



Lunch Break

S E R I E S

Storage Explodes!

With the number of storage solutions growing all the time, find out which is right for you and why the storage “revolution” might not live up to the hype


The Storage Explosion Is Here..... 1

The Software-Defined Storage ‘Revolution’ .. 4

SPONSORED BY

Redmond
MAGAZINE

VEEAM
IT JUST WORKS!™



SHUTTERSTOCK.COM

The Storage Explosion Is Here

There's a dizzying selection of storage solutions out there. Which one is right for you? The answer will often be "more than one."

Suppliers of storage hardware and software are presenting what appears to be a huge list of options. Which ones are best isn't always clear. Furthermore, it's not clear if there's a single option that addresses an enterprise's needs. All of the suppliers, however, promise that their solution is the best, the most

cost-effective and makes the best possible use of <insert the name of your favorite storage technology here>.

Markecture Abounds

Like just about every other area of the IT market, suppliers of storage are always doing their level best to one-up their competitors, seeking

ways to out-gun the others in the areas of storage performance, reliability, scalability and overall cost.

Although there's a great deal of noise in the market, a few things are clear. For one thing, each of the suppliers believes it and it alone is uniquely qualified to be the only source of storage technology. Also clear is the fact that there are a



number of different types of technology from which enterprises may choose, and an overwhelming set of combinations for how this technology may be used together. Finding the right fit largely depends on what question the enterprise is asking.

Increasingly, storage services are being offered by suppliers of managed services, colocation or cloud services.

Why So Confusing?

There are a number of different approaches to storing applications and data, and each is useful in the right place and at the right time. Some approaches require the storage media be directly connected to clients and servers, while others attach storage devices to a storage appliance or a storage server.

If the storage device is directly connected to the client or server, there are many different storage interconnects from which to choose. Each offers a different mix of price and performance, and can control what media options are available.

As with direct-connect approaches, there are several different storage interconnects in use in today's datacenters when the storage devices are attached to a storage server or appliance. These servers may be connected to computing systems using a general-purpose LAN or special-purpose SAN.

The industry is also seeing the increasing use of system memory being used as a special form of storage for computationally intensive, extremely high-performance applications. Sometimes suppliers call this "distributed cache" or "in-memory database."

To add to the confusion, cloud services providers have begun to offer an array of new Storage-as-a-Service products. They're trying to convince enterprises that it's better, less

complex and less costly to use those services rather than purchase, install and operate their own storage.

Different Technology for Different Needs

The industry has used various types of technology over the years, including:

- **Tape.** Different suppliers have offered paper tape, cassette tape, and reel-to-reel tape products. Several suppliers have offered direct access tape devices that could replace rotating media for large-scale storage

There are a number of different approaches to storing applications and data, and each is useful in the right place and at the right time.

applications.

- **Rotating media.** Different suppliers have offered rotating drums and a whole herd of different types of disk storage. While most of these were based on magnetic recording, some were based on optical recording technology.

- **Solid state.** Although solid-state storage has been available for decades, and its access times and throughput made it extremely desirable, the cost was prohibitive for most applications. Recently, however, the introduction of new technologies has resulted in the rapid adoption of flash memory.

As suppliers seek ways to offer flexible and inexpensive storage options, the market is seeing the emergence of distributed cache solutions using the system memory of low-cost, industry-standard servers, blades or distributed NoSQL database solutions using server clusters.

You've Got Options

Like most areas of IT, there are many different types of storage technology, and each has the ability to serve a different set of needs. If the enterprise carefully reviews its application portfolio, it will soon become clear that each application has a different storage profile.

Some applications require the storage and retrieval of huge amounts of data, and longer access times are acceptable. Other applications access huge amounts of data, but the access

time must be kept to a minimum. Still, others require immediate access to data and any delay is unacceptable. Finding the right solution necessitates understanding your environment's unique requirements. Most fall into one of these categories:

- **Long-term storage.** The requirements for this type of storage usually include massive capacity and low cost per megabyte or gigabyte. Applications using this data typically are batch or analytical jobs.

