

EBOOK

SPONSORED BY:



SMARTER STORAGE MANAGEMENT SERIES

By Trevor Pott

IN THIS GUIDE:

- 2** Storage Silos 101: What They Are and How To Manage Them
- 5** RPO/RTO and the Impact on Business Continuity
- 8** Data Storage 101: How To Determine Business-Critical App Performance Requirements
- 11** How Best To Protect and Access Your Company's Data
- 14** Data Fabrics and Their Role in the Storage Hardware Refresh
- 17** The Rise of Hyper-Converged Infrastructure
- 20** A BLUEPRINT FOR SMARTER STORAGE MANAGEMENT
Addressing the Top 3 Challenges Impacting Your Data with Software-Defined Storage



STORAGE SILOS 101: WHAT THEY ARE AND HOW TO MANAGE THEM

Trevor Pott addresses the issues surrounding storage silos and how they can be handled.

If you attempt to research enterprise storage options today, you'll find a great deal of marketing about "storage silos." So just what exactly is a [storage silo](#), why are silos supposedly a problem and what can be done to address them?

A storage silo is any storage that's isolated from the rest of your organization's storage. This is mostly a terrible analogy, because silos contain useful data; they don't arbitrarily separate one thing from another. Because the basic concept of a storage silo is so broad, there are a number of examples.

Let's say that I have two USB keys, each with important data that I require. If one is at home and one is at work, then clearly my storage has been siloed: I do not have access to both the items at the same time, at the same place.

Take both of those USB keys and plug them into a computer, and the storage is no longer siloed. I can move data freely between the storage devices; I can take the data off the storage devices and put onto my computer's primary storage, upload it to the cloud or do whatever I wish.

If I hold both USB keys in my hand, but do not plug them into a computer, is it still a silo? This is a philosophical question, and one I don't have a ready answer to. What I can do is examine this grey area between what is obviously siloed storage and what is obviously not.

If we replace our two USB keys with a SAN and a NAS, at what point does the interconnectivity between them overcome the fact that these are two separate physical devices? And, ultimately, is it useful to even think about storage as being siloed?

Multiple Management Interfaces

If you have a SAN and a NAS, and both live on the same corporate network, in theory you could exchange data between the two. There's the problem that the SAN is block-based and exposes LUNs, while the NAS is file-based and exposes file shares, but in theory it is possible for the SAN to talk to the NAS and vice versa.



Enabling this communication between the two devices may require a third-party piece of software. Alternately, one or both of these devices may contain an appropriate data exchange mechanism native to the unit's software. In any case, NASes and SANs do not speak the same language, and each will have its own distinct management interface. So why is this important?

Let's consider the aforementioned USB sticks. If I plug two USB sticks into my computer, then both of them are managed by a single interface provided by the OS. This interface natively provides the ability to move data back and forth between the two devices. The interface essentially treats each USB stick as if it was equal to any other storage that system had access to; it treats all storage like a commodity.

This unified management interface is an important concept. In order to illustrate why, let's examine virtualization, an IT market closely interrelated with storage in the modern datacenter.

In theory I could set up an unlimited number of individual servers running VMware's free version of ESXi. Without vCenter, each of these ESXi hosts is an island: the hosts are perfectly functional, and they can run VMs. They can address storage to which they are attached. Moving data between hosts, however, requires turning the VM off, exporting the VM from one host's storage to third-party storage, and then uploading that to another host. This is both inefficient and time consuming.

vCenter, on the other hand, provides a single management interface for all hosts. It allows these hosts to use shared storage. This enables the migration of workloads non-disruptively.

Managing five VMs across two isolated ESXi hosts is not generally a big problem. Just as managing a NAS and a SAN within an organization is not generally a problem. Make this 10 hosts—or five NASes and five SANs—and managing these devices is moving past annoying, and into job-impacting frustration. Get into hundreds of hosts or storage systems with thousands of workloads, and management at scale is generally considered impossible without a unified management solution such as vCenter.

Unfortunately, while vCenter is great for virtualization, it's a consumer of storage, not really a metastorage manager in the same way that it is a metaworkload manager. vCenter cannot adequately manage SANs, NASes or storage other than its own vSAN hyperconverged solution.

Organizations trying to manage their compute resources today can turn to VMware. Unified storage management, however, is still something that's unaddressed by organizations of all sizes.

Difficulty Migrating Data Between Silos

Why is VMware such a great example of unified compute management? What makes—or made—VMware special? Personally, I believe that it was vMotion that almost singlehandedly turned VMware from an interesting science project into one of the most critical software companies in the world.

vMotion lets administrators move a workload from one host to the other, with minimum effort. Later versions even made it possible to move workloads with zero impact to the running workload. This was neat, but happened well after vMotion had catapulted VMware onto a trajectory toward

becoming a tech titan. It was the convenience that VMware brought to workload management that really made it attractive, even to those individuals who were opposed to the idea of running more than one workload per physical server. Never underestimate the value of convenience.

Imagine being able to consume multiple storage devices as equals, whether that storage was local to a server, was a NAS, was a SAN or was in the cloud. Imagine being able to consume these multiple storage devices invisibly, as if each of them were part of a single solution.

Imagine being able to consume multiple storage devices as equals, whether that storage was local to a server, was a NAS, was a SAN or was in the cloud.

Some vendors offer such solutions. Most of the vendors offering these solutions require you to throw away all of your existing storage and commit to that vendor as the only vendor that will provide you storage. While some of these single-vendor approaches to the storage management problem can be useful, vendor monoculture is often more trouble than it's worth in IT.

When VMware rose to power, it did not say, "We can do all of these neat things that can make your x86 applications behave sort of like they were on a mainframe, but only if you buy VMware storage, VMware servers and so on." What VMware did was allow organizations to use their existing servers or to buy servers from almost anyone they chose.

VMware turned those servers—and, consequently, the server manufacturers—into commodities. In other words, it broke down the silos that existed between servers. VMware didn't break down the silos between the individual workloads, but rather the silos between the lifecycle management of workloads on those servers.

Many organizations need a similar solution for storage today. Something that allows data to move seamlessly between devices, allows devices to be added or removed as needed without disrupting an organization's storage, which can integrate with or replace an organization's disaster recovery solution.

Consistency

The analogy that links VMware's management of VMs to storage management falls apart when we talk about the maturity of the markets at the time that interconnected management entered the scene. When VMware began commoditizing servers, workload lifecycle management solutions were pretty primitive.

Layering a management solution over the top of your storage isn't particularly helpful if it means you have to give up the storage features for which you paid so much money.

Best-of-breed solutions used scripted installs and imaging, but this was still a highly manual, tedious and long process. There wasn't a lot about how workload management was being done that organizations were particularly interested in preserving. This is not true of storage.

Storage has any number of enterprise features that are the very reason we buy various storage solutions. Data efficiency such as deduplication and compression is one such feature.

Storage pooling, tiering between multiple classes of storage, snapshotting, cloning, and many more features are all important and expensive.

Layering a management solution over the top of your storage isn't particularly helpful if it means you have to give up the storage features for which you paid so much money. On the other hand, if your proposed meta-storage management solution brings these enterprise features to all storage regardless of native support, this significantly enhances the value of your total storage footprint without reducing the value of any individual storage unit.

The storage solution that we all want would allow us to manage all of our storage from a single interface. It would move data seamlessly between storage devices while ensuring that the appropriate number of data copies was being kept. Perhaps most important, it would provide a complete suite of enterprise storage features, regardless of the storage device being used.

Next-generation storage fabrics matching this criteria are just now beginning to see mainstream adoption. There are multiple vendors on the market today, each offering their own take on this sort of meta-storage management. Only time will tell if any of them manage to catapult themselves into becoming the true VMware of storage.

HOW RPO AND RTO IMPACT BUSINESS CONTINUITY

All backups are not created equal. While having any backup is better than having no backup, the different approaches to backups can have dramatic effects on business continuity, should disaster strike.



time it takes that data to be sent to the backup repository.

If you make a backup of a workload every 15 minutes, and you do so to a storage device that can accept this backup instantaneously, then the RPO of this backup solution is 15 minutes: In the worst-case scenario, rolling back to the last backup represents a loss of 15 minutes' worth of data.

