

Managing the information that drives the enterprise

STORAGE



The cloud to the rescue

Cloud-based disaster recovery is inexpensive and relatively easy to implement in companies of all sizes.

AUGUST 2013
VOL. 12 | NO. 6

HOW TO SET UP STORAGE FOR HYPER-V VMs

BACK UP THOSE PHONES AND TABLETS

CASTAGNA:
HYPE, HOPE OR HOO-HA? YOU BE THE JUDGE

TOIGO:
DO VENDORS REALLY KNOW WHAT PRODUCTS YOU NEED?

McCLURE:
SAFER SYNC AND SHARE

MATCHETT:
BIG DATA STORAGE MYTHS

SNAPSHOT:
STALLED TECH REFRESHES PICK UP STEAM



The day you get back, challenges will be solutions

➤ **Storage Decisions** CONFERENCES AND SEMINARS



Storage Decisions events are guaranteed to deliver practical advice and technical guidance that you can put to use as soon as you return to the office. Featuring the world's most renowned storage experts, our events offer tips and vendor-neutral advice for tackling current issues. Our innovative vendor showcases help you avoid falling behind by providing exclusive insight into market trends and emerging technologies. Plus, only you and your true peers in IT can qualify for FREE admission, ensuring an unmatched experience at an unbeatable value.

Visit www.StorageDecisions.com to register for an event near you.



HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

Stop making sense (out of storage)

We need a new rating system to sort out what's really going on in the storage world. Hey, I've got one!

GARTNER HAS ITS magic quadrants and troughs of disillusionment, but that just doesn't seem enough to properly classify everything that's going on in the data storage market these days.

Somebody has to sort these things out if we're going to make sense of all the new [technology developments](#), and how they're being bent and reshaped to fit into various vendors' narratives.

I'll do it.

I'd like to introduce the Hype, Hope or Hoo-ha rating system. This extremely unscientific and almost totally subjective system will rate the latest and greatest buzzwords being bandied about the [storage world](#) so you can figure out what to pay attention to and what you should ignore.

First, let's define our terms. Hype refers to a real, and probably very useful, technology that's getting way too much spin from marketers. In this category, you'll find solid, relatively new entries that are good matches for some shops, but they come with so much hyperbole and hard sell that you can be turned off before you've even tried them. The Hope classification is for those techs that work and you like, but they're accompanied by phone calls and pitches from vendors who hope you'll buy even more. Hoo-ha is just what it seems—bull, balderdash, baloney, bunk, bosh—and refers to imaginary products that address imaginary problems with imaginary effectiveness.

For some examples of the Hype, Hope or Hoo-ha rating system at work, I'll start with my favorite whipping boy: software-defined storage ([SDS](#)). While it lingered between Hype and Hope for a while, the term has been used

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

indiscriminately and so often that it's finally plunged into the depths of Hoo-ha. This was particularly evident when [EMC rolled out its SDS product](#) with the pitch that nobody really knows what software-defined storage means but here's our product that we call software-defined storage. It was a fairly adept attempt to hijack the term and make it its own, so if EMC wants to own that kind of Hoo-hah I guess it's welcome to it.

Contending for the Son of Hoo-ha award is the latest something-defined-something buzzword: Data-defined storage. No kidding, I saw it in a [press release](#). We now have software-defined storage and data-defined storage, but why stop there? How about definition-defined storage? It's kind of a do-it-yourself approach to storage marketing. Apparently, wherever and whatever you decide defines storage is fine, as long as it makes a good acronym.

The Hope category is more civil, mainly because the technologies do most of the talking for themselves. I'd put [solid-state storage](#) in the Hope ranks. No one disputes that it's cool stuff that runs circles around traditional disk-based storage. It's also finding its way into storage shops both large and not so large. Solid-state storage fits snugly into Hope because storage vendors are furiously trying to convince anyone who owns, plans to buy or is even just looking longingly at a storage system that if they don't include some solid-state (heck, supersize my order and make the whole thing solid-state) then why bother? The biggest factor for the Hope classification is

that vendors hope you see how different their products are from everyone else's even though they're the same. As they say, hope springs eternal.

Hype is shaping up to be the most crowded classification in my system. It's almost too easy to tag things with the hype label, but let's start with the mother of all Hype, [big data](#). The ability to analyze tons of seemingly unrelated, unstructured data is a real quest for many firms—mostly big companies with big ideas about all that untapped knowledge they hope to set free. But c'mon storage vendors, it's hard to believe that every product out there is an essential component of an infrastructure to support [big data analytics](#). But the hype is a relentless deluge of big data storage systems, big data host bus adapters, big data cables and big data racks. The sad part is that firms looking for products to support their big data efforts probably can't find them in this haze of marketing zeal.

Now that I have this rating system in place, I'm going to need some kind of graphic representation if I want it to catch on. Quadrants and troughs are already taken. Thumbs up/down? Nah, too Roman coliseum. Smiley/frowny faces? Too cute.

I guess I need to head back to the drawing board, but first I have to read this press release about a solid-state device that enables big-data-defined storage. ■

RICH CASTAGNA is editorial director of TechTarget's Storage Media Group.



HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

Losing my religion

Some storage vendors seem out of touch with the reality of what's actually going on in our data centers, and what kinds of products IT really needs.

AM CONVINCED THAT R.E.M.'s song "Losing My Religion" is an earworm—one of those tunes that, once you hear it, you can't get it out of your head. I found myself humming it recently as I listened to a data storage vendor executive deliver a pitch to a room of resellers. The tune is now stuck on repeat in my mental MP3 player.

The event was a channel partner conference hosted by an outstanding distributor. I spoke there early on, but the next day featured a manager with channel sales responsibilities from a vendor sponsor. He seemed to be churning a lot of vernacular in his attempts to (1) suggest that he had spilled a lot of blood in the same mud as the attendees, and (2) do what he could to energize a room of resellers regarding the struggling company's product list.

He told stories of how vulnerable competitors were

to his company's "[full technology stack](#)," which included switches, servers and storage. His firm, he said, was pushing to replace gear in customer shops not only bearing the pesky branding of competitors, but also older gear sporting his own company's name. He explained that the data storage vendor was spending too much money [maintaining older systems](#) that could be better spent on innovating features on newer gear. He warned attendees that if the replacement of his older rigs wasn't pursued in earnest, he'd use the "nuclear option"—making these "opportunities" available to anyone who wanted the business and the revenue. That was a big stick, but he offered carrots as well, including extra points on replacement deals, to further [incentivize his channel folk](#).

The products the storage vendor now wants to sell are based on intellectual property from smaller storage

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

companies acquired in the past few years. The vendor wants to combine these rigs with home-team-branded servers and switches to build homogeneous, single-vendor infrastructures at customer sites. Listening to the speaker's arguments, and knowing what I do after visiting many IT shops, I wasn't at all convinced that [infrastructure rip-and-replace](#) would sell in the current market. As a courtesy to my hosts, however, I kept those concerns to myself.

But the speaker must have anticipated that kind of reaction because he made a point of suggesting that the biggest obstacle to success were the technical folks at customer accounts whom he called "non-believers."

"They want to know how, exactly, you're going to deliver the value that you're promising with your array or server or switch," he said, looking genuinely pained.

I realized I was experiencing a sensation that's rather unusual for me: *schadenfreude*. I was actually enjoying the fellow's angst regarding those annoying IT managers who refuse to simply take data storage vendors at their word and actually want proof.

It dawned on me that I'm one of them: losing my religion as I listen to vendors trying to promote their one shiny new rig as a panacea for everything that ails storage. At that point, R.E.M.'s track loaded in my brain and began what has become a continuous playback loop.

