WHITE PAPER

Recovering Hyper-V: From the Mundane to Mission Critical

By Nick Cavalancia

UNITRENDS

TABLE OF CONTENTS

UNITRENDS

Organizations are demanding IT look at better ways to not just recover, but actually provide assurance that they can do so quickly.

Not every server in your Hyper-V environment is as important to the organization as the next. Some servers fall into the mundane category, such as a file server, where data changes infrequently and can be forgone in an emergency, while others sit firmly in the mission critical category, such as an ERP system, requiring recoverability as close to real-time with zero loss as is possible. And still others sit somewhere in the middle where they are an absolute necessary part of operations, but the business can do without them for an acceptable period of time.

Given the varying levels of importance the servers in your virtual infrastructure have, your "recovery response" of sorts, needs to align – providing a more and more mature recovery strategy and execution ability as you move up the server food chain. It all boils down to needing an increasingly higher level of recovery assurance – the confidence that you have not just the data or OS image, but a true ability to bring business back into a complete and operational state and an appropriately fast rate of recovery.

On-premises recovery is a staple of every organization, but with the advent of cloud storage, offsite virtual infrastructure, and managed services offerings to provide the most speedy and impactful recovery response, organizations are demanding IT look at better ways to not just recover, but actually provide assurance that they can do so quickly, meeting the recovery time objectives (RTO), recovery point objectives (RPO), while meeting budgetary constraints.

So, what's the right way to recover a Hyper-V environment?

To find out, let's look at three levels of recovery, focusing in on the benefits each brings, and look at what you should expect from each with regards to RTOs, RPOs, and cost.

Data Protection

Data Protection of Hyper-V is the most basic of backup and recovery that you should be doing already. On-premises disaster recovery efforts fall into this category, where backups of your Hyper-V environment are accomplished at a few levels. In some cases, backing up the entire hypervisor makes sense, capturing every guest server OS. In other

cases, backups at a guest OS level are necessary to capture application-specific datasets and services, or shrink the dataset and allow for more frequent, smaller backups and shorter RPO's.

The use of public cloud is also a somewhat mandatory "option" here, and, at this level of recovery, it's strictly for storage. Vendors like Amazon, Rackspace, and Azure (and many others) provide this additional layer of data protection by delivering long-term storage and archiving of both your guest and hypervisor-level backups. It's mandatory, in that without it, you're only able to protect against the lost of data or a server on-premises. But in the case of a loss of location, without cloud storage of backups, you have no easy ability to recover when your backups may be subject to that same loss.

When trying to plan out your RTOs, the biggest unknown is the network and how fast can data be brought back.

What to Expect with Data Protection

At this level, Data Protection simply is about being able to restore backed up data. Whether pulling from on-premises disk or tape, or from cloud-based storage, Data Protection only gives you a level of confidence that you have access to Hyper-V backups, regardless of the level of disaster, in order to begin the recovery process.

So what will Data Protection-level recovery of Hyper-V look like?

Let's review a few important recovery considerations and see just what you should expect should you choose to go the Data Protection route.

Recovery Time Objective

When trying to plan out your RTOs, the biggest unknown is the network and how fast can data be brought back. Assuming the cloud as part of your strategy, when backups are sent to the cloud there's a lot of data optimization used – deduplication, compression, encryption and only sending of changed blocks of your images. So the data going there with each backup is limited in size, making the network not as much an issue.

But when recovering that data you need to rehydrate it completely, so it needs to be brought back on-premises in full, putting the RTO in terms of days, if not weeks, depending on the size of the data set and the

bandwidth available to pull the data down from the cloud. Even if you're backing up at a Hypervisor level and storing it in the cloud, you still need to pull those tens or hundreds of GBs back down to reconstruct the Hyper-V server.

To speed RTOs up, taking a hybrid approach – where backups are synchronized both on-premises and in the cloud – improves your ability to quickly recover. Using fast local disk as often as possible speeds up the recovery process, with the added protection of having another recovery source should your entire site be unavailable.

Using fast local disk as often as possible speeds up the recovery process.

