WHITE PAPER

Resiliency: Future-Proofing Your Data Protection

By Nick Cavalancia

AXCIENT™

TABLE OF CONTENTS

AXCIENT™

While technology moves ahead at rampant pace in feature set and functionality, there is no convergence between technology advances and the needs of data protection.

Data Protection strategies have changed over the years. With humble beginnings of simple tape-based backup and restore have been replaced time and time again with increasingly more complex strategies that both meet the current demands of your organization, AND take advantage of available technology.

So, as you think about your data protection strategy, and how it needs to change over time, to continue to meet the business needs tomorrow, next year, or even five or ten years in the future, it's reasonable to wonder if it's even possible.

Roll back the clock 10 years and just think about the changes that have occurred since 2005. Data protection revolved around a desire to have high-availability via clusters, load balancing, and raid systems. Virtualization was just gaining ground and organizations were still looking to backup & restore mixed with hot or cold site implementations to recover. While there are remnants and evolutions of some of the 2005 recovery technologies, you can see how the tech you commonly use today wasn't, in some cases, even conceived of back then.

Now, think about the changes you'll see in technology 10 years from today. And when you consider applying your current data protection strategy to the tech that will be in place in the future, it becomes clear that just like basic backup and restore of files (once considered a great strategy) isn't necessarily viable today, the same will likely be true 10 years from now with your strategy today.

While technology moves ahead at rampant pace in feature set and functionality, there is no convergence between technology advances and the needs of data protection. Sure, you can take advantage of the latest and greatest tech to improve your data protection implementation, but isn't necessarily designed for data protection. Take virtualization, for example. It was first adopted as a way to optimize compute resources and server infrastructures, which includes the ability replicate VMs. But because backup, as part of a data protection strategy, has long been seen as a back office process that isn't strategic to the business, the result is utilizing replication of physical servers into a virtual world as part of a data protection strategy.

It's obvious that new technology won't be helpful for data protection (it certainly is in the case of virtualization), but you'll always be playing a constant game of catch-up, realigning your data protection strategy to changing technology.

So what steps are necessary to future-proof your data protection?

In this whitepaper, we'll look at three specific actions you can take that will help to make your data protection strategy future-proof.

Introduce Resiliency Into Your Thinking

The emergence of new technologies created a clear distinction between primary workloads, such as your tier 1 applications that keep your business running, and secondary workloads, like backup. And so, when building architectures, you concern yourself with how to deploy and optimize the architecture to support those primary workloads (with virtualization, obviously, being a great way to do that), next look at how you replicate those systems, create a test/dev environment, etc., and then, lastly, decide how you protect the data and systems, archive, maintain compliance, etc.

Instead, you need to unify these avenues of thinking about your architecture and be looking at it holistically – from deployment all the way through to protection, but do so by thinking about it from the perspective of resiliency.

Now, when someone asks you about your data protection strategy, you are likely thinking about it in terms of recovery. That's not a bad thing; it's certainly better than constantly being focused on just having everything just backed up, with no plan for recovery. But, even thinking about only recovery can limit you to consider data protection in very specific terms – an application, a system, a certain kind of outage – and not about the business as a whole.

With resiliency, you're choosing to begin the discussion at a much different point than when you're talking about recovery, backup, or even architecture. You see, with resiliency, you're considering first how the business will bounce back from anything that interrupts it

from functioning – and then working backwards to recovery, backup, and architecture.

So, what exactly does resiliency mean for your organization?

At its foundation, most just think resiliency means to bring systems back up and running. Traditional thinking around resiliency has been done in a tiered application manner, that is, what do tier 1 applications need, then tier 2, etc.—thinking about them separately, setting individual SLAs, from minutes to days. But, because we live in an always-on world, where there is an expectation of accessibility on the part of users, customers, clients, and partners to applications, data, and resources within the organization, resiliency takes on a new meaning. In essence, all elements of the IT architecture need to be resilient and not just specific applications.