If, however, you're sending the backup data off-site, this can get a lot more complicated. Assume for a moment that a backup is taken and sent to a Cloud Storage Gateway (CSG). This CSG acts as a buffer, which absorbs the backup and slowly unspools it to a cloud storage location. If this process takes the whole 15 minutes between backups, then the RPO is in fact 30 minutes: 15 minutes between backups, plus 15 minutes to ship the backups off-site.

To understand what makes one backup approach different from the next, you must understand the concepts of Recovery Point Objective (RPO) and Recovery Time Objective (RTO). Both have been discussed [at length](#) numerous times, but here's a quick recap:

RPO

RPO represents the maximum data loss you can afford as a function of time. Most backups occur periodically, with each backup event being a recovery point. The RPO is a combination of the time between recovery points and the

RTO

RTO is the measure of how quickly workloads can be restored from a backup. If the backup can restore workloads and/or data instantly, then the RTO is zero seconds. No RTOs are realistically zero seconds, but some backup designs can get RTOs down into the single-digit seconds, even for complicated workloads.

RTOs measured in single-digit seconds usually mean that the backup storage is able to mount the data or workloads directly, without having to transfer the data to another storage device. Both on-premises and cloud-based backup

solutions allow this, meaning that near-zero RTO is simply a matter of both design and money.

Not all organizations have the money to accomplish this, however, or are restricted in the backup designs they can use due to regulatory compliance concerns. In these cases, the time it takes to restore from backup becomes part of the RTO calculation.

Time to restore from backup can be due to multiple factors. In extreme cases, backups may need to be physically couriered from a backup vault. Where cloud storage is used, backups may need to be downloaded from the cloud before they can be restored. Even when backups are on-premises, there is often a delay imposed by copying data from the on-premises backup solution to the production environment.

Data fabrics are an emerging solution to both production and backup storage that solves the problems associated with storage, which spans multiple sites and infrastructures.

Local Backup Copies

While this is a good introduction to the basics of RTO and RPO, it's also a gross oversimplification. The design of an organization's backup approach can affect both RPO and RTO in context-dependent ways.

Many organizations follow the 3-2-1 approach to backups: make three copies of your data, storing that data on at least two mediums, with at least one copy off-site. The most basic design that meets this requirement is to back up all of your data to a local backup repository, and have that backup repository spool data to a cloud storage location.

For the overwhelming majority of backup scenarios, this means that RPO and RTO are relatively small: there are copies of data on the same site as the production workloads, so it shouldn't take long to run the backups, or to restore them. Most restore requests are made because of accidental deletions, or similar mundane errors. For these scenarios, local backup copies make abundant sense, and they minimize

the effects of data protection events on business continuity.

The ultimate evolution of the local repository is the immutable snapshot. Instead of having to copy data off of a production storage solution in order to back it up, near-instantaneous snapshots are taken. Here, the time required to take a backup can be reduced to single-digit seconds. This allows for more regular backups to be taken, lowering RPO.

Because the snapshots live on the same infrastructure as the production workloads, RTOs are functionally instantaneous. Snapshots are a good thing, but they come with a caveat: If something happens to that production infrastructure, both the production copy and all the snapshots are affected at the same time.

In order to alleviate this, most storage solutions that make heavy use of snapshots also offer replication. Synchronous replication means that data is written to both the primary and the backup storage device simultaneously. Asynchronous replication means that there's a delay (often heavily influenced by distance and the speed of light) between writes occurring on the primary and secondary storage solutions.

All of this is useful and good, but becomes further complicated when an organization's IT infrastructure spans multiple sites, not all of which are directly operated by the organization. Today, it's not uncommon to see organizations that operate multiple datacenters, make use of infrastructure from one or more services providers, in addition to making use of multiple public clouds.

Data Fabrics

Data fabrics are an emerging solution to both production and backup storage that solves the problems associated with storage, which spans multiple sites and infrastructures. Data fabrics absorb all of an organization's storage into a single solution, and then distribute that storage throughout the fabric based upon rules and profiles applied to each different class or container of storage.

Data fabrics rely on administrators describing storage to the fabric. Administrators attaching storage need to ensure that the fabric understands failure domains. Fabrics may implement failure domains as zones or sites, and administrators need to be careful that all storage that can be impacted by a single disaster is appropriately grouped.

Once appropriately set up, fabrics can be a powerful storage solution with robust data protection. Administrators

may, for example, tell the data fabric that they need a given storage LUN to exist in three copies throughout the fabric, and that regular snapshots be taken every 5 minutes.

The fabric would then ensure that the production site has a copy of the production data, and all relevant snapshots. In addition, two more copies would be distributed throughout the fabric, wherever storage capacity existed. Unless told to keep a second copy on the production site, the fabric would ensure that none of the additional data copies were kept on the same site as the production data.

With data fabrics, the design of an organization's storage and backup solutions is constantly adapting to meet present needs.

Data fabrics move data around as needed to meet performance and resiliency requirements. They typically have a combination of high-performance/high-cost storage, and low-performance/low-cost storage. Cold data—for example, snapshots that haven't been accessed in some time—would be moved to low-cost storage (archival to the cloud, for example). Frequently accessed data would be retained on high-performance storage.

Intent-Based Storage

Data fabrics represent a new approach not only to designing production storage, but to designing data protection solutions, as well. Instead of assigning storage to individual physical devices, or worrying about juggling the RTOs and RPOs of individual workloads, data fabrics provide an intent-based approach to storage.

The basic concerns of RPO and RTO do not go away with data fabrics. Instead, the fabrics abstract the details of these concerns behind an easy-to-use management interface and a set of algorithms designed to manage the common, basic storage tasks in an automated fashion.

Administrators define their intent in the form of storage policies and profiles, and the data fabric takes care of the rest. When more performance is needed, you simply add more performant storage to the fabric, and the fabric rebalances. When more capacity is needed, you simply add more capacity.

With data fabrics, the design of an organization's storage and backup solutions is constantly adapting to meet present needs. This allows administrators to focus on providing the hardware and bandwidth required to keep things functioning smoothly, instead of worrying about how to connect it all together.

DATA STORAGE 101: HOW TO DETERMINE BUSINESS-CRITICAL APP PERFORMANCE REQUIREMENTS

Measuring IOPS, latency and throughput are important for determining the characteristics of physical storage that's added to a modern data fabric.

Storage is not one-size-fits-all, and determining what the right storage is for your needs is difficult. There are many storage vendors out there, each selling a combination of complex engineering, bitter experience and marketing pixie dust wrapped up together into a saleable product. So how do you determine which bit of storage is the best for the job, and what measurements can help you accomplish this task?

As with most things, when you attempt to determine the relative value of storage there are both quantitative and qualitative measures to consider. Quantitative measures of storage are what can be readily determined through benchmarking. Input/output operations per second (IOPS), latency and throughput are the big benchmark measurements.

Qualitative measurements are more subjective. These include things like ease of use, or the importance and utility of various storage features to the organization considering the storage.

Quantitative Measurements

IOPS, latency and throughput are highly interdependent measurements. As a general rule, the higher the IOPS the lower the throughput, and the more stress put on a system (IOPS or throughput), the higher the latency.

With the right benchmarking, it's possible to make some straightforward graphs of basic storage performance. On one side of the scale there's extreme IOPS with minimal throughput, on the other side there's extreme throughput with minimal IOPS. This effect is more dramatic with magnetic storage media, but the basic principle applies to solid-state storage, as well.

Storage obeys certain basic commands: read, write, modify and delete. These commands represent inputs and outputs (I/Os), and IOPS is a measure of how many of these



commands can be performed by a storage solution in a given second.

The commands issued affect the numbers obtained. Flooding a storage solution with 100 percent read requests will result in a different IOPS reading than 100 percent write requests. Magnetic storage media will respond differently to modify requests than solid-state media. The block size of the I/O requests (measured in [kibibytes](#)), the number of simultaneous requests, as well as the randomness of the requests, also impact the results obtained.

At first glance this may make IOPS seem a highly subjective measurement, but it's not. While a configuration of 70 percent write/30 percent read/64K block size/100 percent random will differ greatly from a configuration of 50 percent write/50 percent read/16K block size/50 percent

random, most storage solutions should perform identically under the same configuration.