It should be noted that for those of you looking to utilize Data Protection for your more critical data or servers that fall more into your tier 1 or 2 categories, you should be looking at some form of on-premises Instant Recovery or Bare Metal Recovery to get to data and/or servers up and running in a matter of minutes.

Recovery Point Objective

This is a simpler discussion point, in that your RPO is defined more by how often you backup this data set, than it is by a desired timeframe at recovery. By definition, you're going to be using Data Protection for tier 3 and 4 services, which are neither entirely business critical, nor have massive amounts of data that change rapidly.

Depending on the service in question, you're likely going to be backing up servers at this level every 4-24 hours. This should be your RPO expectation as well.

Cost

This is without question the least costly recovery option. Your only costs are the on-premises backups you're already doing, mixed with perhaps the added monthly operational costs of cloud use for storage and archiving. This is a perfect example of matching the criticality of the servers that fall into this level of recovery with an appropriate level of spend.

While the business is fine with your lower tier data and services being a mostly on-premises function with somewhat longer recovery objectives,

more critical servers will require a prompt recovery response with a degree of recovery confidence – even in the face of the most catastrophic disaster.

Disaster Recovery as a Service (DRaaS)

Your lower level servers merely need to have their data restored, or their images recovered. But as you move to servers that provide more operationally important (and often critical) applications, such as directory services, or email, you quickly realize that taking days to recover isn't acceptable. For these servers, you need an ability to recover with a higher degree of assurance that, regardless of whether you lose a server or an entire site, you know you can have them running and available quickly.

DRaaS is accomplished using a private cloud-based virtual infrastructure, on which you have a set of servers and services lying in wait.

That's where DRaaS comes into play. Since the term gets thrown around in the industry, often having varying definitions, let's first define it and how it benefits your Hyper-V environment.

DRaaS takes all that work you'd need to be doing with Data protection, and offers it to you as a service, where the work is already done and all you need to do is perform a few tasks to get even an entire Hyper-V server with multiple guest OS images on it running very quickly.

DRaaS is accomplished using a private cloud-based virtual infrastructure, on which you have a set of servers and services lying in wait. This recovery infrastructure contains warm standby snapshots of your hypervisors and their Guest OSes that are sent to the provider as part of the backup process, ready to be spun up during disasters or tests.

Because your more critical tier servers would use this level of recovery, you need more assurance of recovery than just having images lying in wait. So DRaaS providers usually include semi-annual or periodic testing of the standby images, as part of the offering to heighten the level of confidence in making an operational recovery.

The recovery process itself is a manual task, involving the spinning up of servers and services, failing over of relevant locations and users to the

virtual infrastructure, and then failing back to an in-house data center once the disaster has passed and locations have been reestablished.

What to Expect with DRaaS

DRaaS is a far more sophisticated offering than just Data Protection. With standby virtual infrastructures, testing of backups, it stands to reason that your expectations of this level of recovery should be higher.

So does DRaaS come with better recovery of Hyper-V than Data Protection?

To find out, let's look at some of the same considerations as before and see how DRaaS stacks up.

The rule of thumb for RTO with DRaaS should be in the 1-4 hour range, with the RPO being in the 0-24 hour range.

Recovery Objectives

We're grouping both the RTO and RPO in together, because they typically aren't as much a concern at this level. You've got Hyper-V servers already lying in wait with the latest images, so the only real unknown is the amount of time it will take to spin up all of the necessary servers in a proper order, ensuring applications are running and available.

In order to provide some guidance from more a "Is DRaaS right for this particular server?" standpoint, the rule of thumb for RTO with DRaaS should be in the 1-4 hour range, with the RPO being in the 0-24 hour range. When evaluating DRaaS services it's important to get a Service Level Guarantee (SLA) for each VM that will be recovered, the time it takes for users to access applications on the warn standby VMs. Additionally, it's important to understand how your data will be recovered back to the on-premises site after the disaster event is repaired. Remember, this could be Gigabytes or even Terabytes of data and could take months if the only option is to recover across a WAN.

