

# The Future of Backup Appliances in a Complex IT World



*By Nick Cavalancia*

## TABLE OF CONTENTS

THE GROWING BACKUP DATA PROBLEM .....	2
THE GROWING COMPLEXITY PROBLEM .....	3
ENTER: THE BACKUP APPLIANCE .....	6
SIMPLIFYING THE COMPLEXITIES OF BACKUPS WITH A BACKUP APPLIANCE.....	7

***There was a simpler time, not so long ago, when backups were thought of as strictly a local process.***

Backups today are anything but simple. There are so many choices around the right way to do backups. There's disk vs. image, local vs. cloud, physical vs. virtual, standby images vs. continuous recovery... and so many other factors that need to be considered.

There was a simpler time, not so long ago, when backups were thought of as strictly a local process; performing backups to tape (and more recently, disk) on-premises. And then if anything related to disaster recovery, archiving, or long term retention was desired, the effort needed was little more than placing those backup copies into a box and shipping them off somewhere to be stored.

The Cloud completely changed that model, shifting parts of backups – such as storage – up to the cloud *because they can exist in the cloud*. While potentially improving backup and recovery through use of the cloud, it isn't necessarily always the right choice. We've also seen improvements in backup and recovery that exist *because of the cloud*. It's a slight differentiation where rather than going to the cloud because it's an option, parts of backup move there because they are the *best choice*. For example, we've taken the manually intensive error-prone process of moving copies of backups offsite and, because of the cloud, transitioned this into something far more automated and reliable. And by moving backups off to a cloud environment, they can be retrieved far more easily.

But the definition of what should be backed up is in a constant state of flux. As your organization grows, so does the data set to be protected, as does the level of application complexity. It's a constantly moving target that requires additional *automation and intelligence* to help keep the backup and recovery process relatively simple.

This necessitates us to have more than just backup software and storage on-premises; to really take advantage of both the storage and compute available in the cloud for recovery, a backup appliance should be part of the consideration. Backup appliances integrate storage that is both local and cloud-empowered, along with simplified backup and recovery automation to help you easily recover in situations that would previously have been challenging.

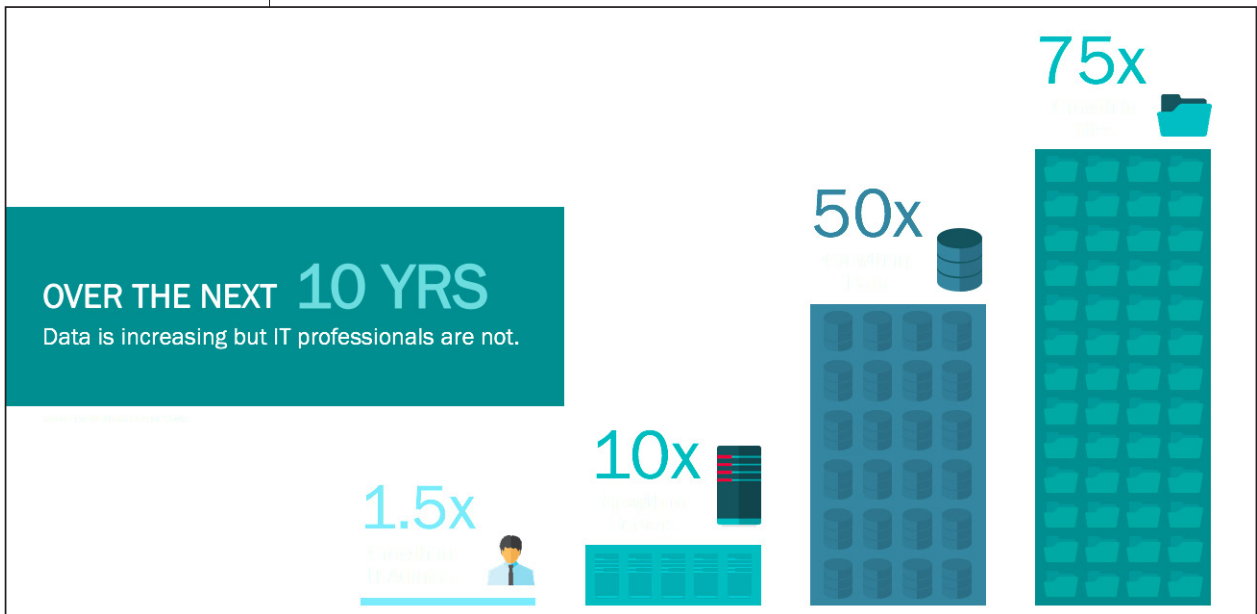
*But, is a backup appliance alone going to provide the right recoverability and protection needed?*

To find out, let's look at two trends affecting backups – *growing data*, and *increasing complexity* – to see how backup appliance can and will play a critical role and what capabilities exist to address both concerns.