I was eager to learn what advice the fellow would offer to help the audience work around those data center

disbelievers. He posited the familiar nonsense from analysts depicting the world as two diverging trends: [increasing amounts of data to store](#) while budgets stay fixed or fall. To fill the void, he offered, you needed technology like his feature-encrusted array that delivers "effortless tiering," "thin provisioning" and other value-add soft-

The speaker must have anticipated that kind of reaction because he made a point of suggesting that the biggest obstacle to success were the technical folks at customer accounts whom he called "non-believers."

ware goodies, and will soon include a tier-zero flash layer (somewhat late to the party with that one, I thought). Such technology would dramatically reduce labor costs (e.g., the need to maintain trained storage geeks on staff), he asserted, although he didn't provide any real Opex or Capex numbers.

Missing from his pitch was any mention of tape. Despite the increasing capacities of tape and the growing popularity of the [Linear Tape File System](#) as a low-cost repository for seldom-accessed files—a solution that would clearly address the aforementioned gap between

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

data growth and budget—nary a word about this tripped from the channel sales guy’s lips.

R.E.M. droned louder in my head. I looked around the room to see a few heads nodding and smiling. Believers, I thought to myself. My question was whether any of the folks were what I would call “solution integrators.”

In the late 1990s, I did a lot of work with solution integrators. These were companies that maintained technically competent and financially astute consulting teams who would seek to understand client requirements and constraints; sift through all the marketecture to identify kit that would most effectively meet current requirements while providing a cost-efficient pathway for growth; and deliver a strategy to the customer that represented the best fit of technology products to needs.

Those guys, true solution integrators, were worth their weight in gold. But they virtually disappeared following the market woes that hit the storage industry in 2001.

For those who weren’t paying attention to market behavior in the post-dot-com era, the response of [many data storage vendors](#) when the bottom fell out of the supposedly unshakeable storage market was to take the most lucrative accounts away from their channel partners and bring them inside to be serviced by direct sales personnel. That made it impossible for resellers to afford to keep qualified and knowledgeable consulting folks on staff. Instead of being solution integrators with

well-defined methodologies and pure vendor agnosticism, many resellers simply became order-takers for three letter companies.

Interestingly, one fellow at this event did raise his hand to ask whether selling hardware with value-add software was the right solution to a problem that seemed to

Those guys, true solution integrators, were worth their weight in gold. But they virtually disappeared following the market woes that hit the storage industry in 2001.

come from data mismanagement. I was floored, and confess that I missed the exact wording of the gobbledygook answer provided by the speaker. We non-believers are like that.

However, reflecting on the query and on what it says about the possibility that true storage integrators may still be out there, I flash on the R.E.M. lyric: “I think I thought I saw you try.” ■

JON WILLIAM TOIGO is a 30-year IT veteran, CEO and managing principal of Toigo Partners International, and chairman of the Data Management Institute.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

SIX KEY STEPS TO IMPLEMENTING CLOUD-BASED DISASTER RECOVERY

Cloud DR puts effective disaster recovery within reach of just about any company. Here's how to get it done.



THE CLOUD MIGHT be the most overhyped technology in decades, but it can be extremely beneficial when it's used as part of an organization's disaster recovery (DR) plan. It's now possible for a company to create a [cloud-based recovery site](#) that can be used if the primary data center is incapacitated. There are a number of alternative strategies for using the [cloud for disaster recovery and disaster prevention](#); we'll describe the key considerations in detail.

1

EVALUATE YOUR DATA PROTECTION NEEDS

The first step in implementing any DR offering is to evaluate your organization's needs. While this might sound simple (and may be), the results of your evaluation will be the major factors that determine the infrastructure and configuration you'll have to put in place to facilitate [cloud-based data protection](#). For

By Brien M. Posey

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

example, some organizations use cloud storage as part of a [disk-to-disk-to-cloud backup solution](#). The primary backups remain onsite, but they're replicated to cloud storage where they'll be protected from things that could disable a data center, such as a fire or flood. Other firms replicate entire virtual machines (VMs) to the cloud so they can be spun up and hosted in the cloud if it isn't possible to continue to host the VMs in the local data center.

2

CHOOSE YOUR CLOUD PROVIDER

Once you've determined your data protection needs, the next step is to [identify cloud providers](#) that can accommodate those needs.

Not every cloud provider is equipped to handle every situation. For example, some providers will allow you to replicate a VM, but won't host it. Similarly, there are providers that offer storage, but little else. If your goal is to build a cloud-based disaster recovery site, you'll need to find a cloud service provider that offers the specific capabilities you need.

On the other hand, if your only goal is to replicate data to the cloud, you might be better off subscribing to a storage-only plan to avoid paying for services and capabilities that you don't need.

Whatever your needs are, it's a good idea to identify several different cloud service providers so you can compare costs and [levels of service reliability](#).

3

ESTIMATE THE COSTS

Once a viable cloud provider has been identified, you'll need to estimate the [cost of using cloud-based disaster recovery](#). Each cloud service provider typically has a unique pricing model, but the total monthly cost usually comprises some combination of the following factors:

- A monthly subscription fee
- The amount of Internet bandwidth used
- The amount of storage space consumed
- The number of VMs (or virtual processors)

Some providers treat the subscription fee as pro-rated payment for service. For instance, the monthly subscription fee might include a specific amount of bandwidth usage, with any usage beyond that level resulting in additional costs.

It's also a good idea to check the provider's policy toward VMs that aren't powered on. Some providers charge based on the number and type of VMs created, regardless of whether those virtual machines are powered on or not. Other providers charge only for actual usage and therefore offer a billing structure that's based on the number of minutes or hours for which a VM is powered on.

The way a cloud provider bills for services can have a major impact on the [cost of your cloud-based disaster recovery initiative](#). That's why it's so important to develop

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

accurate cost estimates before signing on with a provider. Although some providers use very complicated billing formulas, there are tools available to help you estimate the costs. For example, Microsoft Corp. offers a [cost estimate tool for Windows Azure](#). Similarly, Amazon Web Services offers a [cost calculator](#). Also, some third-party tools have built-in cost estimators. Veeam Backup and Replication Cloud Edition, for example, has its own [cost-estimating calculator](#) that works with a number of different cloud providers.



DEVELOP A BANDWIDTH MANAGEMENT STRATEGY

Another important step in taking your disaster recovery initiatives to the cloud is to come up with a strategy for [managing Internet bandwidth](#). Bandwidth management is extremely important for a number of reasons, including:

- Many cloud service providers charge for bandwidth consumption.
- Your own Internet service provider may impose monthly usage caps or may charge for excess bandwidth usage.
- You must provide adequate bandwidth to allow data to be backed up (or replicated) in a timely manner.

- You must ensure that your [cloud backups](#) or replicated data don't consume so much bandwidth that other Internet usage suffers from inadequate bandwidth availability.

The method you'll use for bandwidth management depends on the approach you're using for [copying data to the cloud](#). Some backup software and many cloud storage gateway appliances include built-in bandwidth scheduling features. Such features generally let you limit the overall bandwidth that the data-copying app consumes and may let you increase that limit during off-peak hours.

Similarly, many organizations use quality of service (QoS) to reserve bandwidth for cloud backups and other bandwidth-intensive services. This ensures that each Internet-based service receives the bandwidth it needs without consuming an excessive amount of the available bandwidth.

Regardless of the amount of bandwidth you reserve for your [cloud-based backup or replication service](#), it's important to make efficient use of that bandwidth by also applying data deduplication.