Cost

It's evident that this is a whole other level of recovery from just simple Data Protection, so there's going to be some ongoing operational cost for the provided dedicated infrastructure. This cost is typically a subscription service coming from an operational budget (OPex) versus the capital cost of operating your own data center with equipment, personal and maintenance cost. Remember, the reason you'd use this level of

recovery is because the criticality of the servers to your operation demand it. Also, this level of assurance is married with a concern for recovery in the case of a catastrophic loss of operations.

While DRaaS will provide your more critical Hyper-V servers an ability to be instantly available, the idea of even a 4 hour RTO and the potential for something to go wrong during a manual failover just isn't good enough for tier 1 applications that your business cannot afford to have unavailable. Those applications require an even higher level of recovery assurance.

Recovery as a Service (RaaS)

RaaS is reserved for the most critical of applications, like a reservation system for an airline, or an ecommerce application. RaaS takes DR as a service and elevates it to true managed recovery. Start with the concept of a DRaaS site, but let's change a few things about it to provide a higher level of recovery assurance.

RaaS fits into where you have complex, multi-tiered applications as such as Exchange, SAP, Oracle, SQL, and SharePoint.

First, with RaaS, your application would have a dedicated hot site, ready for failover. Every time you perform a backup of the tier 1 application, the back is tested, verified, and certified to ensure your company can recover from each and every backup made. But the differences don't just stop there – remember, the recovery itself is now provided as a service.

Additional layers of services are included that assure you can absolutely recover. Testing is about more than just the restore processes, but ensure that systems, applications, and data are all in sync, bringing, potentially, your entire network back up and running. Failover and failback becomes automated, connecting users to the failover infrastructure in as seamless a manner as possible and moving servers and data back to their on-premises equivalents once the catastrophe concludes or a new permanent system is recreated.

RaaS fits into where you have complex, multi-tiered applications as such as Exchange, SAP, Oracle, SQL, and SharePoint, where virtualization is a natural fit as multiple VMs can sit on a single physical box. But you also have many services that need to function harmoniously – DNS, AD, web services, backend databases, etc. that all need to spin up in a particular

order with the latest data sets, and all be running to provide what is perceived as a single service to the end user in a very short timeframe.

What to Expect with RaaS

This is obviously for the most critical of infrastructures, so your expectations should be equally as high. RaaS is about integrating recovery with a partner that knows as much about your virtual infrastructure as you do, and has planned, tested, and is ready to automatically failover your environment at a moments notice.

But how different is RaaS from DRaaS when it comes to recovery objectives and cost?

With RaaS in place, RTOs reduce to a range usually less than 1 hour.

Recovery Objectives

Like DRaaS, only the shortest of RTOs/RPOs exist with RaaS, making it less an issue. The difference here is that the real work of recovery – the planning and testing – is done prior to the disaster, leaving only the automated failover to kick in. With RaaS in place, RTOs reduce to a range usually less than 1 hour, with the RPOs remaining similar to that of DRaaS at anywhere from 0-24 hours, but usually are typically 15 minutes or less.

Cost

This is obviously going to be the most expensive option, but still usually available as a subscription service. The value of RaaS is having your users being able to continue to operate as a company – email, ERP, ecommerce, etc. all up and running in a matter of minutes. Given the importance of the application tier appropriate for RaaS and the potential public visibility of a lack of services, the expense of RaaS has to be balanced against the cost of downtime (which Gartner estimates as high as $300K per hour), and reputation loss.

Assuring Hyper-V Recovery

Like most enterprises, you have multiple tiers of applications, each with its own level of importance to the business. As those levels increase, so does your need for assurance that you can get back up and running at increasingly faster rates. Multiple levels of recovery exist – even hybrids of the three mentioned in this paper and no single level of recovery

mentioned in this paper is the right answer. The key is to determine which level of recovery provides the needed level of recovery assurance.

To accomplish this, it's important to identify recovery objectives on a per-application basis, measuring the application's importance to the business. These objectives will themselves dictate the appropriate level of recovery to use, while also self-justifying the expense.

You can take advantage of the right Hyper-V recovery method that will match your business needs.

By looking at your Hyper-V environment as its individual servers, applications, and data, and aligning those components with fitting levels of recovery rather than addressing your Hyper-V environment's recovery as a singular strategy, you can take advantage of the right Hyper-V recovery method that will match your business needs. ■

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