### THE GROWING BACKUP DATA PROBLEM

IT organizations like yours are under constant pressure to go faster, to move to new technologies - whether utilizing cloud, mobile, or big data – pushing you to adopt, in some ways, literally *anything* that will propel the business forward. Meanwhile you still have to do all the same tasks you've always had to do as a responsible IT organization – creating backups, putting up DR environments, patching, etc. And even these aspects of your job are accelerating at a very rapid pace.

One reason for this is that data and servers are growing at a much faster rate than the number of allotted IT staff. As shown in 1, according to IDG<sup>1</sup>, over the next 10 years, it's projected that servers will grow at a 10x rate, data at a 50x rate, and the number of files at an even higher 75x rate. And yet, it's still the same group of guys in IT that have to do all the work.



**FIGURE 1:** The projected increase of servers, data, and files vs the number of IT staff over the next 10 years.

<sup>1</sup> IDC Digital Universe Study (2014)

***You're going to need to focus a material portion of your time on backups, just to ensure the business is protected.***

So as a result, the demand on a backup strategy will be increasingly more complex. Many more applications (with a higher number of them being mission critical), more data to manage between active backups and archiving, and a seemingly impossible demand to get all of this recovered in a very short period of time. If you have the wrong backup solution, you'll likely be dealing with multiple tools to address various backup and recovery needs. And if you're juggling between those different tools, you're going to have to remember how to operate each to tackle issues like how backups of your virtualized environment are handled differently than how you backup your physical servers, separate from still from how you backup your legacy UNIX systems that are still kicking around.

All that that growth in servers, data and files equates to a greater consumption of IT's time and energy, keeping IT from focusing on more critical (and, frankly, more exciting) projects that will propel the business forward. And, should 1 prove to be true, you're going to need to focus a material portion of your time on backups, just to ensure the business is protected. And that's time you don't even have today, let alone tomorrow when there's more to protect.

*But, it's not just the growth of data that will put a burden on backups. With more data comes more complexity.*

## **THE GROWING COMPLEXITY PROBLEM**

1 showed a 10x growth in servers. None of us believe that we're going to have ten times the number of applications to support ten years from now. But we will have far more complex multi-tiered applications running on those servers, taking advantage of the cloud for both storage and compute services. You used to be able to protect an enterprise application by backing up a single server. But even today, the definition of a single application can include back-end database servers, front-end client services, mid-tier business processing, interfaces with credit card processing, and more. To consider an application like that protected, even today, there's a number of systems that all have to be recovered and in sync. Fast forward to ten years from now, and one can only imagine the levels of complexity we'll have to address.

The challenge with this growth is that even today, organizations like yours do not believe they can properly protect the current state of environment complexity. In a recent survey<sup>2</sup>, organizations were asked whether they could recover from their most recent disruption *within a single hour*. Eight years ago, 30% of organizations could. That number has dropped to just 2% two years ago, as shown in 2.



**FIGURE 2:** The percentage of organizations able to recover within an hour is dwindling.

This has a lot to do with two ways IT is facing a complexity problem. The first is a *growing complexity of the environment you need to protect*. Eight years ago, applications were largely single server, on-premises, and non-virtualized. If you were a Microsoft shop back then, the most complex application you had was probably Exchange – and that still only required a few servers at most. On a side note, what’s also interesting is that 30% number from back in 2007 – it’s still a tremendously awful value. It means 70% of organizations weren’t able to recover within an hour back then.

Then there’s that 2% in 2013 – it just isn’t a surprise, is it? There is just so much more going on today - the massive increase in the amounts of data, as well as the shortage of IT talent, more complexity around virtualization, more critical workloads, more application dependency – all with an overarching expectation by the organization as

2 Forrester/DRJ, The State of IT Resiliency and Preparedness (2013)

**Today, at a minimum, you need to know your data set is backed up and you should complete some basic recovery testing in your DR environment.**

a whole for IT to be able to keep this exponentially more multifaceted IT environment running.

That *expectation of continuity* is the second complexity factor IT is facing. We went from a simple expectation of being able to recover a data set (with some degree of wanting business continuity, as much as was possible) back 8 years ago, to one today where discussions around recovery times and recovery points for the aforementioned “more complex environments” are done in terms of *just minutes*.

And recovery isn't anywhere near as simple as it once was. (In some ways, you could even think of the needed recovery process as a *third* complexity factor!) Today, at a minimum, you need to know your data set is backed up and you should complete some basic recovery testing in your DR environment. But even with that being done, the most you know is that a VM will spin up. That's really only half the problem. Just because your VM spun up doesn't necessarily mean that it did so in the exact same configuration that your production environment was operating in, or in the case of multiple interdependent VMs, that they spun up in the proper order. You may have had configurations risk or you may have had changes in one side or the other that don't match precisely. As a result, when you try and spin up and use your DR environment, it's not going to work the way production did.

