery

Software applications and electronic data are the life blood of a business. When they aren't available due to a disaster or outage, business is slowed or stopped altogether. In the short term, outages result in data loss, employee and customer frustration and lost revenue. The long term penalties of an outage can affect a business for a lifetime; lost records, transactions and accounting files can even put a business at risk of regulatory violations.

Protecting business means protect ongoing access to servers and data; traditionally that means backing up the server. However, backing up the server is only part of the equation. If you can't recover the data, the backup is useless.

If your business relies on tape backup alone, restoration is easy only for the simplest failure, and only if everything goes perfectly. If a hard disk fails and all the backup tapes are good and the staff is practiced at doing the repair and restore, then you might be able to simply buy a replacement part and get things up within a couple of hours – though the data will be from last night's backup. If the problem is more complicated and involves a replacement server for instance, you will probably need a day or two to get new hardware in place before you even begin to recover. At this stage, 40% of tape restores fail.

The right way to evaluate the quality of your system and data protection is to evaluate the Recovery Time Objective (RTO) and Recovery Point Objective (RPO). These metrics define how long you think it will take you to get back online and how current the data has to be.

The best protection is off-site, rapid recovery

Every business will have different RTO and RPO goals. If an organization determines it has a four hour RTO and RPO, then the business can tolerate four hours of downtime between failure and recovery – and that they will only have to recreate (or do without) the last four hours of data. Together, this is about eight hours of lost productivity. For most serious problems, it's an optimistic goal for a tape (or disk-to-disk) backup system alone to meet.

Here is a quick list of things that can easily push recovery times into days or weeks:

- Equipment failure, requiring a replacement
- Extended or recurring power outage
- Air-conditioning failure
- Fire
- Flood (water leak)
- Physical damage to the building

The best way to ensure a fast recovery is to have replacement equipment standing by at an off-site location with the necessary software and configuration to quickly transfer users and data. The best practice includes a remote data center with servers, storage, networking equipment and internet access. Restoring to this remote data center from backup tapes will likely take too long, assumes that the tapes were not affected by the original problem and still leaves the risk of only recovering old data. Instead, replication software can be used to keep the backup systems constantly updated.

A four hour RTO and RPO requires:

- Off-site hardware and infrastructure to run servers and applications
- Data updates to the DR site more often than every four hours; preferably real-time
- Continuous updates of the application and OS configuration (without this, recovery may fail after a patch or an upgrade).
- A method to deal with any hardware differences between production and recovery environments.

Cloud Backup and Recovery

Why isn't every server protected?

All the requirements in the list above can be met by current shipping technology. If it is clear that local tape or disk-to-disk solutions do not provide adequate protection and a better solution is available, why isn't every server in the world protected? Usually the answer is "cost". The cost of an off-site, rapid recovery solution comes in a couple of different packages:

- Up front cost
- Technical complexity (requires new IT specialists, or time and budget to train existing staff)
- Operational complexity (managing a new data center and twice as much equipment)
- Project management (complex, expensive projects require lots of planning and management)
- Risk (expensive, complicated projects sometimes fail)

Given all the cost, complexity, time, and risk involved in creating this capability, these projects are often delayed in favor of projects that produce immediate, obvious results such as a web server update or a desktop refresh. For some organizations adding extra servers to an existing off-site location is relatively easy – these tend to be larger organizations with large staff and significant IT expertise. But even in these large organizations, there are still servers that don't make the cut - they are not considered to be critical enough to justify the solution.

"This server doesn't need it, it's not worth it"

If a server is so unimportant that it won't be missed when it fails, perhaps the next question is "why not just turn it off?" The point of this off-site, rapid recovery solution is to preserve as much of the normal operating capability as possible. Customers and business partners don't care that a pipe burst and flooded the data center – they want to know when a business can deliver. If a server is important to meeting a business requirement, it is worth protecting. The question to ask is not "is this server worth the solution"; instead "how do we make the solution practical for every server?"