This makes it important to pay attention to the testing profiles used by vendors or reviewers, and to compare storage solutions using the same testing profiles. A headline benchmark result of 1 million IOPS, for example, may have been achieved using a benchmark profile designed to give a maximum IOPS instead of one designed to mimic “real world” usage scenarios. In addition, each organization’s “real world” usage will differ from the next.

With all of that said, if you take a given storage profile and throw it at multiple storage solutions, the result is a good understanding of the performance characteristics of the underlying storage. Identical storage profiles for benchmarking form the basis for a rational comparison between storage solutions, however, the raw numbers can only tell so much about the overall value of that storage solution to an organization.

Qualitative Measurements

In a perfect world, storage solutions are benchmarked both with all of their advanced storage features off and on. Features such as data efficiency, tiering, caching and so forth all affect how storage will perform under various circumstances, and they’re all highly variable based on the data being used for testing.

Consider two storage arrays: one has hardware-assisted deduplication and compression, the other performs these data efficiency tasks entirely in software. Let us assume that with data efficiency turned off, these arrays perform identically.

Running a benchmark against these systems with data efficiency turned on you might discover that both solutions perform identically up to a specific threshold. Once data volumes reach this threshold one of the arrays will reach the maximum amount of data it can perform data efficiency operations upon per second, and a demonstrable difference between the two arrays will have appeared.

Similarly, the arrays could use completely different approaches to different data efficiency. Different approaches to data efficiency have different performance costs, and produce different results in terms of data reduction. The effectiveness of data efficiency at the scale of a single storage array is also often quite different when compared

to solutions that span multiple devices and perform data efficiency tasks across the entire solution, instead of just at the level of an individual array.

Different types of storage have different levels of reducibility. RAW images, for example, compress more readily than JPEGs. Virtual desktop infrastructure (VDI) VMs generally provide a much higher level of deduplication than a collection of VMs all running different OSes.

In addition to this, storage features such as data efficiency typically compete with one another for array resources. Features like data tiering take processing power; if a storage solution is busy working on data efficiency, it might have to delay tiering tasks or vice versa. Which advanced storage features are enabled—and how they’re used—can make a noticeable difference on the performance delivered, even when comparing two arrays that would perform identically with those features off.

In a perfect world, storage solutions are benchmarked both with all of their advanced storage features off and on.

The Advanced in Advanced Storage Features

At first glance, advanced storage features would seem to be quantitative. Features like data efficiency are affected by multiple variables, but if you can control enough variables they should hypothetically perform the same every time. There is some truth to this, but it’s also more complicated than that.

Fifteen years ago, how a storage solution performed data efficiency probably would’ve been a quantitative. The only data efficiency most storage used was compression, and an individual storage array would apply the same compression algorithm to all data, all of the time. That was then, this is now.

Today, storage tries to be smart. Some solutions will test a small piece of data to see how compressible it is before deciding whether or not to compress the whole data stream, or which algorithms to apply. Deduplication comes in flavors, and may only be applied on certain tiers of data, when the array is idle or in response to other parameters.

Tiering of data between different storage media—or even between different arrays, sites or clouds—can occur based

on any number of criteria, and the criteria themselves can change as the array learns storage patterns and adapts which features it implements under which circumstances. The smarter storage becomes, the harder it is to predict.

Profiles, Policies and SLAs

Storage solutions are no longer discrete, self-contained items. An individual SAN or NAS is often just one part of a larger whole. When multiple arrays are joined together with server-local storage, cloud storage and who knows what else, an organization's storage becomes a data fabric.

A single data fabric today could easily join multiple on-premises sites' worth of storage, services provider storage and public cloud storage.

Data fabrics store data on multiple devices. These devices can use multiple physical storage media, be located across multiple sites, and even across multiple infrastructures, where the different infrastructures are owned and operated by different organizations. A single data fabric today could easily join multiple on-premises sites' worth of storage, services provider storage and public cloud storage.

Data fabrics typically have the ability to add or remove storage in a non-disruptive fashion, meaning that the overall physical composition of the data fabric is itself a variable. This changes how you must measure data fabrics, both qualitatively and quantitatively.

Hyper-converged infrastructure (HCI) and scale out storage are examples of simplistic data fabrics. HCI and scale out storage both take storage located inside individual servers, lash all that storage together and present it using a single interface. In the case of HCI, workloads are run on the same nodes that supply storage to the fabric, while scale out storage dedicates nodes to storage alone. Data fabrics can get much more complicated, however, and consist of any or all storage that an organization uses.

Because data fabrics are a logical construct instead of a fixed physical asset, you would rarely attempt to measure the

performance of the whole fabric. Instead, profiles, policies and service-level agreements (SLAs) are set in the data fabric, and tests are performed to see if the fabric can deliver.

What is the maximum ingestion rate of data into the fabric? Can the fabric deliver a LUN with x IOPS, y throughput and z latency? How many of these can it deliver simultaneously? Does the fabric warn when it has been asked to deliver performance beyond its capabilities, and how does it do this warning?

The response of fabrics to change is a critical measure. If you manage to ask more of a fabric than it's currently capable of delivering, how smoothly will it absorb additional storage into the fabric, and how quickly will that help meet the storage demands, especially if the fabric is overburdened? Does the fabric perform parallel I/O, or is everything funneled through a single orchestration node?

Fabric Softener

The existence of data fabrics doesn't make classic qualitative measurements of storage irrelevant. Data fabrics are an orchestration layer that smoothes together multiple storage devices, layers a universal set of advanced storage features on top, and then presents a single storage interface for storage consumers. This makes storage easier to use, but it does not magically solve performance problems.

In order for data fabrics to do their thing, they must be supplied with adequate amounts of task-appropriate storage. If the fabric is supplying storage with unacceptably high latencies, you might consider adding NVMe solid-state drives to the mix. If capacity is a problem, but performance is fine, then perhaps a big box of magnetic hard drives is the ticket.

Measuring IOPS, latency and throughput are important for determining the characteristics of physical storage that's added to a modern data fabric. That said, extended proof-of-concept testing with real world workloads and copies of real data are equally important for teasing out the qualitative performance characteristics of a storage solution, especially when the data fabrics used start doing spooky things like tiering solid data to the public cloud, or dynamically changing data efficiency approaches based on load.

HOW BEST TO PROTECT AND ACCESS YOUR COMPANY'S DATA

As much as preserving our data matters, access to that data is equally important.

In 2018, you have to work to find an organization that isn't utterly dependent on computers. Data has become as important as any tangible asset. As much as preserving our data matters, access to that data is equally important. Uptime matters.

One of the most important assets any organization has is its data. Unfortunately, as we are all aware, many organizations ignore data protection until it's too late. Ransomware became such a problematic epidemic specifically because of the habitual neglect that organizations have for basic backups.

In 2016 Dell EMC ran [a survey](#) on data protection. This survey reported that the average cost of data loss was \$900,000 USD, and the average cost of downtime was \$555,000 USD. As the cost of both data loss and down time have only gone up throughout my lifetime, it's reasonable to assume that if you were to run the same survey today, those costs would only be higher.

Unfortunately, backups are easy to see as a nice-to-have

instead of a must-have, and this is exactly the approach both organizations and individuals seem to take by default. For this reason I've written a seemingly endless number of articles about backups. I'm sure I'll write many more.

Yet backups aren't the only aspect of data protection that organizations ignore. Storage availability is all too frequently neglected, often with punitively expensive results.

That Speed of Light Thing

Backups exist to protect against data loss. This loss can be because data is deleted or overwritten, or because the storage device upon which it rests has been destroyed. Because backups primarily exist to protect against statistically unlikely events, restoring from backups is rarely expedient.

In addition, there often exists a gap in backup data called the recovery point objective (RPO). RPO is a measure of the time between backups. More important, RPO is a measure of the maximum amount of data that might be lost if you have to roll back to the last backup.

Data is lost when rolling back to a backup because it takes time for data to travel between the production system and the backup location. Even if you're using continuous data protection (CDP), which is essentially real-time streaming of writes from the production location to the backup location, the speed of light says that there will always be a delay in getting data from A to B.

That delay represents the theoretical minimum RPO of any backup solution. In many cases the delay is only a few milliseconds, but it only takes a few milliseconds to store dozens or even hundreds of transactions on a Web



site, or a critical update to a human resources record.

Within a single datacenter, the distance that data would have to travel between any two systems is so short that the travel time is less than the speed at which storage devices perform transactions. In many cases, real-time replication of data is possible between two systems within the same city, with a distance between the two systems of 100km being the generally accepted maximum distance.