IT resiliency is so much more than a disaster; it's about anything that can interrupt your business. Therefore, to properly introduce resiliency, you need to look at all of the interruptions that can impact resiliency...

Then, how do you plan for resiliency, when you don't know what form a disaster will take?

First off, the concept of a disaster is really a misnomer. It means many things to many people. And a disaster, in the "act of God" sense—hurricanes, tornados, etc.—is not likely to happen. In fact, natural disasters are not the most common cause of IT interruption.

IT resiliency is so much more than a disaster; it's about anything that can interrupt your business. Therefore, to properly introduce resiliency, you need to look at all of the interruptions that can impact resiliency; not just an outage of tier 1 services, but any type of interruption, including:

- **Human errors** – corruption, individual data loss, virus, unintentional deletion
- **Application failures** – corrupt databases, failed services, disk space issues
- **System failures** – OS or hardware issues impacting a server or cluster being down
- **Environment** – everything from office fires to true disasters

When you begin thinking holistically about resiliency and the interruptions that can cause outages and a lost of productivity it creates thoughts all the way back to a very different architecture that actually incorporates data protection. And it's this thinking in terms of resiliency (and the work product that comes because of it in the future) that will start you down the part of future-proofing your data protection.

Now, very few of you are in a situation where you have the immediate opportunity to re-architect your entire environment, so let's take a look at what you will require in terms of data protection to achieve resiliency.

Determine What You Need from Data Protection

Providing resiliency for your entire IT architecture is no simple feat, especially when you are, at least for the time being, stuck with the environment and architecture you have. So, to provide some guidance, begin with some basic steps that will help take you from the purely conceptual "let's make our business resilient!" to an actual plan of action.

1. **Start with the interruption** – You can't protect your data without first knowing what you're protecting it from. So, defining the interruptions to your business that you want to protect against (such as loss of data, system, application, location, or operation, and even cyber threats or ransomware) is the initial key step. Do note that you may be defining interruptions differently for each tier of applications, for a specific location, for the entirety of IT operations, or something in-between. No matter how you slice up your protection, build a list of ways that part of your business can be interrupted as the basis for a new data protection strategy.

2. **Assess where you are today** – For each interruption you've defined, you next need to determine whether you can achieve resiliency. For some of you, backup and recovery may not even be a secondary workload in your environment (it may be tertiary, or quaternary in priority). So it's important to not just

think about whether it's technically possible to recover against a given interruption, but actually determine if what you have will actually provide resiliency. For example, you may think you have an ability to recover multiple servers from failure within, say, an hour because you have backup snapshots of each VM stored in the cloud, but when you play the plan forward and consider how those all will boot up in order of service dependency, and how some applications will need additional database restores to bring databases into a consistent state – and you quickly realize you're not ready to protect from your specified interruption. What you thought would take an hour is more like a day to days, at best.

Today's modern data center already provides many technologies that can be leveraged to achieve resiliency including server, desktop, network, and storage virtualization.

3. **Define a business impact threshold** – Bringing the previous two steps together, you have on the one hand a set of interruptions and, on the other, the current state of resiliency with a given recovery time and recovery point. Use these combinations (done repeatedly for each interruption you've defined) to understand the business impact and spell out what you'll likely lose in a given interruption scenario. By comparing what kind of loss the organization can actually handle, as well as what IT can provide from a resiliency perspective, you'll begin to see where the threshold is, and where your current data protection strategy is lacking.

4. **Determine how will technologies facilitate resiliency** – This is probably one of the most critical points in future-proofing. It's here that you take those deficiencies in your ability to achieve resiliency and begin to think about data protection in the same way as every other consideration put into architecting environments and applications. Today's modern data center already provides many technologies that can be leveraged to achieve resiliency including server, desktop, network, and storage virtualization.

These four steps should leave you in a place where you have a clear definition of what situations you need to protect the business from, whether you can actually provide protection today, and what

technology is missing to achieve resilience and future-proofing your data protection. Once you've gone through the four-step exercise, you should be recognizing that you can't simply rely on your own modern data center meet the resiliency goals. You're also going to need some help beyond the four walls of the organization.