Where the pain really is

Most of the cost and complexity of this solution comes not from the specialized tools for replication and recovery. Instead the pain comes from, ironically, the extra facilities and equipment, both of which will sit relatively idle most of the time. Specifically:

- Selecting, acquiring, and building out a second data center (or the high cost of renting one already configured).
- Selecting, acquiring, installing, and configuring the standby equipment
- Managing and maintaining the facility and equipment
- Integrating all the parts into a reliable solution

This creates a peak-versus-average problem, where time and money are spent building a redundant data center than can meet the peak capacity of the IT department, but the average utilization of that capacity will be very small. You pay for peak, but only get the benefit of a very low average utilization.

Cloud Backup and Recovery

New Technology opens New Doors

Tape backup was the preferred recovery solution of the 1970s computing era. Disk-to-disk technology is a more recent, but minor improvement. Modern technology and business practice have caused heavy dependence on IT systems, and therefore creates greater business risk from IT failures.

Cloud Computing to the Rescue

Cloud computing gets a lot of press from both supporters and detractors. Some say it is the future of information technology, others say it is the same old thing, repackaged. The only real interest is whether it can be a means to an end. The term 'cloud computing' is so generic (and sometimes misused) that it is nearly worthless in a practical discussion. Two aspects of cloud computing that are interesting to this discussion are Infrastructure as a Service (IaaS), and Software as a Service (SaaS). With IaaS, a service provider delivers raw resources, like virtual machines, storage, and network bandwidth, as a service. With SaaS, a provider layers a specific software solution on top of those raw resources, and delivers that.

Isn't this just Managed Services or Hosting?

Despite the press releases, there are key distinctions between hosting companies and managed service providers and cloud providers. The key differences have to do with how much of the infrastructure (and therefore, cost) is dedicated to the solution. MSPs and hosting companies generally provide dedicated hardware, software, and storage systems to each customer. This requires that the customer specify, pay for and commit to specific capacities in advance. Cloud providers generally provision a customer's current demands out of a pool of capacity. Think of the difference between buying a generator and getting an account with a power company.

The best cloud providers will let you buy capacity in very small chunks, change your usage on the fly, and bill you only for what you consume; all without a long-term commitment to any specific usage pattern or cost. This new financial arrangement breaks the peak-versus-average problem mentioned above.

The opportunity

If an IaaS cloud provider can offer a complete data center, with great enough capacity to meet peak needs (i.e. during a production outage), but only bill for the average usage during normal operations, there is clearly an opportunity to redefine the cost and complexity of an off-site, rapid recovery solution.

By partnering with the right cloud vendor, a manager planning a disaster recovery solution gets access to :

- Unlimited spares for computers
- Disk capacity on-demand
- Free idle bandwidth, with (nearly) unlimited burst capacity
- A data center that is highly optimized and managed for low cost, high reliability, and high security
- Data centers located in multiple countries, to best meet geographic and regulatory requirements

All that remains is to integrate a suitable disaster recovery solution to this IaaS capacity.

Cloud Backup and Recovery

Don't be fooled!

Sound good? It should. There are now real solutions to real problems - but be careful. Many managed service providers and hosting companies are touting their solutions as "Cloud" and many offer disaster recovery or online backup solutions, but do they meet your Recovery Point and Recovery Time Objectives? It helps to ask a few key questions:

- Can they protect all of the servers and applications?
- Can they protect the OS and applications as well as the data?
- Can I actually failover to the Cloud and stay up and running?
- Can I test the failover process to ensure the servers are recoverable?
- Do they provide a mechanism to recover the data/servers without lots of downtime?
- Can I just pay for what I use or do I need dedicated servers in the Cloud?

Once you find a solution that ticks all the boxes you can look to protect every server in the infrastructure. It should be so cost effective that you can just sign up, set it and forget it. Set a reminder for six months to test failover, and ensure you haven't added any new servers.