100km isn't very far, in disaster avoidance terms. A power failure, hurricane, earthquake or other event could easily affect both locations. For this reason, backups are typically stored at much greater distances from the production system, even though this means that rolling back to the latest backup could mean data loss: data is sacrosanct, and it must be protected against even statistically unlikely events.

When the primary storage has a little lie down—regardless of the cause—avoiding the costs of downtime requires that there be a secondary storage array to take over.

High Availability vs. Backups

Backups are incredibly useful for recovering deleted files, but—because of that RPO gap—far less so for running workloads. That being said, the equipment on which running workloads operate can, does and will eventually fail. Nothing lasts forever.

Because the cost of downtime is so high, investment in technologies that allow for real-time replication of data between two or more storage devices with no data loss (an RPO of zero) is usually called for. This is high availability (HA). As already discussed, the distance between storage devices in an HA solution is largely dictated by the speed of light.

For many organizations, HA consists of two storage devices located physically one on top of the other on the same rack within a datacenter. For some organizations, HA will be performed between storage devices located at two different datacenters that are situated on a Metropolitan Area Network (MAN). In the finance industry in particular,

it's common to have two datacenters on a MAN providing HA, and a third datacenter in another geographic location as the disaster recovery site.

Data Migration Affects Uptime

It's easy to understand HA in terms of device failure. Bobby Breakit trips over his untied shoelaces, takes a header into the storage rack, and makes the primary storage array go boom. Bad Bobby, do not break it.

There are, however, a number of other perfectly routine and innocuous reasons to invest in HA; updates, for example. Everything in IT needs to be updated eventually. Many HA storage solutions can be set up to seamlessly switch between physical devices, allowing administrators to apply updates to the secondary device while the primary one continues to serve workloads. Fail back over to the other device, and you can apply workloads to that one, too.

When the primary storage has a little lie down—regardless of the cause—avoiding the costs of downtime requires that there be a secondary storage array to take over. In a perfect world, the handoff between the two is seamless, and running workloads never notice that their storage is now being served from a different device.

Not all storage solutions are capable of this kind of HA. Some solutions can replicate data between two devices in real time. However, failover between the two devices is not seamless, and workloads will crash if a failover event is forced. These less advanced HA solutions don't lose any of the data written to the storage devices, but require workloads to be restarted if a failover is forced, which causes at least some downtime.

This latter failover scenario is quite common when data movement is occurring between two dissimilar storage devices. Data migration activities due to datacenter upgrades, or due to space saving or archival efforts are common causes of outages. These outages are often a "hidden cost" of storage upgrades and maintenance: in addition to the cost of the equipment, you must factor in the time necessary to switch over to new devices, or repoint workloads to data that's been moved.

Avoiding Downtime

While all of this is still true for traditional storage arrays, storage technology has come a long way in the past decade.

Data fabrics—often nebulously described under the nearly useless header of “software-defined storage”—remove the need for downtime, even when adding or removing storage devices from an organization’s datacenter.

The short version of how data fabrics work is as follows: a highly available cluster of servers acts as a data presentation layer. In turn, this data presentation layer controls any and all storage that it’s fed.

Backups are a good and necessary thing, but because of the RPO gap, most organizations prefer to avoid rolling back to them if at all possible

The data presentation layer presents storage to running workloads. Because this presentation layer is itself an HA cluster, the presentation layer can survive the loss of a node that’s performing data presentation activities.

The data in a data fabric is spread across all storage devices made available to the fabric. This storage could be traditional storage arrays that have offered some or all of

their storage to the data fabric. And it could be in the form of whitebox servers full of disks.

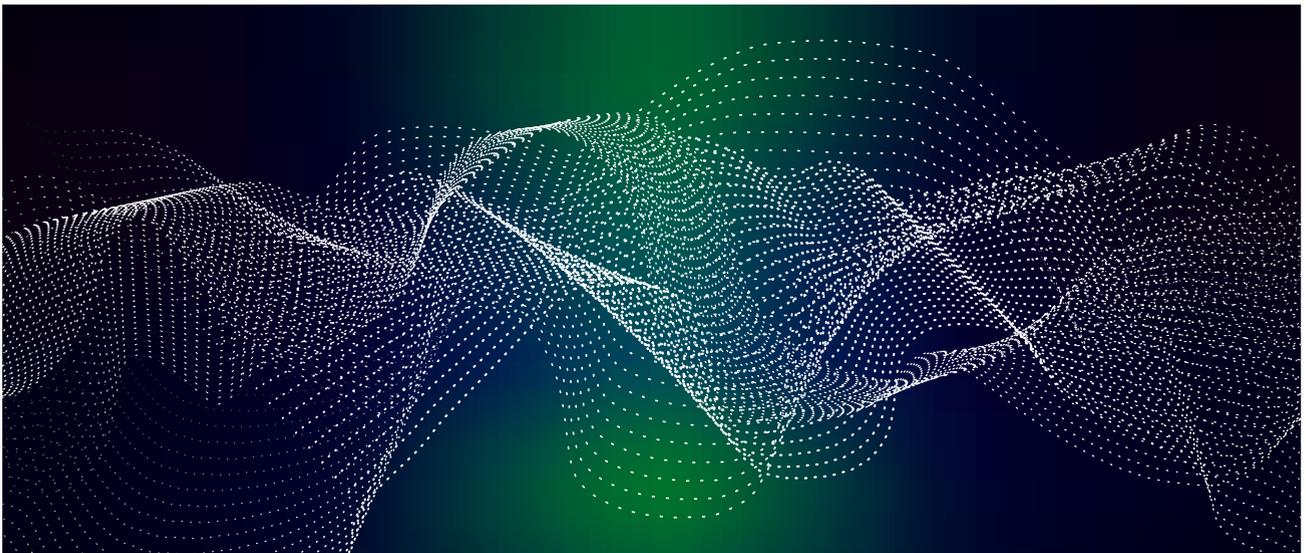
The storage could also be drives attached to the nodes hosting the data presentation software itself. When a data fabric is designed in this manner, storage industry nerds call it “scale out storage.” When a data fabric not only puts storage drives in the data presentation nodes, but also allows administrators to run production workloads on those same nodes, it’s generally referred to as hyper-converged infrastructure (HCI).

Backups are a good and necessary thing, but because of that RPO gap, most organizations prefer to avoid rolling back to them if at all possible. Doing so requires HA, and data fabrics are the best solution currently available to provide HA.

The true beauty of data fabrics is that you don’t have to throw away your existing investment in storage hardware to take advantage of them. There are numerous solutions available that can marry your traditional arrays with whitebox servers, and even blend this with HCI. When your next storage refresh crops up, talk to your vendors about data fabrics: They could help avoid costly downtime, and they’re quickly becoming the mainstream storage solution of choice.

DATA FABRICS AND THEIR ROLE IN THE STORAGE HARDWARE REFRESH

Data fabrics allow organizations to go a step further and spread data across multiple arrays such that entire storage arrays become redundant.



How do you refresh storage? Or, more specifically, what happens to perfectly good storage when the support contracts are up? Should organizations extend the support contract at potentially punitive rates in order to save on capital expenses, or can the previous generation's storage see valid use even outside support?

These questions all have different answers depending on which person you're asking. Vendor employees will generally provide a very earnest response that includes a number of reasons that nobody should ever run anything in IT without a support contract. There are also apparently very good (from whom?) reasons why storage array support contracts need to get dramatically more expensive after the second or third year.

It's in the interests of the customer, you see, to be not-so-subtly nudged toward refreshing their storage every three years; even if that storage happens to be working just fine, and meeting the organization's needs. While this approach may work in organizations where money is no

object, throwing away perfectly serviceable storage and then spending hundreds of thousands or even millions of dollars to buy replacement equipment seems to be a sore point for most other organizations.

Drive Locking

Regardless of the specific flavor of storage in use, there are always two components to any datacenter storage. There's the bit that stores the data (drives), which can be either rotating magnetic media (spinning rust), or solid-state drives (SSDs). There's also the bit that chooses which drives will hold what data, typically referred to as a controller.

The drives in an enterprise storage array are usually not very different from the drives in a consumer network-attached storage (NAS). There's usually a custom firmware applied to the enterprise drives, but how much this alters the performance or reliability of the drives is still, after decades, mostly unknown.

