WHITE PAPER

Going Hybrid: Cloud-Based Disaster Recovery Without Barriers

By Nick Cavalancia

UNITRENDS

TABLE OF CONTENTS

UNITRENDS

Disaster: Many Things to Many People

The term disaster usually invokes fairly dramatic imagery akin to a Hollywood blockbuster movie. But disasters to your business are far less impressive; they simply represent losses in business continuity that need to be made operational again. Some "disasters" are nothing more than a loss of data, an enterprise application, or an entire server. But in other cases disasters can take a bit of a spectacular turn where losses include a location from flooding, fire, tornados or hurricanes.

Are Disasters Expected, Probable, or Both?

Hardware failures, configuration changes, data corruption and human error are far more likely than floods, tornados, and hurricanes. But you need to be prepared for every kind of disaster that may come.

According to a 2014 survey of over 2000 companies, 36% of companies had critical failure of an application that lasted an hour or more within the last year.

So, what's the likelihood of a "disaster"?

According to a 2014 survey of over 2000 companies, 36% of companies had critical failure of an application that lasted an hour or more within the last year, while 20% of companies had a failure that lasted more than a day. Does this mean you will? Not necessarily. But it does give you some idea of what you should expect and be preparing for.

Disaster, in Business Terms

From a business standpoint, situations like a fire, hardware failures, or hurricanes are merely the cause of a disaster. But the actual disaster is the loss of business continuity, loss of revenue, loss of reputation and even potential loss of customers.

Your company is concerned about concepts like a Service Level Agreement (SLA), which defines how quickly you will respond and, at a higher level, what % of uptime the business can expect. So it's that one-tenth of 1 percent of downtime this year that is the disaster. No one cares how it happened; just that it did.

To prepare, you need to take your SLA and dig deeper and put definitions around what needs to be recovered and by when. The first is your Recovery Point Objective (RPO), which defines an acceptable loss of data in terms of how old the data being recovered can be. For example, your most critical customer-facing applications may only be able to lose

5 minutes of data. The second definition is the Recovery Time Objective (RTO), which defines acceptable loss based on how long recovery can take to get back to the RPO.

Like SLAs, RPOs and RTOs need to be defined on a per-server, per-application, and per-data set basis. Once you've done this, if you think about it, you've actually defined what a disaster is, to your company; anything that falls outside these definitions is the disaster.

Dealing with Disaster

Today, there are two traditionally accepted recovery methodologies, on-premises and cloud-based backup and recovery. Each provides a set of unique benefits, while still not providing a complete set of options. On-premises backups (presumably using a NAS device) provide fast access to local storage, but does little when the disaster impacts the entire location. Cloud backups provide highly-available access to backups anytime, from anywhere. But take away an Internet connection and you're in real trouble.

There's one option worth looking at – hybrid cloud – that combines the features of on-premises and cloud backups, promising better, faster anywhere, anytime recovery.

No one solution has all the answers. But there's one more option worth looking at – hybrid cloud – that combines the features of on-premises and cloud backups, promising better, faster anywhere, anytime recovery.

Hybrid-Cloud Backup and Recovery: A Quick Primer

Hybrid-Cloud backup and recovery takes the best of both these worlds, and creates an environment where you have the speed of local backups, with the convenience of cloud-based storage, all working together to provide a unified recovery front. For those new to hybrid-cloud backup and recovery, there are three basic components:

- **An on premises storage appliance** – As a true appliance (that can be implemented physically or virtually), this is an "all in one" kind of device, retaining backups, archiving data, and providing immediate recovery. But it doesn't just stop there. This same appliance is responsible for replicating data with a secondary site, our next component.
- **A secondary storage site** – While this is usually cloud-based storage from an external partner, this can be customer co-location.
- **WAN optimization** – Because you need to get local data up to the

cloud matching as close to local performance as is possible, the storage appliance utilizes technologies such as deduplication, compression, scheduling, replication, throttling and security in transit, to optimize transmissions to and from the secondary site.

When utilizing a hybrid approach, you mature your backup efforts to true Disaster Recovery as a Service. Even if you're part of an internal IT department, you augment and elevate the offerings of how and when you can backup and recover various parts of your IT environment.

With some of you firmly in the Cloud camp, and others strictly using on premises solutions, there are challenges you each currently face when it comes to recovery. Let's look at 5 key barriers to recovery when only using one or the other, and see how moving to a hybrid-cloud model can help.

The whole point of backups is to have them when you need them.

Recovery Barrier #1: Availability

The whole point of backups is to have them when you need them. But let's roll the clock back just a bit and talk about availability from a few other angles where the concept of availability is equally as critical. When planning your disaster recovery strategy and implementation methods, availability should be thought of in terms of backups, data, and archives.

Availability of Backups

The greatest barrier to recovery is not having the right backup set in the first place. Both on-premises and cloud backups can have difficulty meeting RTOs and RPOs, depending on how stringent the definitions are. If an RPO is defined as being just a few minutes ago, and the system in question has lots of changes, it may be too large of a data set for, say, cloud backups to handle in the timeframe given. Or if perhaps, you're one of the few still utilizing tape as your primary backup medium, given its serial recording nature, it simply may be impossible to find a recurring window of time to backup the needed data set.