DETERMINE THE LOGISTICAL REQUIREMENTS

If your company is using the cloud solely for its storage capabilities, there will only be a

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

minimal amount of logistical planning that will have to take place. However, organizations that wish to perform full-blown failovers to the cloud will need to take a number of considerations into account.

The actual logistics that must be planned can vary considerably depending upon your company's existing

infrastructure, the cloud service being used and the desired end result. Even so, there are some aspects of the logistical planning process that are especially common.

The first consideration relates to [how you will copy data](#) from your on-premises data center resources to the cloud service. If you're using a public cloud, the

Disaster recovery: A new app for the cloud

CLOUD-BASED DISASTER RECOVERY (DR) usually means replicating data or even entire virtual machines to the cloud. However, for companies that already have a secondary data center for DR in place, it might make more sense to use cloud services as a mechanism for facilitating the DR process rather than using the cloud as a data repository. Microsoft Corp., for example, has introduced the Hyper-V Recovery Manager on Windows Azure.

The Hyper-V Recovery Manager is a hybrid service that allows you to use Windows Azure to manage the replication process between your primary and secondary data centers.

Hyper-V Recovery Manager is intended to replace storage vendors' proprietary SAN-to-SAN replication. Instead, replication is performed at the hypervisor level using the native Hyper-V 3.0 replica feature. Although the virtual machine replication process occurs between data centers, Windows Azure is used as a cloud-based solution for managing the replication, failover and DR testing process. Administrators are able to define a series of recovery plans within the Windows Azure interface, and use the interface to perform failovers and failover testing. Azure communicates with on-premises System Center Virtual Machine Manager (VMM) server deployments (one in each data center). VMM, in turn, performs the heavy lifting by instructing individual host servers to replicate data, perform a failover and so on. This can be an especially effective solution for organizations that want to take some of the cost and complexity out of failovers to an alternate data center. ■

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

replication process will most likely be software based. In any case, you'll have to use a replication mechanism that's supported by the cloud provider and also compatible with the on-premises resources you plan to replicate.

Another important consideration is [Active Directory \(AD\) synchronization](#). Windows-based clustering solutions require cluster nodes to be members of a common Active Directory domain. This same concept also holds true for many other Microsoft fault-tolerant offerings (such as Exchange Server database availability groups) because the technologies are built on top of failover clustering components. That means that if you want to extend a cluster to the cloud for DR purposes, you will most likely need to extend an on-premises AD domain to the cloud.

The actual method used to accomplish this can vary from one cloud provider to the next. In the case of Windows Azure, Microsoft provides a [cloud-based directory service](#)—called the Windows Azure Active Directory—that can be synchronized to an on-premises Active Directory and an on-premises DNS server. The biggest trick to making the synchronization work is that the on-premises network and the virtual network that exists within the Windows Azure cloud must be able to communicate with one another. The easiest way to facilitate this communication is by deploying an on-premises Routing and Remote Access Server (or a VPN server) and configuring Windows Azure to attach to your

on-premises network through the VPN.

Another big issue you'll have to address is that of the cluster maintaining quorum. Windows failover clusters are based on the Majority Node Set model. This means that for a cluster to maintain quorum, half plus one of the cluster nodes must remain online. For example, if a clus-

You'll have to use a replication mechanism that's supported by the cloud provider and also compatible with the on-premises resources you plan to replicate.

ter has seven nodes, then four nodes must remain functional for the cluster to continue to function.

The problem with this concept is that Windows can't tell the difference between a multinode failure and a WAN link failure. Placing the majority of the nodes on the premises would prevent a cloud failover from occurring in the event of a WAN link failure, but it wouldn't protect you against a data center-level failure because an insufficient number of nodes would exist in the cloud. Similarly, if the bulk of the nodes existed in the cloud, then a WAN link failure would trigger a failover to the cloud. One common solution to this problem is to place

HOME

HYPER-HYPE OR HOOP-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

an equal number of nodes on the premises and in the cloud and to then place a “tie breaker” node in a third location. Doing so ensures that one of the two sites will always be able to maintain quorum, and reduces the chances of link-related failovers. It’s important to make sure that all three locations can communicate with one another at a speed that adheres to Microsoft’s latency requirements for clustering.



VIRTUAL MACHINE REPLICATION

Clustering isn’t appropriate for every situation, and not every application or virtual server can be clustered. An alternative technique for using the cloud for disaster recovery is to simply [replicate VMs to the cloud](#). If an organization uses this approach, they must determine what they hope to gain from the replication process. For example, VM-level replication can provide the following benefits:

- Point-in-time-image-based recovery
- The ability to mount a cloud-based copy of a VM and extract data

- The ability to redirect users to a cloud-based VM replica in the event of an on-premises failure

If the goal is to redirect users to a cloud-based VM in the event of a failure, then the biggest challenges you’ll face are related to IP address injection and DNS record modifications. To be usable, the VMs will need IP addresses that are local to the cloud-based virtual network subnet on which they will reside during a failover. DNS record modifications are required so the virtual server can be found when it’s running in the cloud.

The major hypervisor vendors offer features for performing these tasks and redirecting the user workload, but some third-party backup vendors offer similar capabilities that can be used without admins having to configure IP address injections and DNS modifications. ■

BRIEN POSEY is a Microsoft MVP with two decades of IT experience. Before becoming a freelance technical writer, Brien worked as a CIO for a national chain of hospitals and health care facilities. He has also served as a network administrator for some of the nation’s largest insurance companies and for the Department of Defense at Fort Knox.

[HOME](#)[HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE](#)[DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?](#)[THE CLOUD TO
THE RESCUE](#)[HOW TO SET UP STORAGE
FOR HYPER-V VMs](#)[BACK UP THOSE PHONES
AND TABLETS](#)[SAFER SYNC AND SHARE](#)[BIG DATA STORAGE MYTHS](#)[STALLED TECH
REFRESHES PICK UP
STEAM](#)

CONFIGURING STORAGE FOR HYPER-V

New storage features in Windows Server and Hyper-V make storage configuration and maintenance easier than ever for Microsoft-based virtual environments.

MICROSOFT HYPER-V SERVER offers numerous, and powerful, storage features that provide flexibility, availability and (potentially) a reduced total cost of ownership for data centers. Hyper-V Server is available in a number of different editions, both paid and free. You'll find [Hyper-V 3](#) in the commercial version of Windows Server 2012 and in the free Hyper-V Server. No matter which edition you use—with the exception of Client Hyper-V in Windows 8—all the storage features discussed below are available now.

DATA CENTER SOFTWARE IS VORACIOUS

In Mark Andreessen's famous article "Why Software Is Eating the World" (*The Wall Street Journal*, August 2011), he discussed how software companies' products are disrupting traditional industries around the world. Hyper-V Server is one of those pieces of software with the ability

By David Davis

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

to replace what has previously been provided by numerous dedicated pieces of custom hardware and software and, thus, “eat the datacenter.” That means [Windows Server 2012 Hyper-V](#) can potentially replace numerous servers (through consolidation), network devices (with virtual switches and firewalls), and storage resources such as SANs and network-attached storage (NAS) arrays leveraging some of the storage features described later in this article.

The storage features in the latest version of Hyper-V

Server, and those promised for the upcoming version of Hyper-V Server, can supercharge the storage infrastructure of your data center.