Leverage the Cloud

The vast majority of organizations have already embraced the cloud for storage. But the use of cloud for compute is still relatively new. One of the benefits of the cloud is it's elastic nature, leveraging a vast pool of storage, networking and compute resources. And when you're looking at an interruption that includes a loss of application, location, or operation, you require that elasticity to provide you exactly the resources necessary—whether a single server, or an entire environment—at the moment you require it. With resiliency defined in a very small period of time—usually within minutes or hours, depending on the level of interruption—the cloud as a not just a recovery point, but a resiliency point is critical in the success of achieving resiliency

You should be looking at the cloud for two very specific uses to create the necessary level of resiliency that places you firmly on the path to future-proof data protection. The first is the use of virtualization, and the second is the use of cloud services.

Virtualization

While it's probably safe to say it's already a part of your product environment, virtualization may not yet be a part of your cloud-based data protection strategy. What has made virtualization so attractive in production is that it allows you to optimize a pool resources across a variety of apps and virtual instances. The same benefit that drove adoption of virtualization in production is the very same one that is driving use of the cloud for data protection. The simple ability to replicate the state of a VM to a cloud-based provider at a frequency that meets your resiliency goals, allowing you to recover services within minutes with little loss of data makes virtualization a no-brainer, and defines why it should be an integral

part of both your resiliency plan overall, as well as your use of the cloud.

So, if you're a big believer and user of virtualization, and yet your resiliency plan involves building out dedicated infrastructures lying in wait, you've somewhat forgotten the very reason you're using virtualization in the first place. After all, it's impossible to achieve resiliency using nothing but physical servers; it's simply too costly and not fast enough to achieve the goal.

Trends in technology have always led us down a path where we focused on the cool things we could build, with how to protect it as an after thought.

Services

When you consider interruptions that require use of the cloud to provide resiliency, you need to consider whether your IT organization has the internal resources and expertise to perform recovery itself, or whether it should be provided as a service. Sure, a simple failure of a single server is easy enough, but if you're properly thinking about resiliency in terms of the entire IT architecture, you may quickly realize you can't do it alone.

Your IT organization is already overtaxed, underfunded, and pulled in many directions. And with all this thought leadership on how to approach your data protection strategy to achieve resiliency, it's almost like you've been given a recipe to build a replicated infrastructure in the cloud, test out the, say, 9 interruption scenarios you've decided upon and build an automated resiliency plan for each one.

The reality? You just don't have time to do that.

Like any problem, there are multiple ways to solve this. Those of you that can afford to build out and staff the process of resiliency, there are tools and solutions that can be used together to do it completely yourself. You can outsource service elements of your recovery as well to the cloud.

For example, you may outsource the infrastructure to be maintained by another organization, using the cloud for infrastructure as a service (IaaS), but leaving the data protection and recovery processes

to your internal IT team. Some of you will look to a cloud provider to perform the actual recovery process itself, thereby providing more disaster recovery as a service (DRaaS). And still others will look to outsource the entire resiliency process, looking to a cloud provider to certify backups, create automation playbooks for various interruptions, perform recovery testing, perform the failover, as well as perform the failback—all known better as recovery as a service (RaaS).

Becoming Future-Proof

Trends in technology have always led us down a path where we focused on the cool things we could build, with how to protect it as an after thought. Future-proofing your data protection requires a convergence of data protection and technology—in both the way you think about technology before you adopt it, as well as the way you use it once adopted to protect the organization.

This convergence leads to resiliency being a primary driver as technology changes—along with performance, availability, usability, and ROI—leading you to break the cycle of adopting the latest technology first and then realizing you need a way to protect it. By placing a focus on resiliency within your organization, data protection will become a primary workload—equal to that of the very systems and applications it protects—affirming your ability to protect the organization, no matter what new technology may be coming next. ■

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.