What the custom firmware does do is give storage vendors a way to prevent organizations from simply replacing the drives in their storage array with generic, off-the-shelf replacements. Many storage arrays—especially the more expensive ones—will reject drives that do not have the vendor’s custom firmware, even if the drives purchased are “enterprise” or “datacenter” drives.

With consumer storage devices, if the drives are starting to fail, or if free space is getting a little low, you can simply log on to Amazon and buy a few dozen replacement drives. The vendor of a consumer storage solution makes its money in selling the controller, and it uses unmodified consumer hard drives. If consumer array vendors can do this, why can’t enterprise storage vendors?

Storage Refresh Rationale

The rationale for the enterprise storage vendor approach to drive replacement is complicated. Vendors who are selling enterprise storage can’t afford to sell storage arrays that aren’t rock solid. Consumers and small businesses can tolerate failure rates and downtime that enterprises won’t. This means, among many other things, that vendors don’t want anyone using drives that haven’t been 100 percent fully tested for use with the controller’s hardware and software.

This is complicated by the fact that consumer drives aren’t homogenous. A Western Digital Gold drive, for example, is considered to be an enterprise-class hard drive by Western Digital. If you wanted a top-notch hard drive to install into a server, or put in your consumer NAS, this is the drive for you. But Gold is a brand, it’s not a homogenous line of identical drives.

Even within a specific capacity, say 12TB drives, there are multiple production runs of the drives, each with their own quirks and foibles. In addition, Western Digital is constantly tweaking the design and firmware of the drives. The result is that over the full run of 12TB Western Digital Gold drives there may be dozens of different models of drives that all share the same brand.

If consumer storage can cope with this, why can’t enterprise storage? The answer is that enterprise storage absolutely can cope with this, but that storage vendors don’t want to deal with the hassle, and for good reason: Sometimes failure rates can get ridiculous. The three big cases of high failure rates that immediately come to mind are (in reverse

chronological order), the [Seagate ST3000DM001](#) (dubbed the “Failcudal”), the OCZ [debacle](#) and the [IBM Deskstar 75GXP](#) (dubbed the “Deathstar”). The OCZ debacle is notable because it shows us that SSDs can be just as awful as spinning rust.

There are rumors of other drive models with failure rates that are worthy of inclusion in that list, but getting data on failure rates is hard. Storage vendors (with the notable exception of Backblaze) keep such data a closely guarded secret, and it usually takes a judge ordering disclosure to get access to it.

Because this data isn’t shared throughout the industry, figuring out which drive models are likely to cause problems is next to impossible. What vendor wants to be on the hook for the reliability of a storage array running mission-critical workloads if the customer can just start putting in fail-class drives?

In addition to worries about drive failure rates, vendors worry about controllers. The physical hardware in most controllers is no different from what’s in your average x86 server, and that means that most of them will last for at least a decade. But a controller designed to support 1TB through 6TB drives may not perform appropriately if you start cramming 12TB drives into it.

Storage arrays have RAM that they use for caching, as well as for executing storage controller functions such as deduplication, compression, encryption and so forth. In enterprise storage arrays some or all of this RAM may even be non-volatile RAM (NVDIMMs), which is designed to allow the controller to use write caching without having to worry about data loss due to power failures.

Installing larger drives than the array is designed for could mean that the controller no longer has enough RAM to do its job appropriately. In the case of controllers with NVDIMMs, part of the reason for three-year refresh cycles is that both the batteries and the flash that enable the NVDIMMs to be non-volatile have lifespans much shorter than the rest of the controller.

Data Fabrics

From the point of view of the customer, vendor drive locking is annoying. It appears to be nothing more than a cynical attempt by the vendor to squeeze money out of a captive audience, and milk its customer base those vendors absolutely did. Somewhere around 2009 a large number of startups began to emerge with new approaches to storage,

most of which were aiming to capitalize on frustration with the status quo.

Complicating the simple David-versus-Goliath narrative, however, is that storage vendors absolutely did have good reasons for engaging in drive locking and other practices aimed at driving short refresh cycles. At least, they were good reasons from a certain point of view.

Storage vendors sell their storage assuming that they will be the only storage solution in use for a given workload. Customers are expected to perform backups, but the storage array itself must always be up, always delivering rock-solid storage. Lives may depend on it. This approach to storage was established before data fabrics existed, and entire generations of engineers, storage administrators and storage executives were raised with this thought process.

Today, however, we do have data fabrics. A storage array has redundant drives to protect against drive failure, and even redundant controllers to protect against controller failure. Data fabrics allow organizations to go a step further and spread data across multiple arrays such that entire storage arrays become redundant.

Most data fabrics offer much more functionality than simply writing data to multiple storage arrays in order to protect data against the failure of an entire array. Data fabrics also continuously reassess the performance of all storage in the fabric, and then place data on appropriate storage, based on the profile settings for the workloads using that storage.

If an array is underperforming (say, because you stuffed it full of higher-capacity drives than it was designed for), the data fabric treats that array as a cold storage destination. This means that the data fabric would only put data blocks on that array that had been determined to be unlikely to be frequently accessed, or which are part of a workload whose storage profile says that workload is latency-insensitive.

Frankenstorage

The beauty of data fabrics is that you can use any storage you want. Do you want to build whitebox storage out of Supermicro's [utterly ridiculous lineup](#)? Go right ahead: You'll find units allowing you to cram [90 3.5-inch](#) spinning rust drives into 4u. Or how about [48 NVMe drives](#) in a 2u server?

Supermicro is easy to point to for extreme whitebox solutions, but most server vendors are providing

[Opencompute solutions](#), and Opencompute Storage is now a thing. There's also the [Backblaze Storage Pod](#) for those looking to grind every last penny out of their whitebox storage designs.

The basic premise of data fabrics is that you put together a whole bunch of spinning rust for bulk storage purposes, and then add SSDs (increasingly NVMe SSDs) for performance. If you need more capacity you add more spinning rust. If you need more performance, you add more SSDs. You let the fabric figure out where to put the data, and just by adding what you need, when you get the job done.

But data fabrics don't have to be whitebox storage. Many data fabrics allow administrators to add any storage they happen to have access to into the fabric. Your old storage array that's out of support? Add it to the fabric, scour eBay for replacement drives and run that array right into the ground. Cloud storage? Punch in your credentials and voilà: additional capacity.

If and when your no-longer-supported array finally dies, that's totally OK: the data that was stored on it is replicated elsewhere in the fabric, and the fabric will continue serving workloads uninterrupted. Most data fabrics will sense the loss of the array and begin replicating additional copies of the data to compensate for the failed storage.

Traditional storage administrators often refer to data fabrics derisively as "Frankenstorage." Data fabric vendors are terrified of anyone using the term because they feel it has negative connotations, but in fact it is the perfect analogy. Data fabrics allow organizations to take any number of different storage solutions from any number of vendors and weld them together into something that is far more than the sum of its parts. Like Frankenstein's monster, data fabrics are deeply misunderstood, and the subject of unwarranted Fear Uncertainty and Doubt (FUD).

Any new technology in IT faces opposition, but data fabrics are no longer just a science project. There are multiple vendors offering proven solutions, and they're starting to have real-world impacts on how organizations handle storage refreshes. If you're looking at yet another round of forklift upgrades, it's a technology worth considering.

THE RISE OF HYPER-CONVERGED INFRASTRUCTURE

What it means to your data storage plan.

Many of the big-name hyper-converged infrastructure (HCI) vendors were founded in 2009, meaning we're coming up on a decade of HCI. Everyone knows that it's storage and compute in a single system, but what does that actually mean? How does HCI fit within a messy, non-optimal, brownfield environment, and has it lived up to the hype?

To understand HCI, you need to understand a bit about the evolution of storage—or at least storage buzzwords—over the past decade. When HCI emerged in 2009 the buzzword of the day was “software-defined storage (SDS).”

SDS is the perfect buzzword because it's approximately as meaningless as saying “human-built structure.” Yes, human-built structure narrows the field a little. We're clearly not talking about a bowerbird's bower, or a termite mound, but “human-built structure” is still so broad a definition that it has very little real-world utility.