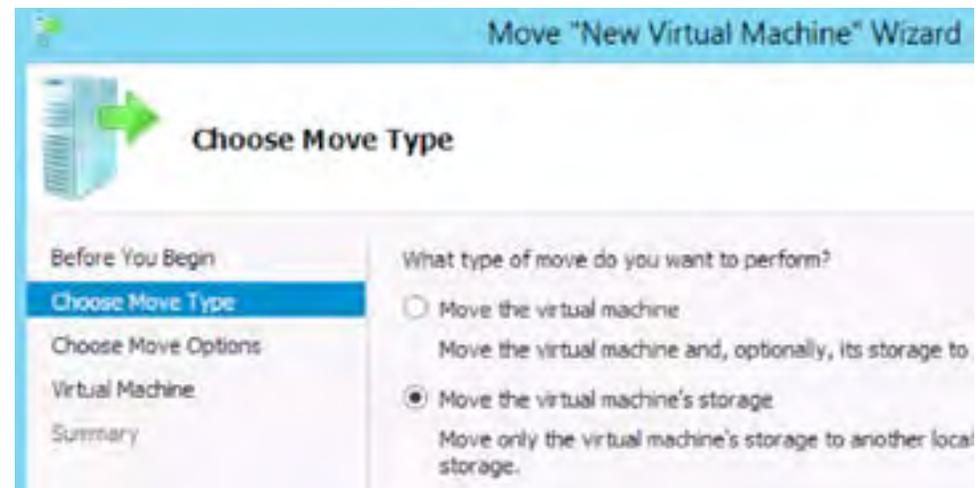
LIVE STORAGE MIGRATION

The ability to move the storage data of a running physical server from one SAN to another without downtime is an advanced SAN feature. Hyper-V Server (even the free edition) includes this capability for virtual machines

FIGURE 1:

Hyper-V Live Storage Migration

The Hyper-V Live Storage Migration wizard steps users through moving a running virtual machine from local to shared storage, from shared storage to other shared storage, or from a shared storage resource to local storage.



HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

(VMs). With [Hyper-V Live Storage Migration](#), you can follow a wizard's instructions to move a running VM from local storage to shared storage, from one shared storage to another, or from shared storage to local storage.

As with any technology that moves a large amount of data across a network, the time that the storage migration takes and the number of migrations you can do at once depends on the size of the data and the available network bandwidth. For example, Hyper-V Servers on a 10 Gb Ethernet network will be able to move many more VMs at a time, and faster too.

SHARED NOTHING LIVE MIGRATION

Initially a feature that only Hyper-V offered, shared nothing live migration lets you move a running virtual machine's memory, CPU cycles and configuration/registration from one Hyper-V Server to another with no downtime. With live migration, the VM's disk file stays in the same place. With previous editions (and competing hypervisors' previous editions), the VM was required to be on shared storage (a SAN or NAS). With [Hyper-V 3.1 shared nothing live migration](#) allows you to move a running VM from server to server, even if the VM is using local storage. Effectively, the VM is migrated with both a live migration and live storage migration at the same time. For small virtual infrastructures, this negates the need for costly shared storage.

Top seven storage enhancements to Hyper-V Server

- 1. Live Storage Migration** moves running virtual machines' (VMs) virtual disk from one location to another.
- 2. Shared nothing live migration** lets you move running VMs even without shared storage.
- 3. Data deduplication reduces storage capacity requirements**, especially for VMs.
- 4. SMB 3.0 storage for VMs** stores Hyper-V VMs on Windows Servers, and uses advanced features without shared storage.
- 5. Hyper-V Replica** offers cost-effective data protection.
- 6. Storage Spaces** is a new Windows Server virtual SAN offering to pool storage resources.
- 7. Offloaded Data Transfer** can speed common VM storage operations by offloading virtual machine storage operations to the SAN.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

DATA DEDUPLICATION

[Per-volume data deduplication](#) is built into Windows Server 2012 (not the free version of Hyper-V). According to statistics from Microsoft Corp., by enabling dedupe, you can achieve an 80% reduction in storage capacity utilization of VM libraries. Windows Server deduplication isn't just for Hyper-V VMs. It can be used on any data stored on NTFS volumes (it doesn't deduplicate boot drives, live VMs or live SQL databases), and offers greater than 20% storage capacity savings on user home directories and 50% savings on general file shares. This allows

you to not only save on the number of disks and arrays you have to buy, but the amount of data you have to back up.

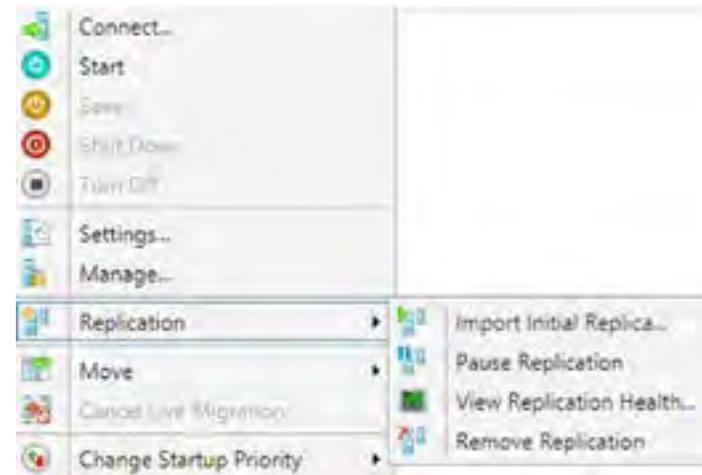
VIRTUAL MACHINE STORAGE USING SMB 3.0

With Windows Server 2012, Microsoft supports using [Windows Server SMB 3.0 file shares for storing your VMs](#). In other words, you can get by without costly shared storage and still use the advanced features. The Windows 2012 Server that stores the VMs could be anything from a single host to a cluster of Windows Server

FIGURE 2:

View Hyper-V Replica Status

With Hyper-V Replica, it's easy to configure replication on a single virtual machine (VM) or several VMs with just a few point-and-click steps.



HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

2012 systems in a failover cluster that provides highly available shared storage.

HYPER-V REPLICA

One advanced feature offered with many storage arrays is LUN replication. While it's a nice (but usually expensive) feature, most Hyper-V administrators want more granularity. The ability to replicate virtual machines at the VM level in Hyper-V 3 (including the free edition), allowing admins to select the VM they want to replicate, is therefore a welcome feature. Shared storage isn't required (even local storage can be used) and no additional replication licenses have to be purchased. An ideal feature for small and medium-sized businesses that can't afford expensive hardware-based replication for disaster recovery (DR), today's [Hyper-V Replica may grow](#) into an enterprise-grade DR tool in time. Best of all, configuring replication of a VM (or group of VMs) is done in a few point-and-click steps.

WINDOWS SERVER 2012 STORAGE SPACES

To take the idea of software-defined storage for the virtual infrastructure (or the entire data center) even further than just SMB 3.0 file sharing, Windows Server 2012 includes [Storage Spaces](#). Storage Spaces is a storage abstraction layer that runs atop the NTFS file system and uses

the SMB 3.0 protocol to create a pool of storage out of local disk, shared storage or local flash drives in a server. In other words, you can think of Storage Spaces as a virtual SAN that pools different types of storage across various servers into a single unified storage resource. It may or may not be used for Hyper-V. Storage Spaces is a data center-wide storage offering that Microsoft continues to improve and, at some point, may begin to offset the number of dedicated SAN arrays that are purchased.