- **Medium-term storage.** The requirements here lean more toward finding a good balance between performance and cost. This often means storing applications and data for remote or VDI desktops, servers or even handheld applications. Enterprises are often willing to compromise on storage performance to reduce overall cost.

- **Short-term storage.** The requirements for transactional or



business intelligence applications often include the need for very short access time and medium levels of throughput. Enterprises are often willing to compromise on cost to obtain performance.

- **Storage for high performance or technical processing.** The requirements for this type of workload often include extreme storage performance (seek performance and access time or latency), extreme needs for data throughput, and huge volumes of data. Shared cache, clustered NoSQL databases and in-memory databases are often used to address these requirements.

- **Flexible storage.** When the enterprise faces ever-changing, dynamic requirements, it will often turn to some form of distributed or hybrid storage. A local cache made of high-performance storage is deployed to improve the performance of off-site or cloud storage. The off-site storage may be at another enterprise-owned site, at a site managed by a managed services supplier or in the datacenter of a cloud services provider.

How Are Suppliers Addressing These Requirements?

Although each supplier is addressing these storage requirements differently, there are some common threads:

- **Slow, but reliable storage for huge amounts of data.** A form of tape or optical technology often satisfies these requirements. Some cloud services providers are suggesting their Storage-as-a-Service offerings might be a replacement for this type of storage. It's not clear what type of storage they're actually using to address this need. Low-cost, low-performance rotating disk storage is very likely part of the cloud services provider's offering.

- **Fairly fast and inexpensive disks.** These can be used for client-side applications or server-side applications for small to midsize businesses.

- **High-speed, expensive disks.** These target server-side applications that need both a larger amount of storage and high levels of performance.

- **Flash and other forms of solid-state storage.** They're packaged as storage devices that address the needs of applications requiring very low access times or high levels of throughput. Typically

Low-cost, low-performance rotating disk storage is very likely part of the cloud services provider's offering.

these devices are much faster and more expensive than traditional disks, and offer less capacity.

- **Internal solid-state memory.** It's packaged by storage virtualization technology so that it appears to be a storage device.

What's the Best for Me?

The enterprise must take the time to survey its portfolio of workloads to learn the answers to the following questions:

- How much is the enterprise willing to pay for storage? High-performance storage typically is expensive.
- Does the enterprise really need massive storage capacity? There are many ways to address this type of need. The best answer usually is a compromise between cost, performance and storage capacity.
- Is the enterprise willing to use off-site storage? Enterprises in

regulated environments may only be able to use on-site, locally controlled storage for regulated applications. Collaborative applications, e-mail and other non-regulated applications might be candidates for off-site cloud storage.

The Golden IT Rule

Most enterprises, by necessity, rely upon many different types of storage. This is partially due to the different needs of each application, and also due to the application's age. Older applications, for instance, are likely

to be using older types of technology. Enterprises often follow the golden rule of IT, "If it's not broken, don't fix it," when dealing with these applications.

The more enterprise decision makers know about their applications and their requirements, the easier it is for them to select the right storage technology, storage location and determine if cloud storage is even a reasonable option.

Daniel Kusnetzky, a reformed software engineer and product manager, founded Kusnetzky Group LLC in 2006. He's literally written the book on virtualization and often comments on cloud computing, mobility and systems software. He has been a business unit manager at a hardware company and head of corporate marketing and strategy at a software company. In his spare time, he's also the managing partner of Lux Sonus LLC, an investment firm.



The Software-Defined Storage ‘Revolution’

They’re not always what they’re cracked up to be.

The press and analyst community have been amping up the rhetoric for the past few years, assigning the term “revolutionary” to just about every technology introduced (or in many cases resurrected) into the market, and making a lot of IT folk concerned that they may be missing out on an important trend. A lot of folks have reached the saturation point -- the place where we immediately doubt the credulity of any such claims and the integrity of those who make them.

I happen to combine a long career in IT with a couple of degrees in political science and international relations -- a mix that provides, perhaps, a hybrid perception of both technology and revolution. From my perch, it seems that, in both the fields of contemporary politics and contemporary technology, the propagandists -- er, marketing folks -- have taken control of the dialog. Whatever merits there might be in the case for revolutionary change, they often get diluted, distorted or perverted by the marketing hype -- i.e., the propaganda -- around the effort.