Almost all data storage today is software-defined. Even a lot of physical data storage is software-defined: I myself have implemented filing solutions for hard-copy papers that use barcodes, QR codes or NFC to track where documents are stored, and even help automate lifecycle management of those documents.

What SDS was originally supposed to mean was “storage software that commoditizes storage vendors in the same way that VMware commoditized server vendors.” It eventually came to be used to mean “any storage that isn't a traditional storage array,” which is meaningless.

A storage array in 2018 is an entirely different animal than a storage array from 2009, but won't get called SDS by most IT practitioners because of its lineage. An HCI solution provided by a vendor that requires customers to only use



a narrow set of nodes provided by that same HCI vendor, however, will usually be considered SDS.

Being more clear about this: SDS is a buzzword that, in reality, was coined by startups as a polite way of saying “EMC, NetApp, IBM, HP and so forth are too expensive, so buy storage from us instead.” It has no other real-world meaning.

Data Fabrics

What SDS was supposed to stand for—the commoditization of storage—is still a relevant concept. For a time, the term “storage virtualization” was used. This was intended to draw a parallel to VMware’s commoditization of server vendors, as well as the idea that storage could be moved between storage solutions as easily as one might move virtual machines between hosts in a cluster.

This never took off in large part because VMware was better at actual storage virtualization than most startups. vSphere has a feature called “storage vMotion” that makes moving storage between different solutions simple, assuming your workloads are virtualized with VMware. There’s also VVOLs, which is supposed to make managing LUNs easier, the success (or failure) of which is a discussion best left for another day.

The original concept of SDS has re-emerged in the form of a data fabric. A data fabric is a distributed application that combines all storage donated to it into a single pool of storage. The software then cuts the storage up according to whatever parameters it’s given and serves it up for consumption.

With open data fabrics, storage can be anything: cloud storage, whitebox storage, local server storage, storage-area networks (SANs), network-attached storage (NAS), you name it. That storage can then be made available in any manner that the data fabric vendor chooses to allow: iSCSI, Fibre Channel, SMB, NFS, Object or even appearing to nodes as local attached storage. A truly open data fabric can consume any storage, anywhere, and present that storage for consumption as any kind of storage desired.

While most data fabrics could be as open as described earlier, most aren’t. There’s nothing preventing any data fabric vendor from incorporating any storage source they choose, or emitting storage in any format they wish. The hard part of building data fabric software are the

bits that choose where to put different data blocks, and applying enterprise storage features such as compression, deduplication, snapshotting, rapid cloning, encryption, and so forth. Compared to that, consuming a new type of storage, or adding an emission shim so storage can be consumed in a different format, is easy.

Vendors, however, restrict the capabilities of their data fabrics for all the same reasons that traditional storage vendors engaged in drive locking and pushed for regular [three-year refreshes](#). This brings us to HCI.

HCI-Plus

The vendor of your data fabric software matters, because the vendor determines how flexible that data fabric will be. Let’s consider HCI. HCI is essentially a data fabric consisting of servers that have local storage, which donate that storage to the fabric, and which use their spare compute capacity to run workloads.

While there’s no reason that most HCI solutions couldn’t incorporate other storage sources into their data fabric, many HCI solutions are restrictive. Customers are only allowed to add a narrow class of nodes to a cluster, cluster sizes are constrained to a small number of nodes and the nodes themselves are often not particularly customizable.

More open data fabrics allow for an “HCI-Plus” approach. With open data fabrics customers can buy servers filled with storage and build clusters that behave exactly like HCI. They can also add their old storage arrays or cloud storage to the mix, using them to provide cold or archival storage, snapshot storage, and so forth.

Real-World Data Fabric Use

It’s reasonable to wonder why an HCI+ approach to data fabrics is even entertained. It sounds great on paper, but if you’re going to build a data fabric, why not just build it all out of whitebox servers, use a traditional HCI approach and be done with it? Maybe tack on some cloud storage as an off-site destination for snapshots or cold storage, but is there really a point behind welding on anything else?

The answer to this is complicated. If you’re designing a brand-new deployment, then an HCI+ approach has very limited appeal. Traditional HCI does the job, with even basic cloud replication or backup capabilities showing tepid uptake from HCI customers.

Greenfield deployments, however, don't represent the majority of environments. Most organizations are a messy mix of IT. Different tiers or classes of infrastructure from different departments all refresh at different times. Mergers and acquisitions can often mean organizations are running multiple datacenters, each with their own entirely distinct design and approach to IT.

Many organizations aren't yet ready to embrace HCI—let alone open data fabrics—for mission-critical workloads. For these organizations, a decade isn't enough time for a technology to meet their standards of reliability, they'll stick with the tried, true and expensive for years to come.

Even in traditionalist organizations, however, there's usually a push to find efficiency wherever possible. Non-mission critical workloads, dev and test, as well as archival storage are all areas where administrators are more open to trying "new" technologies. This is where data fabrics—and more specifically the HCI+ approach—are gaining the most traction.

The Problem with Classic HCI

The problem with classic HCI solutions is that they're very restrictive regarding cluster composition. The ratio of storage capacity to number of CPU cores to RAM isn't particularly flexible. Additionally, individual nodes within the cluster can only have so many SSDs (and so many NVMe SSDs), limiting the performance of any given node's storage.

So long as you know exactly what you want to put on a given cluster before you buy it, you're probably good. You can specify the cluster's capacity and performance to meet those exact needs. In the real world, however, priorities change, new workloads are introduced and the workloads assigned to most clusters at their retirement look nothing like the original design.

A common reaction to this scenario is for administrators to overprovision their clusters. Overprovisioning is a time-honored tradition, but it sort of defeats the purpose of using HCI in the first place. HCI is usually sold on the basis that you can grow your cluster as needed, when needed, saving money as you go.

Unlike more open data fabrics, however, classic HCI solutions don't let you just add a pair of all-NVMe nodes for performance, or a great big box of spinning rust for capacity. This is where HCI-plus comes in.

Not only will an HCI-plus approach let you create HCI clusters with diverse node composition, it allows for taking all the old storage arrays kicked loose from the mission-critical tier, and giving them a second life by adding them to the non-mission-critical tier's data fabric.

Buzzword Bingo

Like it or not, nomenclature matters. Few (if any) vendors that bill themselves as HCI vendors take an HCI-plus approach. Vendors that do offer a more open data fabric than classic HCI haven't yet collectively settled on a means to describe what they do, in part because there's a lot of room for variation between classic HCI and a completely open data fabric.

The inability of vendors to settle on nomenclature has made educating customers about what vendors do—and how they differ from their competitors—rather more difficult than it really needs to be. This, in turn, makes it difficult for customers to find the best fit for their needs. In part, this is why the storage market is such an unruly mess.

What's important to bear in mind is that, while many HCI vendors bill themselves as "the new hotness," pointing to traditional storage arrays as being an archaic solution, HCI is already almost a decade old. It's not new, and in those 10 years many HCI vendors have been shown to be just as prone to lock-in, overcharging and needlessly restricting customers as the vendors they sought to displace.

Open data fabrics offer the same promise as the storage revolution we were promised a decade ago; they're what HCI should have been. It remains to be seen if existing HCI vendors will open up their products to at least offer an HCI-plus approach, or whether they'll fade away, ceding territory to a new generation of open data fabric vendors eager to define their own buzzwords, and make their own mark on IT history.

A BLUEPRINT FOR SMARTER STORAGE MANAGEMENT

Addressing the Top 3 Challenges Impacting Your Data with Software-Defined Storage

Significant IT challenges exist in the industry. Increasing pressure to maintain uninterrupted service, scale with increasing capacity demands, deploy high-performance architectures, and reduce cost, rank high among them.

The Core Issue: DATA

Whether it is complying with new regulations, deploying new applications, or dealing with rapid growth, your data is the core of all operational and service functions. Business critical data must always be accessible and immediately retrievable. Failure to do so could result in a significant loss of business.