WINDOWS SERVER 2012 HYPER-V OFFLOADED DATA TRANSFER

Virtualization is all about efficiency, but some things still aren't done efficiently with software within a server (like moving massive amounts of data). It's best to work with data where the data lives. For those who have SANs in their data centers, Hyper-V can use some help in manipulating VMs; otherwise, the VM data has to be moved to the server, manipulated and then put back on the SAN. [Offloaded Data Transfer \(ODX\)](#) is included with Hyper-V, but you can only use it if your storage array supports it. For example, it has been reported by some media that with an ODX-capable array, a new VM deployment took only two minutes using a 30 GB template. However, with a traditional array and a 1 GB Ethernet connection, the deployment took 40 minutes, and used 99% of the network bandwidth and 1% to 5% of the host CPU.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

NEW FEATURES IN WINDOWS SERVER 2012 R2 HYPER-V

While it may still seem new to many in the IT world, Hyper-V 3 (included with Windows Server 2012 and released in September 2012) is about to be [upgraded with Windows Server 2012 R2](#). Windows Server 2012 R2, announced at the Microsoft TechEd 2013 North America conference, offers numerous new features; among the storage-related features are:

- **Tiered Spaces**, which adds tiering intelligence to Storage Spaces. Tiered Spaces knows to keep the most frequently accessed data on solid-state storage (or whatever the fastest drives in a pool are) and to move less frequently used data to the slowest spinning disks.
- **Live migration with compression** brings compression to the live migration process, [enabling Hyper-V live migrations](#) to occur in roughly half the time. If you add RDMA controllers to a Hyper-V Server, you can reduce live migration times (comparing Windows Server 2012 to the R2 version) by as much as 70%.
- **Hyper-V Recovery Manager** is Microsoft's follow-on to Hyper-V Replica, which has been one of the most popular features of Windows Server 2012 Hyper-V. Because of the interest in Replica, Microsoft says it will add a new centralized Hyper-V Recovery Manager application with Hyper-V R2 that allows managing replication in a single place, across multiple sites.

WHAT THIS MEANS FOR STORAGE MANAGERS

Storage managers should be aware that the software takeover of all the pillars of the data center is coming thanks to faster servers and software advancements. Features included with VMware vSphere and Microsoft Hyper-V will soon be able to provide capabilities similar to what modern storage arrays provide for the virtual infrastructure. Thus, if you can get a 100% virtualized infrastructure (or close to it), these up-and-coming software-based storage offerings may be all that you need. ■

DAVID DAVIS is the author of the VMware vSphere and Microsoft Hyper-V video training library from www.TrainSignal.com. He has written extensively, is a vExpert, VCP, VCAP-DCA and CCIE #9369 with more than 20 years of enterprise IT experience.

[HOME](#)[HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE](#)[DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?](#)[THE CLOUD TO
THE RESCUE](#)[HOW TO SET UP STORAGE
FOR HYPER-V VMs](#)[BACK UP THOSE PHONES
AND TABLETS](#)[SAFER SYNC AND SHARE](#)[BIG DATA STORAGE MYTHS](#)[STALLED TECH
REFRESHES PICK UP
STEAM](#)

THE STATE OF MOBILE DEVICE BACKUP

With the growing bring-your-own-device movement, backing up data on mobile devices is quickly becoming an issue for most companies.

THE [bring-your-own-device \(BYOD\) movement](#) means that in most companies a lot of employees are using their own laptops, smartphones and tablets to do company work when they're on the road or even when they're in the office. People use these devices, particularly tablets, for content creation and content consumption. That means they're not just reading documents, but making notes, editing files and otherwise working with data that must be protected.

Employee-owned laptops are being brought into companies, but phones (and increasingly tablets) may represent the [bigger data protection challenge](#) for IT because they're more easily lost or stolen, and incorporating them into a traditional backup plan isn't as simple as it may be for other endpoint devices. There are fewer iOS and Android device options available from traditional backup application vendors and non-traditional alternatives, such as cloud services for backup, sync and file sharing, may

By Eric Slack

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

not meet IT requirements for protecting and controlling corporate data.

THE BYOD BATTLE

One alternative would be to provide company-purchased smartphones to employees and prohibit the use of personal devices for company business. Aside from the cost of this option, many companies see the use of mobile devices as a way to enable people to do their work whenever and wherever they want, which can promote collaboration and better productivity. It's harder to attract good employees if your [company is perceived as too restrictive](#), and requiring a separate cellphone for company calls could fall into that category.

If there's no getting around the [use of personal devices to handle company data](#), then IT will have to assume the responsibility of adding smartphones and tablets to their list of endpoint devices that must be backed up.

THE RISKS OF MOBILITY

The first consideration around [data protection for portable devices](#) is the potential loss of corporate data if the device is lost or stolen. Regular, automatic backups can address this, but a bigger risk may be from potential intrusion to the larger corporate infrastructure if one of these mobile devices is compromised. Part of the

reason for this is that users often neglect to set passwords for their phones and tablets, and when they do use a password it's often the simple four-digit variety that's relatively easy to crack.

This may be considered more of a security issue, but while we're focusing on backup there are other services corporate IT should take into consideration to address security concerns. It's important to understand what some of these security functions are since they may be part of a comprehensive data protection strategy that includes backup. [Data loss prevention \(DLP\)](#), [mobile device management \(MDM\)](#) and [remote wiping](#) are security offerings that are being deployed in conjunction with backup to protect corporate assets on employee-owned smartphones and tablets.

- **DLP** is a comprehensive software solution that discovers, monitors and protects confidential data, typically on [endpoint devices](#). Much more than just protecting data, it monitors devices while in use and can block transmission into and out of them when it discovers a potential breach. DLP products typically route all Internet traffic through a corporate DLP server that controls what gets into and out of the endpoint devices.
- **MDM solutions** are software products that help to monitor and manage mobile devices to ensure they're

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

being protected with updated backup and security applications. In that regard, backup is just one of the functions these products can manage. MDM software also reduces the overhead on IT administration associated with deploying and updating applications on mobile devices.

- **Remote wiping** allows IT to delete some or all of the data, profiles and personal settings stored on a mobile phone or tablet that has been lost or stolen. This type of protection is limited, however, because the device must be connected to the Internet for the remote wipe command to be executed. Many data protection and security applications designed to run on mobile devices include a remote wiping function. In addition, Apple has a [Find My iPhone service](#) that can wipe settings and personal information, and Android tablets and smartphones can be similarly wiped using Google Sync.

There are a number of [products that provide mobile device backup](#), three of which are described below. It's surprising to see how many backup programs don't back up data on tablets and smartphones. Of all the traditional enterprise backup vendors, CommVault Systems Inc. is the only one that can back up mobile devices.

It's interesting to note that although many vendors say their products "support" mobile devices, they don't

actually back up data from that device. They only allow backed up data to be viewed on or restored to a mobile device, but that data was previously backed up from servers or desktops in a traditional backup infrastructure.

- **Asigra Inc.** Asigra is the software infrastructure application used by many cloud backup providers. Asigra's DS-Tablet Client and DS-Smartphone Client install on the appropriate mobile devices and collect data

Many vendors say their products "support" mobile devices, but they don't actually back up data from that device—they only allow backed up data to be viewed on or restored to a mobile device.

for transmission to the cloud provider's network at specific scheduled times. This automated process includes block-level deduplication and compression to reduce WAN bandwidth. Users can also do manual, one-click backups or initiate a restore themselves without involvement by the company or their cloud provider. iOS clients are downloaded from the Apple App Store, while Android clients can be downloaded from the Android Market or Amazon Appstore.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

- **CommVault.** Simpna Edge Data Protection is an optional module for CommVault's enterprise backup application that supports direct backup of Android, iOS and BlackBerry devices. The product has built-in, policy-driven capabilities for client discovery and automated backup agent deployment. Simpna provides HTTPS protection, built-in Secure Sockets Layer and data encryption to help secure data transfers without a virtual private network (VPN). Users can be authorized to control backups and restores for their own devices, as well as other backup clients in the environment. Simpna Edge Data Protection uses global, source-side deduplication and can leverage "opportunistic" scheduling rules to minimize the impact of backup activities on users.
- **Druva.** Druva's inSync product is designed for endpoint data protection for laptops, desktops, smartphones and tablets. It provides full support for iOS and Android devices, and partial support for Windows mobile devices. It's the most complete and feature-rich product of the three described here. Companies can deploy inSync in their data centers or use Druva's cloud service to host the application. inSync leverages global deduplication and bandwidth throttling with auto-resume to optimize the backup window. Users can self-deploy a mobile app using an ID and a password, allowing them to add devices without IT

involvement. Druva offers a DLP module add-on with on-device data encryption, remote wiping and geo location of lost devices.