It's important to also mention that if you only have either cloud or on premises backups, you have a single point of failure – the medium. If the Internet goes out, the cloud is useless as a backup medium. And if your NAS or tape system is not working, on premises is useless.

Regardless of whether you currently use an on premises or cloud-based backup solution, without having the needed backups in place, you'll bring your DR efforts to a screeching halt well before a disaster ever occurs.

Availability of Data

Assuming you do have backups running, when it comes time to recover, you need to be able to have access to your backed-up data. If you're in the on premises side of the house, where are you storing your backups? If on tape, are they truly on premises, or at an off-site facility? What's your plan if a disaster affecting your location occurs? Will your NAS or tape system be accessible or even functional?

Some of your data has very specific retention needs outside of normal backups.

And those of you backing up to the cloud, you don't get off easily either. You have a myriad of issues that can keep you from accessing your backed up data – routers or telco issues, even downtime of cloud providers (which does happen) can all keep you from similarly gaining access to critical data at the time you need it most.

Availability of Archives

Some of your data has very specific retention needs outside of normal backups. External influences such as meeting compliance requirements, legal holds, and the need for eDiscovery all may put certain data in a somewhat separate category when it comes to retention. Regardless of the increase on storage requirements, and complications added to your backup strategy, if you have data like this, it's going to be necessary to keep archives of this data accessible beyond default retention periods.

Removing the Barrier

To achieve true availability of backup mediums, data and archives, you need to remove any single point of failure within these three aspects of recovery. By moving to a hybrid implementation, backups and archives have redundant destinations, and recovery has redundant sources. When Internet connectivity is down, local storage is available. When local storage is impacted by a disaster, recovery in the Cloud comes to the rescue.

Recovery Barrier #2: Speed

When it comes to backup and recovery, speed is a tricky one. Speed depends on many factors – how often do you need to backup, what data sets, how often they change, how much is changing, what is the size of your WAN pipe in the first place, and whether to use the cloud or a secondary site.

The general view is on premises backups run at a predictable rate of GBs/second, but the Cloud appears to be an unknown when it comes to speed. Understandably so, Cloud speeds are largely out of the control of the business, once you've hit the Internet.

Cloud speeds are largely out of the control of the business, once you've hit the Internet.

Those using the Cloud today are likely taking advantage of some of the most state-of-the-art technologies created to drive the size of data to be backed up and restored to an absolute minimum. And we're talking far more than just compression.

Cloud backups analyze the changes made to data, systems, and applications to identify those pieces of data that need to be backed up. Data is then deduplicated – after all, there are only so many copies of the word "the" that need to be transmitted up to the Cloud. The value of deduplication depends on the expanse of it's efforts; if comparing the global data set within your company, you'll only have a single "the" ever. But if backups are deduplicated on a per backup job or per application basis, there will still be duplicate data (from a global viewpoint) sent to and from the Cloud. Once deduplicated, data is compressed, encrypted, and transmitted.

And when it comes to recovery, speed becomes an even more critical factor. In addition to deduplication and compression, the issue of what needs to be recovered impacts speed greatly. There's a vast difference between recovering a few files and recovering an entire server, which is why defining your SLAs, RPOs, and RTOs is so important.

Removing the Barrier

Backup solutions designed for a hybrid approach focus on creating the smallest data set possible – regardless of where the data will be stored. By taking advantage of technologies designed to reduce the size of data

to be backed up and recovered, you will not only gain faster use of the cloud for backup and recovery, and a smaller storage footprint within the cloud (which reduces storage costs), you'll also see faster local backup and recovery with better utilization of your local storage. Add this to the availability and scalability of a hybrid recovery solution, and the only thing left to address is how to recover when the disaster takes away what you were planning on recovering to.

Recovery Barrier #3: Scalability

If you were a DR "prepper", you'd have every single file, folder, application, server, workstation, and network device backed up. In actuality, the prepper mindset really isn't too far from where you should be thinking. In order to resume business, you need to have backups of everything and anything that could be impacted by the disasters you are preparing for. And when you start planning in the general "everything and anything" direction, you're going to need to consider how scalable your current solution is.

You need to have backups of everything and anything that could be impacted by the disasters you are preparing for.

Scalability may be something you're proactively planning for, as in the case of measuring storage needed to add additional employees. But it can also be somewhat reactive and dynamic, in the case of an impending hurricane coming in two days causing you to decide to backup all your workstations just in case. In either case, at that point in time, you need more storage space.

Local storage is the least scalable when comparing it with the Cloud. Yes, local storage can be scaled, but it comes at a price. Two, actually. The first is the price of capital expenditures. Your company needs to budget and plan for large purchases. And that day you walk into Finance wanting to upgrade your 5TB NAS to 10TB, it's not going to be an easy conversation. The second price is the price of utilization. Even if Finance does approve the 10TB storage, you are not even at 5TBs yet, so you'll be wasting much of that additional (and unutilized) 5TBs of storage.