*And you're left scrambling to figure out what went wrong.*

To address some of these complexities, you've likely already adopted best practices like the 3-2-1 rule (3 copies of backups, across 2 mediums, with at least 1 offsite). But, given the lack of organizations being able to recover within an hour, you're still trying to figure out what backup solutions are necessary to ensure a recovery strategy that actually keeps your business operational.

You've got a backup solution on the one hand that does, technically, backup all your data, systems, and applications. Then you've got compute resources in the cloud that can act as a recovery environment in the event of a loss of location. It's bridging the two that is the tough part.

*Without a bridge, you're going to remain a part of the 98% who can't recover quickly.*

## **ENTER: THE BACKUP APPLIANCE**

The challenge here is to extend the concept of backups into the cloud in a way that both allows you to host critical data, applications, and systems in the cloud (when recovery there makes sense), but still have an ability to keep recovery on-premises when you need to. And with data needing to be backed up and the complexity of your environment increasing, you need to actually reduce complexity in the backup process or you'll never have time to do anything but backups and recoveries ever again.

It's obvious you can't do this efficiently yourself with scripts and backup jobs. You need something more intelligent to bridge this gap. That's where a backup appliance comes into play. Think of this type of appliance as part backup software, part storage, part traffic controller, and part cloud interface. With IT headcount growing only marginally in comparison to the amount of information growth over the next number of years, you're going to need some help that addresses the areas your current appliance-less backup and recovery strategy simply can't tackle.

1) **Empowered use of the cloud** – While you know you need to keep a copy of backups offsite, the question is how to get them there. By having a backup appliance that is aware of the cloud and is connected to a cloud-based storage provider, you automate the process of getting your data offsite while creating multiple copies of your data.

2) **Local protection** – You don't want to go “all in” on the cloud and only have backups there, because not only is it inefficient to have to pull terabytes of data back down from the cloud, but you also may have an issue when there is no Internet connectivity, leaving you stranded without a backup. A backup appliance retains copies of backups locally to ensure you always have a copy for recovery.

3) **Optimized WAN usage** – Backup software only works to compress and optimize the size of a backup. But because cloud-empowered backup appliances are aware of their interaction with the cloud, they include

***You need to reduce complexity in the backup process or you'll never have time to do anything but backups and recoveries ever again.***



***You need a copy of everything if you are to protect yourself from just about any kind of “disaster.”***

additional optimization to both minimize the actual transmission of data in backup and recovery operations, but also do so securely using encryption.

4) **Efficient storage** – You need a copy of everything if you are to protect yourself from just about any kind of “disaster.” And while storage in the cloud is theoretically limitless, your local appliance has its limits. So you have a problem – how do you ensure you have what you need locally, without having an entire copy of cloud-based storage on-premises? A backup appliance can intelligently keep only what you need to accomplish local recovery of critical systems and applications, and push everything else that is kept more for long term retention to the cloud.

5) **Speed of Recovery** – By having an appliance in place that knows where backups are stored and whether to pull from the cloud or from local storage – and by the very fact you have both storage locations as recovery source options – you increase the speed at which you can recover, regardless of whether a disaster is a loss of data, system, application, or location.

6) **Business Continuity** – Depending on the appliance and cloud provider used, the potential exists for the appliance to not just use the cloud for storage, but actually pass backups of virtual machines up to cloud-based recovery infrastructure automatically, providing you the ability to bring up as much as a complete environment and temporarily run operations from the cloud.

While there are many more benefits to using a backup appliance, this list represents some pretty powerful reasons to look to automatically bridge backups and the cloud. No one’s saying you absolutely can’t accomplish recovery without an appliance, but there will come a time when the growth in data and environment complexity will envelope your time and no longer afford you an ability to do anything other than backups.

## **SIMPLIFYING THE COMPLEXITIES OF BACKUPS WITH A BACKUP APPLIANCE**

As we continue to increase our reliance on progressively more complex applications utilizing more data and more servers, the demand to

**Your current strategy needs to embrace the cloud for both storage and compute.**

achieve a state of recovery quickly is only going to amplify over the next few years. Your current strategy needs to embrace the cloud for both storage and compute, but do so in a way that automates as much of the backup and recovery process as is possible. Remember, there's only going to be an extra half of you to do the work ten years from now.

By utilizing a backup appliance, now and in the future, the complexities that are quickly overwhelming IT, turn into assurances that regardless of data size or system criticality, they are being backed up, efficiently stored, and are properly available for recovery when the time comes.

*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 tradeshow around the world.*

Sponsored by

**UNITRENDS**