MOBILE DEVICE BACKUP NOT A STANDARD OPTION (YET)

As mentioned earlier, it's somewhat surprising to find how few backup applications actually support iOS and Android devices. There are a couple of reasons for this. First, larger, traditional backup companies may not be thinking such support is essential, since most of these devices have some sort of embedded backup/sync service, like Apple Inc.'s iCloud. While these are largely consumer services that put data outside of corporate IT control, they may still be seen as a free alternative to a backup application's mobile device agent.

Another reason is that backup isn't the foundational data protection tool for mobile devices that it is for other endpoint systems, such as desktops and laptops connected to the corporate LAN or VPN. The [risks associated with mobile devices](#) are more than just the potential loss of the files a person had on their particular tablet or smartphone. Security is at least equally important, so companies are looking at DLP and MDM offerings because backup isn't all they need.

There are other ways to protect files that users access, create and modify on their mobile devices. Cloud file

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE

DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?

THE CLOUD TO
THE RESCUE

HOW TO SET UP STORAGE
FOR HYPER-V VMs

BACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

sync-and-share services are becoming popular choices as data repositories for remote access; they provide IT with a way to protect these company data assets as long as users keep their files uploaded and there are sufficient security measures in place. The bring-your-own-cloud problem is becoming the next “BYO” issue for IT, but there are many available options, including having the company stand up its own private cloud.

THE MOBILE BACKUP BOTTOM LINE

While few traditional backup vendors currently support iOS and Android devices, there are backup products that can run in the corporate data center and/or in the cloud to provide most of the traditional backup protection that

IT is looking for. Larger companies that simply want to add smartphones and tablets to their list of endpoints backed up with their existing enterprise software may be surprised at the slim pickings. But for companies willing to use a dedicated backup product for their mobile device backup, there are several to choose from, especially among smaller vendors.

However, IT may want to address issues beyond backup by taking a broader look at ameliorating the risks brought on by employees using their own mobile devices at work to do company business. Security in the form of DLP and MDM solutions should be part of the mix for many companies. ■

ERIC SLACK is a senior analyst at Storage Switzerland.



Security concerns drive online file-sharing innovation

As the online file-sharing market grows, so do concerns about security and the opportunity for users and hackers alike to damage the data and reputation of companies.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE

DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?

THE CLOUD TO
THE RESCUE

HOW TO SET UP STORAGE
FOR HYPER-V VMs

BACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

HAVE WRITTEN QUITE a bit in recent months about the increasing adoption of (and challenges associated with) online file-sharing (OFS) and collaboration solutions. Because of its prevalence in the consumer market, the most well-known product on the market today is Dropbox. But over the past two years, many offerings targeted at businesses have arrived on the market. Those targeted at business use cases differ from consumer products in that they have a central administration console to provision, manage and monitor individual user accounts, as well as centralized billing and often a shared storage quota. At the same time, traditional consumer-focused vendors, like Dropbox and SugarSync, now have [offerings built for business use](#). However, enterprises are still hesitant to adopt these solutions, primarily because of security concerns.

SECURITY EMERGES AS TOP CONCERN

The increasing use of mobile devices by employees is a key driver behind the [adoption of online file-sharing and collaboration solutions](#) at the corporate level. Corporate use of OFS solutions is growing rapidly. According to Enterprise Strategy Group (ESG) research, 28% of organizations have established a corporate OFS account and 61% expect to do so within two years.

For organizations that haven't adopted online file-sharing solutions or have no plans to adopt one, "security concerns" is the most commonly cited reason behind this lack of interest. These security issues include data leakage, Web-based threats and application-layer vulnerabilities. Organizations also grapple with ongoing questions about data ownership and regulatory [compliance in an online file-sharing environment](#).

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

BIGGEST THREATS: EMPLOYEE BEHAVIOR AND ATTACKS ON PROVIDERS

A variety of security challenges are causing headaches for current users of online file-sharing solutions, and deterring other organizations from deploying online file sharing. The primary concern is simply that the online file-sharing service provider itself will be attacked, potentially leaving customer data vulnerable to theft. An example of this vulnerability was the [breach of Dropbox last summer](#). Fortunately, that incident was not disastrous, resulting mostly in spam and inconvenience for Dropbox customers. Indeed, the spam issue seems to be continuing into this year. Nevertheless, this was a wake-up call for current and planned OFS customers. Online file-sharing platforms could well be the target of sophisticated attacks in the future, based on the volume and value of the customer data passing through their gates.

The aggregate amount of business data stored in these solutions presents an extremely rich target for hackers. To protect against such attacks, most online file-sharing providers ensure data is encrypted and that the encryption keys (if the service provider is holding keys) are stored at a different location than the data so that attackers would, at the very least, have to break into two data centers to gain access to usable data. And most (including Dropbox) have [introduced optional two-factor authentication](#). Moreover, attacks designed to steal an administrator's password could be achieved as easily within a

corporate data center as they could at a service provider. For IT professionals, it's important to understand how and where encryption keys are held and protected, as well as the implications of an administrator password theft for any online file-sharing and collaboration provider under consideration. You should also ask about and understand which employees in the provider's organization have access to their passwords and data.

STAYING COMPLIANT, PLANNING FOR ACCIDENTS

Organizations that have adopted online file sharing are still concerned about data leakage, whether accidental or intentional, by their own employees. Whether subscribing to a service or hosting an OFS solution internally, the risk of having an internal employee either accidentally or willfully causing a data breach is one to be taken seriously, though not likely to be elevated simply through the use of a cloud service provider. The proliferation of corporate [information on a greater number of employee-owned devices](#) could certainly increase risk, and is a good reason to deploy a corporate OFS solution to monitor file usage patterns and detect anomalies, as well as wipe corporate data if a device is lost or stolen.

Organizations also worry that they'll have a tough time remaining compliant with industry regulations as a result of their online file-sharing usage. Service providers can't be fully compliant on their own—the Health

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGEDO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?THE CLOUD TO
THE RESCUEHOW TO SET UP STORAGE
FOR HYPER-V VMsBACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

Insurance Portability and Accountability Act (HIPAA), Federal Information Security Management Act (FISMA) and other regulations are a shared responsibility between the service provider and the subscriber. More and more service providers are providing tools and controls to enable compliance across many industries. Service provider Box recently announced a [comprehensive HIPAA/HI-TECH-enabled solution](#), and most well-known vendors, such as [AirWatch](#), Citrix (ShareFile), Egnyte, [EMC Syncplicity](#), Intralinks, WatchDox and Workshare, claim to support HIPAA; some even support the Federal Risk and Authorization Management Program, FISMA and PCI compliance.

Of course, there's always the option to host data in-house. There's a different cost model, of course, but the tradeoff is peace of mind. There are options to leverage the benefits of a Software-as-a-Service (SaaS) model with in-house storage—vendors such as [Citrix ShareFile](#), EMC Syncplicity, [Egnyte](#) and Signiant offer the ability to store data on site while the software platform runs in the cloud. There's also a fully private option in which the software and data are installed on the premises. [Accellion](#), Acronis, Airwatch, OpenText and [WatchDox](#), among others, all support on-premises installation.