The Sound of Inevitability

True revolutions are inevitable. Revolutionary thinkers will tell you that conflict reaches a point where issues can no longer be resolved through conventional institutions or processes, and a straw is finally introduced that breaks the camel’s back. Conflict ensues,



SHUTTERSTOCK.COM

according to the theory, resulting many times in the triumph of “reactionary forces” -- that is, the existing order prevails. Sometimes, very rarely, revolutionaries win the day and become the new order.

The thing about revolutions is that you can’t make them happen. They just do. They happen as a result of inevitable and immutable forces that cannot be directed or diverted or contained. They happen because they have to.

Usually, revolutions occur when the price, the cost, the downside of revolution doesn’t seem as terrifying or insufferable as the continuation of the status quo. Ideally, the revolution promises better than the current state of affairs, better outcomes, and meaningful improvements in the way things are. Unfortunately, most 20th- (and 21st-) Century revolutions have been characterized not by the advancement or progress of organiza-



tions or groups toward a better outcome; instead, the so-called revolution has been cover for a changing of the guard, the shift of power from one group of corrupt so-and-so's to another.

SANs: The Revolution that Wasn't

So it is with most technology revolutions. The storage area network (SAN) was supposed to bring about a kind of nirvana in which all storage vendor gear participated in a common network infrastructure and a common management scheme designed to bring new value and order to the IT universe. That was the vision of the Enterprise Network Storage Architecture (ENSA) that came out of Digital Equipment Corporation, via Compaq Computer Corp., in the 1990s. ENSA was supposed to end the old storage model -- the hegemony of monolithic storage arrays -- that caused storage infrastructure to be so costly and so difficult to administer with any sort of efficiency.

Only the revolutionaries who dreamed that up were squelched by Compaq (and later HP) management. They were afraid that gutting the proprietary differentiators in their gear and providing a mechanism for common interoperability and manageability would enable the Chinese to come into the US market with their monolithic arrays, which would be loaded with proprietary value-add software features, and clean our proverbial clocks.

Managers vs. Innovators

The difference between the managers and the innovators was that the former lacked the faith in the consumer articulated by the latter. Consumers were simply not sufficiently aggravated by the cost and inefficiency of monolithic storage to

actually change their infrastructure model. The time wasn't ripe for revolution. Another key difference was that the managers held the purse strings, which in politics or technology has a tendency to shape outcomes. So, ENSA went nowhere.

In the end, HP tried to disappear it the way that certain Banana Republic dictators disappear their opposition following an election. Instead of ENSA, we got SANs. Storage Area Networks weren't networks at all, only a bunch of monolithic arrays with simple

The difference between the managers and the innovators was that the former lacked the faith in the consumer articulated by the latter.

physical layer attachment plumbing and protocols -- Fibre Channel.

Sticking with storage technology, as the outcome of the ENSA revolution was becoming evident (the reactionaries won), another revolutionary surge was shaping up between the traditionalists and the advocates of revolutionary change in the form of virtualization.

We saw this movement first in the storage world, with several upstart vendors appearing in the market at about the same time with different strategies for aggregating storage capacity and storage services from heterogeneous storage arrays, then serving as a software-based uber-controller that could divvy out storage to any app that needed it from shared pools (sort of an ENSA at the software level). DataCore Software continues to fight this fight, and IBM is also dusting off its SAN Volume Controller kit to deliver similar functionality.

The Virtualization Revolution

This revolution, however, failed to gain momentum at the time it was introduced, possibly because consumers were too busy trying to digest and make sense of SANs that weren't really SANs. Meanwhile, a similar conception of virtualization did become a meme in the server community, where the hardware components of competing server gear from different vendors were just as identical as the hardware components of storage kits from different vendors. When hardware becomes commoditized, virtualization

advocates argued, it was time for revolutionary change.