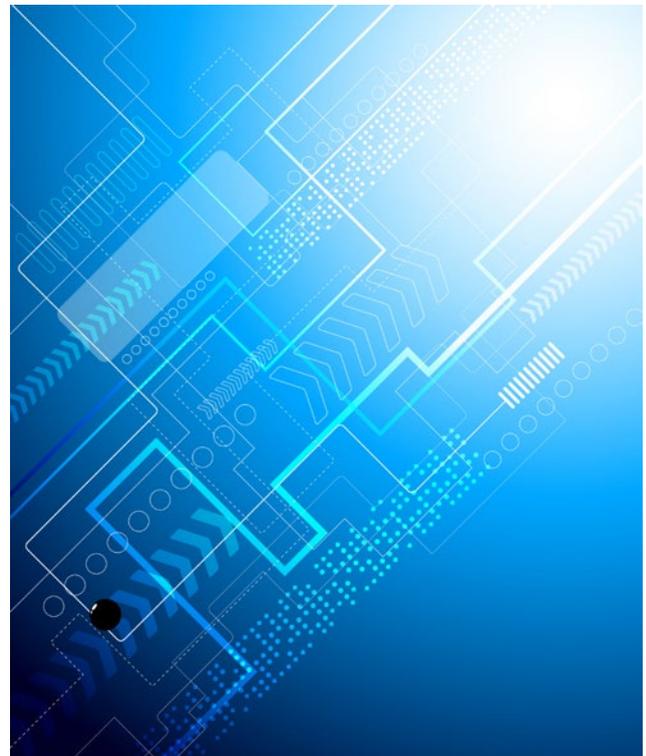
Additionally, current IT data retention policies make data storage and management a monumental task. IT storage systems need to scale like never before. And to top it all off, IT managers need to work within aggressive IT budgets.

Architecting a traditional storage solution that provides enterprise-grade data accessibility, scalability, and performance can be challenging enough. But you also need to be able to meet changes in these demands on a moment's notice. This agility would require significant financial and human capital to implement, not to mention, to maintain.

The DataCore Solution: Software-Defined Storage

DataCore™ SANsymphony™ software-defined storage is the most robust enterprise storage services platform on the market delivering uninterrupted data access, massive architectural scale-out, and game-changing application performance. The software forms a transparent virtualization layer across diverse storage systems to maximize the availability, scalability, and performance of all storage resources.

Let's explore how SANsymphony solves all of these challenges, simplifies storage management enterprise-wide, and gives you back control of your storage.



Challenge #1: Data Availability

Data availability is focused on maintaining data accessibility, even in the event of a catastrophic failure within the storage architecture due to hardware malfunctions, site failures, regional disasters, or user errors. SANsymphony accomplishes continuous data availability by providing features such as synchronous mirroring, continuous data protection (CDP), multipath communication channels, and remote site replication for disaster recovery operations. In addition, DataCore provides guidance on how to set up your storage infrastructure and processes to ensure you have a complete solution.

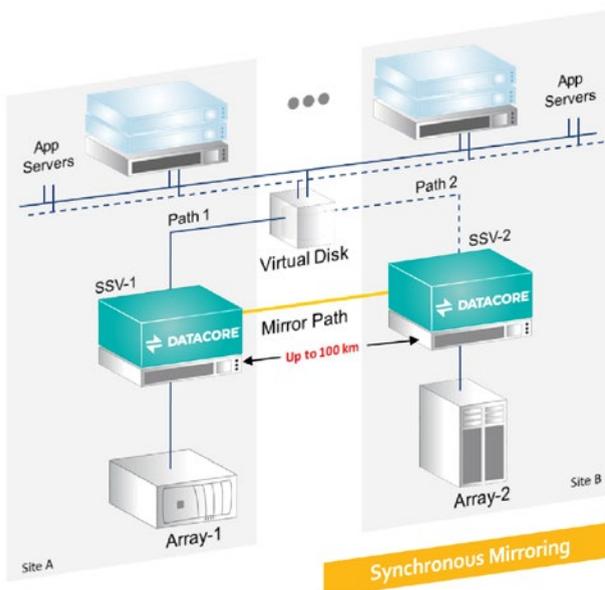
Synchronous Data Mirroring

The need to safeguard business-critical data is extremely important for your organization. Before we talk about how this is accomplished, let's review the distinction between fault-tolerance and high-availability.

Fault-tolerance provides protection against component-level hardware failures in a storage system such as hard disks or power supplies. All production IT systems should have reasonable fault-tolerance implemented to protect against the certainty of hardware failures. But if fault-tolerance is all that has been achieved, then the most important component has been overlooked: Your DATA!

SANsymphony provides protection against catastrophic storage system failures such as would occur with environmental failures (i.e. power outages, fire, loss of computer room air conditioning, etc.). Fault-tolerance measures cannot adequately protect against these. High availability provides data-level redundancy on top of the hardware component-level redundancy to maximize your data availability and keep applications running without interruption.

SANsymphony software-defined storage accomplishes high availability by synchronously mirroring every data block to a second fully active SANsymphony node. These nodes maintain redundant copies ensuring that even if one entire side of the storage infrastructure becomes inoperable, the applications will not experience any interruption, and the data is always available.



There are four principles to consider when deploying a highly available solution:

- › **End-To-End Redundancy:** The underlying storage systems need to have a reasonable amount of fault-tolerance (hardware component-level redundancy) combined with High Availability (data-level redundancy).
- › **Subsystem Autonomy:** The underlying storage systems need to work independently and are not aware of each other.
- › **Subsystem Separation:** The underlying systems can be separated up to 100 km apart to prevent an incident at one site from affecting the delivery of storage services at the other site.
- › **Subsystem Asymmetry:** The underlying storage systems on each side can vary in manufacturer and/or model to ensure that a bug or software/hardware issue at one site doesn't also impact the other.

Continuous Data Protection

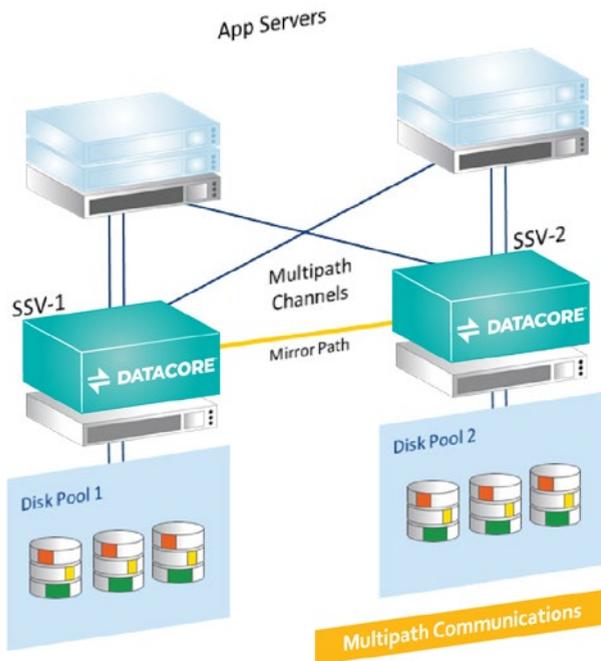
Another important consideration to safeguard business-critical information is to leverage Continuous Data Protection (CDP). The CDP feature provides up-to-the-second roll-back capability for any CDP enabled volume. This delivers the best possible recovery point objective (RPO) and recovery time objective (RTO) in the event of unintentional data loss due to malware or user error. CDP can also be used to roll back volumes to a previous point in time in the event that a backup is missed. CDP is very useful for those organizations under strict data retention and recovery compliance policies.

Multipath Communication Channels

In addition to ensuring that the data is safeguarded through synchronous mirroring and CDP, data availability is also concerned with providing continuous accessibility through multipath communication channels.

If a communications channel becomes unusable for any reason, the data should still be technically accessible. However, there is now a corresponding reduction in the total channel bandwidth used to access the data. With the high demands of today's application workloads, this could have a substantial impact on performance and user experience.

SANsymphony is a fully abstracted independent software-based storage virtualization solution. This means that you have complete control to adapt and modify your storage architecture to maintain service-levels under normal and abnormal operating conditions.