Double-Take Cloud

Double-Take Cloud is a combination of technologies from Double-Take® Software and Cloud Service Providers (CSP) such as Amazon Web Services. This combination allows you to replicate data and system state information from one or more Windows servers in your production environment into a virtual server – called a Repository - running at the CSP. From this Repository, you can restore entire servers to virtual machines – also housed at the CSP – to resume normal operations quickly and effectively.

Scenario 1: Backup Operations

Maybe the most common use of the Double-Take Cloud solution set is creating an off-site, rapidly recoverable backup system for a production server set. Many organizations do not have the technical staff or physical space required to create and host an off-site datacenter, but still need the ability to recover critical servers and services in the event of a loss at their primary location. Justification of an off-site location becomes even more difficult if only a few servers out of the production server set are critical; many companies simply hope for the best with tape backups carried physically off-site. They're betting everything that alternate space can be found, servers provisioned and restoration performed fast enough for the business to resume with minimal loss. Instead of attempting to provision an entire alternate datacenter ahead of time, or hoping to provision one quickly at the time of an emergency, the cloud can be used to provide the best of both worlds. Double-Take Cloud installs on your critical production servers and replicates data and system state information into a cloud-based Repository. Since only the Repository is running on a day-to-day basis, the amount of infrastructure needed at the CVS is minimal and more easily justifiable. Replication is performed at the byte-level and occurs as quickly as either your internet links or the limits you set within the Double-Take Cloud software allow. This provides WAN-efficient replication of all critical information from multiple production servers, providing near-real-time replication if sufficient bandwidth is available. Even if throttling must be used due to bandwidth constraints, Double-Take Cloud still allows you recover up to the last byte of information that was successfully transmitted to the Repository.

When recovery of one or more critical servers is required, the Repository can be used to push all system state information, Windows and application binaries and all data to a cloud-based Recovery server. This server is provisioned only when needed and so does not need to be maintained, patched or otherwise tended to. It is brought online via your CSP's utilities at the time of restoration and then converted into a replica of your production server at the time of recovery by the Double-Take Cloud wizard. As the production server is not required for this recovery to occur, you can restore your normal operational environment – in the cloud – even in the case of a total loss of production infrastructure.

Cloud Backup and Recovery

Scenario 2: Cloud Migration

The cloud is a viable platform for testing, development, QA and other back-office operations. Many organizations even use the cloud as a primary business infrastructure, creating agile systems that can move within regions and scale on demand. The issue has always been how to safely migrate current back- or front-office servers from the machines they reside on in production to the virtual systems they will occupy with a CSP. While some downtime could be absorbed by the back-office teams, it must still be minimized, and extensive downtime for front-office applications is generally considered unacceptable. There is also the question of application and system compatibility, as the production servers are typically located on physical or virtual hardware that may not match with that of the CSP.

Double-Take Cloud provides a simplified method for migration to a cloud infrastructure with minimal downtime. Current servers are installed with the Double-Take Backup software, and begin to replicate system state, applications and data into a CSP-based Repository. Once the initial mirror has been completed, new information will continue to be replicated right up until the moment you are ready to stop the developers from working locally and have them begin work remotely. During the initial mirror and ongoing replication, the production servers are available and accessible. Prior to the actual cutover, new servers are provisioned within the CSP and brought online with temporary server identities.

At the point of cutover, the production systems in the local environment are shut down, and Double-Take Cloud creates copies of those servers based on the replication information set on the Repository. What would normally take days of coordination and downtime can be accomplished in a matter of a few hours from start to finish, with the hardware differences properly mitigated by the Double-Take Cloud solution set. Once each migration is final, users can immediately begin using the new systems remotely, with the same data, server identities and information that they had while the same systems were local.

Summary

Application and data availability can make or break a business. These days backing up to the Cloud is a simple, cost effective way to get the advantages of a datacenter without the expense of creating and managing a physical center. Double-Take Cloud is a simple way to migrate, backup and restore data and applications to the cloud with minimal downtime and complete control over your expenses.

QUESTIONS?

1-888-674-9495

info@doubletake.com

www.doubletake.com



© Double-Take Software, Inc. All rights reserved. Double-Take, Balance Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.