On the other hand, the Cloud easily scales as an operating expense that is dynamically tied to your specific usage… as long as you can get to it.

Removing the Barrier

If you're moving to a completely outsourced Hybrid-Cloud backup and recovery implementation, generally even the hardware is part of the operating expense. This means you get the benefit of obtaining enough needed local storage, while paying for it like it's a service. Scalability is achieved through the Cloud – and locally – through the very same technologies utilized to enhance the speed of backup and recovery, making scalability not only easy to achieve, but less costly.

Recovery Barrier #4: Recoverability

We've touched on the single point of failure from a backup perspective. It equally applies here. The Cloud is only viable if there's Internet access after a disaster. And on premises solutions are useful as long as the hardware and the location are intact. The reality is using either backup methodology independently creates a single point of recovery failure.

Remember, disasters will take many forms.

You can gamble around the likelihood of any given type of disaster, but as data sizes are increasing, and internal systems become more critical to operations, the time a company can be down has decreased. And that means greater pressure on IT if you haven't planned for the specific kind of outage. Remember, disasters will take many forms. This means you need as many recovery options as possible to ensure business continuity.

Beyond the simple recovery of data to its' source server, planning for disasters will require thinking about recovery in terms of what you wouldn't have after a disaster. That standby server? Burnt to a crisp. The entire server room? Flooded like a kiddie pool. Now that you see the disaster possibilities, you need to look for additional recovery methods that make your business operational again as quickly as possible. Here's just two examples of recover methods you're missing:

- **Instant Recovery** - (also known as Virtual Disaster Recovery) facilitates recovery of a previously existing physical or virtual system to a standby vSphere to Hyper-V environment. In some cases, the DR virtual environment is hosted (since the original system no longer exists, it's safe to assume neither does the rest of the server room), or can be hosted on the local storage appliance. Once the business

has fully recovered, the recovered VM can be migrated into a permanent virtual environment.

- **Bare Metal Recovery** - recovers an entire OS, system state, applications and data to similar (or with some vendors, dissimilar) bare metal hardware. BMR makes a recovery process that would normally take a day or more, far quicker and easier with little more than a single click.

If you're using Cloud backups, you probably have some of these capabilities in place, but are relying on recovering entire servers over your Internet pipe. If you're using an on-premises solution, you may have the ability to restore the data sets needed to create an entire system, but you're still in the "as long as my NAS works" corner.

If you're really going to plan for every possible disaster, you need every recovery benefit in your corner.

Removing the Barrier

If you're really going to plan for every possible disaster, you need every recovery benefit in your corner. A hybrid-cloud solution not only gives you the recovery options you need, such as instant recovery – on-premises or in the Cloud, and BMR, but does so with the greatest speed, and with the loss of anything from a single server to an entire location in mind, minimizing your downtime while meeting those hard to hit RTOs and RPOs.

Recovery Barrier #5: Security

For those of you already using the Cloud, security isn't a concern because you've already addressed it. But for those of you thinking about adding the Cloud in order to move to a hybrid recovery approach, it's a concern – and a big one. After all, you're taking your most critical and sensitive data that has resided on-site and putting it into the hands of another party somewhere out there, on the Internet.

Removing the Barrier

So, what security should you be looking for in the Cloud? Simple – security needs to be implemented both when the data is in motion and at rest. A proper hybrid cloud solution encrypts data before its transmitted, sends it across a secure channel, and stores it encrypted, ensuring that in the multi-tenant environment that every Cloud storage provider utilizes, the data is only seen and accessible by you, the customer.

There are various levels and methods of encryption used, so the onus is on you, as the Cloud customer, to identify the ones that meet your company's security standards.

Conclusion

Making the jump from just on premises or from only the Cloud to a hybrid model is a big one. Those using on-premises don't trust the cloud to be fast, secure and dependable. Those using the Cloud are worried that local storage will fail or be unavailable in a disaster. It's only when you stop looking at the other one as an alternative, and start thinking about combining the two that you begin to see the possibilities.

The onus is on you, the customer, to identify the levels of encryption that meet your company's security standards.

By taking a deeper look into the barriers that keep you from looking at a hybrid backup and recovery solution, you find that with the Cloud and on premises storage together, you not only overcome the availability, speed, scalability, recoverability and security limitations of each individual backup methodology, but you also gain the added benefits each uniquely brings to the recovery table. ■

With nearly 20 years of enterprise IT experience, Nick Cavalancia is an accomplished consultant, speaker, trainer, writer, and columnist and has achieved certifications including MCSE, MCT, MCNE and MCNI. He has authored, co-authored and contributed to over a dozen books on Windows, Active Directory, Exchange and other Microsoft technologies. He has spoken at conferences such as the Microsoft Exchange Conference, TechEd, Exchange Connections, and on countless webinars and at tradeshows around the world.

UNITRENDS