Remote Site Replication

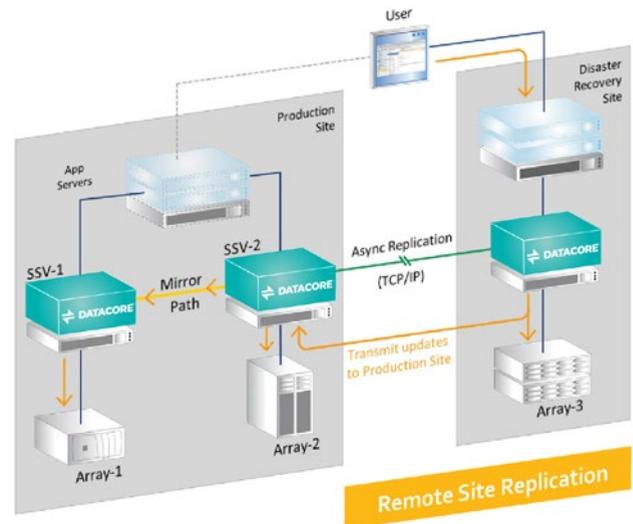
The last line of defense for data availability is remote site replication in the event a natural disaster impacts a large geographic region. Using backups to restore an entire site is unrealistic due to the time needed to recover, that is, if there is anything to recover to. SANsymphony asynchronous replication, combined with the Advanced Site Recovery (ASR) feature, turns a potential IT nightmare into a manageable and testable process. Together, these will make the difference between staying in business and going out of business.

With asynchronous replication, the data is copied to the remote site at the block-level and becomes quickly accessible should the production site become unavailable. Asynchronous replication provides the highest integrity and the most flexible disaster recovery solution for operating systems and business-critical applications and data.

A disaster recovery plan is incomplete without the ability to test the failover process and ensure data integrity at the remote site. With SANsymphony software-defined storage you can test site-to-site failover as often as necessary without impacting the production environment and without interrupting the remote site replication process. This is accomplished by enabling the test mode feature for the volumes you want to verify. Once the remote site

test is complete, disable test mode and SANsymphony will automatically return the remote site configuration back to normal operating mode.

There are many ways to replicate data to a remote site, that part is easy. But what if a failover and then a subsequent failback become necessary? This is much more difficult. ASR automatically reverses the asynchronous replication direction during a real failover or failback operation to ensure that the alternate site receives all updates once services are restored. Additionally, a full resynchronization of all the data is not necessary to bring the original site back online. This prevents unnecessary downtime during the transition.



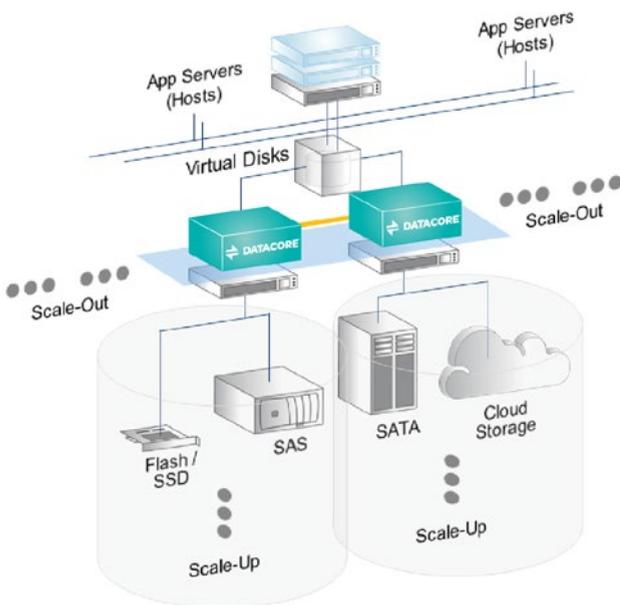
Challenge #2: Capacity Utilization and Scale

SANsymphony ensures your data is stored efficiently, maximizing utilization of storage resources. It accomplishes this through thin-provisioning which dynamically allocates space only when data is written. The result is an increase in storage utilization and capacity reclamation by as much as 4X. Thin-provisioning also eliminates logical volume resize operations since there is no penalty for presenting large logical volumes to application servers. And over time, SANsymphony can reclaim zeroed-out space that occurs as a result of removed or deleted data.

The DataCore solution also enables you to rapidly and safely scale your storage infrastructure up and out to stay ahead of changing storage demands. Whether it is adding

additional nodes or adding additional capacity to existing nodes, you can do this quickly and without interruption to the production environment.

There are different reasons for scaling out versus scaling up. Depending on the server hardware deployed as the SANsymphony node, you may choose to simply scale up by adding storage capacity (both internal to the server and external), high-speed cache, communication channels, processors, etc. If additional IOPS are needed and the existing hardware is at capacity, then additional nodes can be added to accommodate the application workloads. Either way, scaling additional resources can be done without taking your production systems offline.



Challenge #3: Storage Performance

IT storage systems require extreme performance to keep up with the rate of data acquisition and the increasing number of applications. SANsymphony provides cutting-edge performance through two primary methods: high-speed caching and automated storage tiering.

High-speed Caching

The first method is through the use of high-speed RAM as a cache. SANsymphony software leverages the node's processors, memory and I/O resources to execute

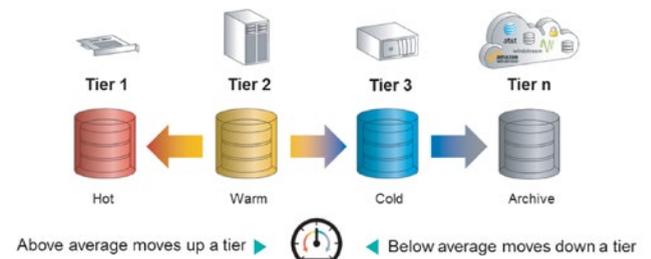
sophisticated multithreaded caching algorithms across all of the storage devices under management. The software can reserve up to 1 Terabyte (TB) of RAM per node forming SAN-wide mega-caches.

High-speed caching is critical for maintaining application performance since RAM is many orders of magnitude faster than the fastest flash technologies and it resides close to the CPU to minimize latency. It is the fastest storage component in the architecture, delivering up to 10X performance boost to applications and freeing up application servers to perform other tasks. It also extends the life of traditional magnetic storage components by minimizing the stress experienced from disk thrashing.

Automated Storage Tiering

The second feature responsible for storage acceleration is automated storage tiering. Automated storage tiering is designed to move data in real-time to higher performing storage tiers based on access frequency. When combined with high-speed storage devices such as flash, it yields markedly better overall application performance.

Automated storage tiering provides seamless integration of flash storage or other high-speed storage systems into the overall storage architecture. Simply add flash storage into the SANsymphony disk pool alongside existing magnetic disks and you're all set. SANsymphony will automatically move data back and forth across storage technologies at the sub-LUN block-level, yielding the best performance for the applications that require it.



A very compelling reason to use storage tiering is that you don't need to move to an all-flash architecture to get all-flash performance. Typically between 8-10% of the total storage pool should contain high-speed solid state devices. SANsymphony will ensure that the hot data (frequently

accessed) is resident on the fastest tier and the cool data (less frequently accessed) is resident on the slower, less costly tiers.

Industry statistics show that, on average, only 15-20% of a company's entire data footprint benefits from flash storage since the remaining percentage is considered largely dormant. According to an industry analyst report:

"The access patterns, value and usage characteristics of data stored within storage arrays varies widely, and is dependent on business cycles and organizational work processes. Because of this large variability, data stored within storage devices cannot be economically and efficiently stored on the same storage type, tier, format or media."

SANsymphony supports up to 15 different tiers of storage technology including cloud storage for infrequently accessed long-term archival data. As more advanced technologies become available, existing tiers can be modified as necessary and additional tiers can be integrated to further diversify the architecture.

The end result is superior application performance and lower TCO across the entire storage architecture.

Conclusion

IT is expected to maximize the return on every investment. But, there is also a great need for high-performance, reliable, enterprise-grade storage systems, which are expensive. With SANsymphony, you can have the best of both worlds.

SANsymphony is a fully abstracted independent software-based storage virtualization solution. It grants you complete control over your storage architecture while providing the most robust enterprise storage functionality available today.

Consider the impact on ROI and TCO with a purpose built solution that allows you to leverage any make and model of storage system, all while providing blazing performance and the most advanced forms of data protection to safeguard your data. And as conditions change over time, expand without interruption the resources needed to fulfill the demands of the new requirement.

THAT'S THE POWER OF DATACORE!

Trevor Pott is a renowned technologist with 20-plus years of experience serving organizations of varying sizes as an IT practitioner, adviser, consultant and writer. A prolific tech writer and highly sought-after vendor consultant who knows both technology and the people who apply it, Trevor is also a systems administrator with eGeek Consulting Ltd.

Check out DataCore today at datacore.com.