WHAT THIS MEANS FOR YOU

OFS is becoming an increasingly important tool in the

IT team's toolkit, helping organizations to reduce storage and administration costs and [improve employee collaboration](#), workspace flexibility and productivity. But organizations continue to struggle with security and governance concerns. In a mobile world, data lives on many devices in many places, even when using an on-premises OFS solution.

There's no such thing as a sure thing when it comes to security; data is at risk whether within the four walls of the enterprise or at a [service provider](#). Service providers realize the scrutiny they're under and most have invested much of their venture capital or parent company funding in beefing up security. It's amazing just how far these companies have come from a security and control standpoint in just the past 18 months. I was recently told by an IT administrator that he had to wait for Active Directory integration before deploying a corporate OFS solution—most vendors overcame that hurdle a year ago and have moved on to much deeper reporting and controls. Many are even integrating data loss prevention and [information rights management controls](#). With the plethora of offerings available (ESG is tracking roughly 60 at the moment), there's likely something for everyone, no matter what your level of threat or concern. ■

TERRI McCLURE is a senior storage analyst at Enterprise Strategy Group, Milford, Mass.



Big data storage myths

Big data is a big deal for storage shops, and a clear understanding of what it means—and doesn't mean—is required to successfully configure storage for big data apps.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE

DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?

THE CLOUD TO
THE RESCUE

HOW TO SET UP STORAGE
FOR HYPER-V VMs

BACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

LOVE THE IDEA of changing the world through big data technology. [Big data](#) promises we'll all be IT superheroes just by storing more raw data than ever before and then using parallel processing techniques to yield great new insights that will catapult our company to the top. Good storage is costly and the rate that interesting new data is produced increases daily, but the Apache Hadoop product calls for leveraging scale-out commodity server nodes with cheap local disk.

Of course, there's more to it. Conceptually, big data products bring new ways to store and [analyze the mountains of data](#) that we used to discard. There's certainly information and insight to be mined, but the definitions are fuzzy, the hype is huge and the mining technologies themselves are still rapidly evolving.

Adding to the confusion, big data technology has been

enthusiastically marketed by just about every storage vendor on the planet. But despite the marketing, I believe it's just a matter of time before every competitive IT shop has a real big-data solution to implement or manage, if only because of staggering data growth. For those just [setting out on a big data journey](#), watch out for these common myths.

MYTH NO. 1: JUST DO IT

A sure way to waste a lot of money is to aggregate tons of data on endlessly scalable clusters and hope that your star data scientist will someday discover the hidden keys to eternal profit.

To succeed with any IT project, big data included, you need to have a business value proposition in mind and an

[HOME](#)[HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE](#)[DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?](#)[THE CLOUD TO
THE RESCUE](#)[HOW TO SET UP STORAGE
FOR HYPER-V VMs](#)[BACK UP THOSE PHONES
AND TABLETS](#)[SAFER SYNC AND SHARE](#)[BIG DATA STORAGE MYTHS](#)[STALLED TECH
REFRESHES PICK UP
STEAM](#)

achievable plan laid out. Research is good and those “aha” moments can be exciting, but by the time big data gets to IT, there needs to be a [more practical goal](#) than just a desire to “see what might be in there.”

MYTH NO. 2: STORE EVERYTHING

One of the problems caused by big data hype is that unrealistic expectations are often built on the premise of “keeping it all.” It may seem plausible for a company to [use a big data platform](#) to keep all its data forever. In fact, Cloudera, the most widely adopted Hadoop distribution among enterprises, markets directly to that point. But is it true that accumulated data will become more valuable over time?

Storage experts, at this point, might want to make a few comments along the lines of, “Is all that data going to be actually accessible, usable, reliable, verifiable, available, secure, protected and, certainly not least of all, affordable in the long run?”

For most organizations, far less than “all” data will prove to deliver potential value. And most data declines in relevance as it ages. The faster you can get to an understanding of where your valuable data “subset” is, the more you can direct your resources and attention to what is likely to be most successful. Somewhat ironically, the less data you store, the more efficient and cost-effective you can be with big data.

MYTH NO. 3: BIG IS SIMPLE

The Apache Hadoop Distributed File System ([HDFS](#)) makes it easy to store lots of high-volume, high-velocity and highly variable data across a scale-out cluster. It does it in a way that makes it easy to process using highly parallel [MapReduce-style algorithms](#) that farm the heavy-lifting compute tasks out to each data chunk. HDFS also provides for in-cluster replication mainly to improve cluster availability.

Hadoop 1.0 currently doesn't support snapshots, mirroring or remote replication. And there are no easy ways to further optimize space or tweak I/O performance.

But as suggested above, HDFS doesn't natively provide advanced enterprise storage features that might be needed to support good data protection or disaster recovery. Although evolving, Hadoop 1.0 currently doesn't support snapshots, mirroring or remote replication. And there are no easy ways to further optimize space (deduplication, compression) or [tweak I/O performance](#) (targeted caching, judicious use of flash or highly parallel streaming).

If you have lifecycle data management or governance

[HOME](#)[HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE](#)[DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?](#)[THE CLOUD TO
THE RESCUE](#)[HOW TO SET UP STORAGE
FOR HYPER-V VMs](#)[BACK UP THOSE PHONES
AND TABLETS](#)[SAFER SYNC AND SHARE](#)[BIG DATA STORAGE MYTHS](#)[STALLED TECH
REFRESHES PICK UP
STEAM](#)

requirements for [data stored in a big data environment](#), you might need to consider an enhanced Hadoop distribution like the one from MapR that provides a full-featured storage service layer that transparently replaces HDFS.

MYTH NO. 4: SERVE EVERYONE

Hadoop represents a new way of processing certain types of data in certain parallel ways. And there are some exciting advances coming (e.g., YARN) that enable Hadoop to become a more universal processing platform. But HDFS doesn't provide a universal data storage service. It's designed and optimized for high read-throughput batch processing, and HDFS has no way to target or deliver I/O performance by dataset or workload.

Data has to be specifically loaded into HDFS. It can be difficult to get new data into and results out of it for immediate use or direct access by other applications using other protocols (e.g., NFS, CIFS). And [Hadoop's combined compute/storage node](#) makes it challenging to grow compute and storage on different vectors.

Breaking the HDFS "local" storage paradigm can make a lot of sense. For example, an enterprise scale-out array like EMC's Isilon provides "remote" HDFS storage to a Hadoop cluster, while actually hosting data in its native storage array file system with multiprotocol access and all its other enterprise array features.

MYTH NO. 5: BIG AND FAST

A common misconception about Hadoop is that it's fast. Actually the core design is all about high-throughput "batch"-style processing, and avoiding the impact of common hardware failures that in many larger-scale computing designs (i.e., supercomputers) limit their ultimate efficiency. Hadoop just wasn't originally intended to be an interactive or real-time system.

Due to demand, there are a lot of projects aimed at ramping up Hadoop's performance and expanding the application "footprint" of Hadoop to better support more interactive workloads.

However, due to demand, there are a lot of projects aimed at ramping up performance and expanding the application "footprint" of Hadoop to better support more interactive workloads. Some of these involve integrating traditional database, streaming data or in-memory processing products. There are also high-performance hardware offerings like [DataDirect Networks' hScaler](#) that take an "appliance" approach with compute nodes running in the same rack as their SFA series storage with a customized Hortonworks Hadoop distribution.

HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE

DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?

THE CLOUD TO
THE RESCUE

HOW TO SET UP STORAGE
FOR HYPER-V VMs

BACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

BIG DATA WILL GET BIGGER

Some people may think big data technology is past its peak stage and is crossing a “chasm” of disappointment, but I think we’ve just seen the beginning of its potential and the start of the evolution of the real value proposition of big data to enterprise IT. Those who have approached it realistically are gaining valuable results.