Virtualization of workload wasn't anything new, of course. Mainframes had been doing it since the late 1970s. But most IT operators hadn't worked in DP (data processing, the previous moniker for the activity) and didn't know what a mainframe was, so it all seemed new. Good propaganda convinced everyone that instantiating applications and operating systems as virtual machines (VMs) atop commodity hardware was the next big thing, the revolution that would drive cost and complexity out of client-server computing. Adoption was encouraged by capabilities for supporting multi-tenant computing added by Intel to its CPU chips and by an economic disaster that forced firms to use any strategy they could find to bend the CAPEX cost curve in IT.

This revolutionary zeal around



server virtualization put new pressure on storage, of course. Aggregating VMs onto fewer servers changed traffic patterns on networks, fabrics and storage. Hypervisor vendors, the revolutionary leaders in the virtual IT infrastructure, found storage to be an easy target to blame for all that ailed their programs and strategies: Applications were slow, blame legacy storage. I/O was randomized, slowing reads and writes, blame legacy storage. IT costs had not decreased, but rather increased with virtualization, blame legacy storage. Clearly, evil legacy storage vendors were the reactionary forces that needed to be brought into line with the new order. Rip and replace became the order of the day.

Software-Defined Storage Takes the Stage

Out with the old. In with the new. Enter software-defined storage. SDS was to storage what server virtualization was to application hosting, according to hypervisor vendors. It was a way to replace expensive, complex, hard-to-manage commodity infrastructure with something more elegant, simpler, and much more automated. Perfect for those shops that didn't have rocket scientists on staff to administer the storage resource or manage its operation, and better suited to handle the new I/O demands imparted by revolutionary VMs.

It sounded great to those who lacked the skills and knowledge to measure or understand that application performance issues rarely had anything to do with legacy storage, or that I/O logjams were the result of hypervisor computing itself.

The hypervisor vendors gave firms someone to blame for their own inefficiency, and in SDS (requiring an

expensive overhaul of storage infrastructure) they offered a solution. Like contemporary revolutionary leaders, they demanded a little more sacrifice in order to realize the IT utopia.

I Just Can't Wait To Be King

Interestingly, before it even appeared as a storage model, SDS had been hijacked by the hypervisor vendors. Not surprisingly, each promoted their

From where I'm sitting, there's nothing revolutionary about the current crop of IT revolutions.

own flavor of SDS infrastructure in an attempt to ensure that their solution couldn't be shared by data from VMs created by rival hypervisors. One lesson that the server virtualization folks had learned from their historical precursor -- 1970s IBM -- was that it was good to be the king.

The SDS models advanced by the hypervisor vendors often reflected a lack of understanding of storage itself, of the impact of random I/O from multiple VMs all sending I/O down a common pipe, of cost-efficient strategies for replicating data between storage nodes, or even of the right way to use Flash memory to slow wear rates.

They especially eschewed storage virtualization as part of the SDS model, treating it like a bastard child they didn't wish to acknowledge. Such a technology would enable the development of a common storage resource pool, preserving investments in "evil legacy storage," that could be shared between competing hypervi-

sor vendors and with non-virtualized workloads, too. That didn't pass the revolutionary litmus test, it seemed.

A Revolutionary Concept: Use What Works

This brings the story up to date. From where I'm sitting, there's nothing revolutionary about the current crop of IT revolutions. Counterrevolutionary that I am, it seems to me that the smart choice is

to deploy whatever technology works to meet workload requirements in a manageable way, just as we've been doing all along. Not jumping on the bandwagon of every "new and improved" technology or "revolutionary meme" doesn't make you less competent or time-bound in your thinking or uncool. It makes you smart.

And that's what we need most of all in IT today.

*Jon Toigo is a 30-year veteran of IT, and the Managing Partner of Toigo Partners International, an IT industry watchdog and consumer advocacy. He is also the chairman of the Data Management Institute, which focuses on the development of data management as a professional discipline. Toigo has written 15 books on business and IT and published more than 3,000 articles in the technology trade press. He is currently working on several book projects, including *The Infrastruggle* (for which this blog is named) which he is developing as a book.*

COMING SOON



Veeam Availability Suite v9

Availability for the
Modern Data Center



Learn more and preview
the upcoming v9 release

vee.am/v9