Big data, in the form of Hadoop, is but the start of a broader change in how data processing will need to be approached, and how future data centers will be designed. Data will continue to increase, processing technologies are in high flux, and the most competitive organizations will strive to wring as much value out

of as much of that data as they can. Today, most enterprises haven’t yet invested in [game-changing big data technology projects](#) intended to move the bottom line, although many have deployed Hadoop as an extract, load and transform (ELT) “ingest” platform for their more traditional data warehousing/business intelligence offerings.

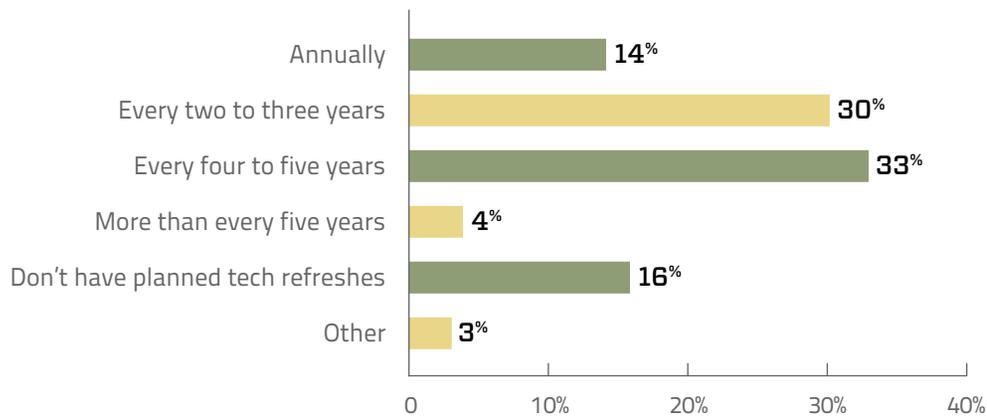
Big data projects are as much storage projects as they are parallel compute. Tell us about your adventures with big data as an enterprise storage solution. ■

MIKE MATCHETT is a senior analyst and consultant at Taneja Group.

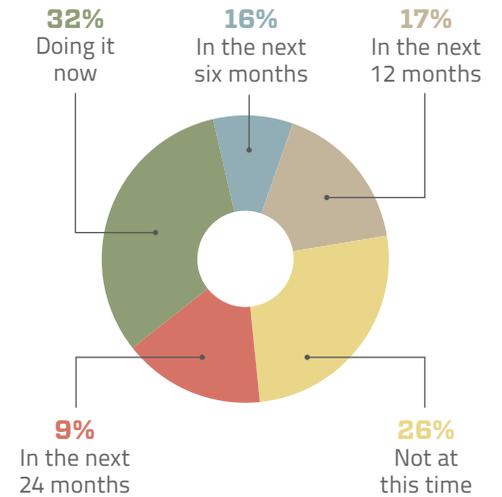
Stalled tech refreshes pick up some steam

NEARLY HALF (44%) of our survey respondents said a tepid economy has either significantly or very significantly affected the frequency at which their companies [refresh their storage technologies](#). But things must be getting better as 32% are undergoing a technology refresh update right now, with another 33% planning to upgrade in the next six to 12 months. Sixteen percent of companies [don't have planned refresh cycles](#), 30% perform them every two to three years, and 33% are on a four or five year rotation. More than half will do DIY refreshes, while 45% will seek outside help from vendors, VARs or consultants. Those vendors will be happy to hear that a [storage technology refresh](#) means replacing older arrays for 72% of respondents—and a whopping 79% who say they need more capacity will add to their currently installed capacity by buying brand-new systems. On the software side, 57% are interested in storage management apps, 43% are looking at storage virtualization software and 41% are sizing up DR/BC software. For storage infrastructure refreshes, 44% plan to add FC switches, 28% will opt for Ethernet switches and 25% plan to shop for FC HBAs. —*Rich Castagna*

HOW FREQUENTLY DO YOU TYPICALLY DO TECH REFRESHES?



DO YOU ANTICIPATE ANY STORAGE SYSTEM TECH REFRESHES FOR YOUR COMPANY?



86

PERCENT OF RESPONDENTS WHO SAY ADDING STORAGE CAPACITY WILL BE PART OF THEIR NEXT TECH REFRESH

- HOME
- HYPE, HOPE OR HOO-HA? YOU BE THE JUDGE
- DO VENDORS REALLY KNOW WHAT PRODUCTS YOU NEED?
- THE CLOUD TO THE RESCUE
- HOW TO SET UP STORAGE FOR HYPER-V VMs
- BACK UP THOSE PHONES AND TABLETS
- SAFER SYNC AND SHARE
- BIG DATA STORAGE MYTHS

STALLED TECH REFRESHES PICK UP STEAM



HOME

HYPE, HOPE OR HOO-HA?
YOU BE THE JUDGE

DO VENDORS REALLY
KNOW WHAT PRODUCTS
YOU NEED?

THE CLOUD TO
THE RESCUE

HOW TO SET UP STORAGE
FOR HYPER-V VMs

BACK UP THOSE PHONES
AND TABLETS

SAFER SYNC AND SHARE

BIG DATA STORAGE MYTHS

STALLED TECH
REFRESHES PICK UP
STEAM

STORAGE MAGAZINE

EDITORIAL DIRECTOR Rich Castagna
SENIOR MANAGING EDITOR Kim Hefner
EXECUTIVE EDITOR Ellen O'Brien
CONTRIBUTING EDITORS James Damoulakis, Steve Duplessie,
Jacob Gsoedl

SEARCHSTORAGE.COM

EXECUTIVE EDITOR Ellen O'Brien
SENIOR NEWS DIRECTOR Dave Raffo
SENIOR NEWS WRITER Sonia R. Lelii
SENIOR WRITER Carol Sliwa
SENIOR MANAGING EDITOR Kim Hefner
ASSOCIATE SITE EDITOR Sarah Wilson

SEARCHCLOUDSTORAGE.COM

SEARCHVIRTUALSTORAGE.COM
EXECUTIVE EDITOR Ellen O'Brien
SENIOR MANAGING EDITOR Kim Hefner
ASSOCIATE SITE EDITOR Sarah Wilson

SEARCHDATABACKUP.COM

SEARCHDISASTERRECOVERY.COM

SEARCHSMBSTORAGE.COM

SEARCHSOLIDSTATESTORAGE.COM

SENIOR SITE EDITOR Andrew Burton
MANAGING EDITOR Ed Hannan
FEATURES WRITER Todd Erickson

STORAGE DECISIONS TechTarget Conferences

EDITORIAL EVENTS MANAGER Jacquelyn Hinds

SUBSCRIPTIONS

www.SearchStorage.com

Storage magazine
275 Grove Street
Newton, MA 02466

editor@storagemagazine.com

TechTarget Inc.
275 Grove Street
Newton, MA 02466
www.techtarget.com

©2013 TechTarget Inc. No part of this publication may be transmitted or reproduced in any form or by any means without written permission from the publisher. TechTarget reprints are available through [The YGS Group](#).

About TechTarget: TechTarget publishes media for information technology professionals. More than 100 focused websites enable quick access to a deep store of news, advice and analysis about the technologies, products and processes crucial to your job. Our live and virtual events give you direct access to independent expert commentary and advice. At IT Knowledge Exchange, our social community, you can get advice and share solutions with peers and experts.

COVER IMAGE AND PAGE 8: RETROROCKET/ISTOCKPHOTO