

Can your data storage ecosystem meet your virtualization needs ?

Everyone knows a chain is only as strong as it's weakest link. We are learning that when it comes to virtualized environments, the position of the weakest link in the data storage chain is shifting.

Disk spindles and interconnecting physical links are no longer the "most likely candidates to slow data down". That honor now goes to the translations between the virtual and the physical layers.

Where Virtual Meets Physical

Physical data storage has to exist somewhere in the chain; it can be as close as local disk on the same backplane, or, as far away (electrically speaking) as one of the spindles in a RAID set behind a storage gateway presenting virtualized volumes across a switched network.

Sure local storage is an option, but we just spent the last decade abstracting our data from our servers because of the benefits. You can see where this is going; pretty much everything between the bits and the pixels is going to be virtualized.

In a nutshell, every word that we read and every image or video that we watch will be delivered through a series of translations. To paraphrase a recent observation, "we are talking about a series of pipes" and increasingly, a lot of "twists and turns in those pipes."

Consider the simplest and most direct method of transmitting data, the good old copper wire; no caching, no queuing, no latency, and above all, no translation.

Next, think back to how disk drives were first addressed; the operating system talking directly to the disc using SCSI; one translation with immediate local feedback.

Then came RAID sets (lots of bits) and the challenges of delivering screenfulls of information to multiple users (lots of pixels).

The only way to meet these challenges was with a tightly coupled "ecosystem" wherein every aspect was developed, delivered, and supported by a single vendor named either IBM, or DEC, or HP.

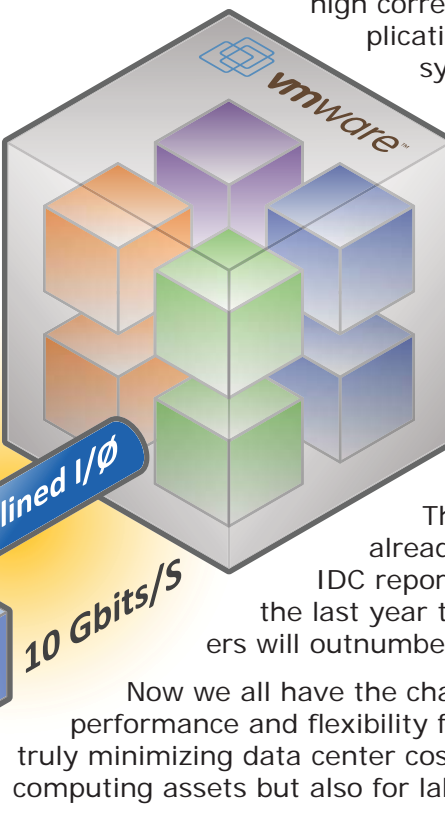
The resulting "information silos" seriously impeded the creative and productive process. Enter distributed computing with - distributed storage. This in turn created lots of opportunities for independence and productivity; usually in standalone applications with progressively more sophisticated GUIs.

Eventually those same GUIs methods were applied to the shared information repositories and it was possible to have both the personal productivity of distributed architectures, and some of the benefits of centralization by moving servers to a data center using shared SAN and NAS storage.

Notice I said "some". This is because although the hardware at the data center is now being managed by a relatively smaller team in a more secure environment - there is still a high correlation between applications and dedicated systems.

The next step in fully realizing the benefits of centralization is to make hardware assets multi-tenant so that they are supporting many applications on the same platform.

This process is already underway with IDC reporting that 2009 is the last year that physical servers will outnumber virtual machines.



Now we all have the chance to optimize performance and flexibility for our users while truly minimizing data center costs; not just for computing assets but also for labor and energy.

That is if we don't make the same mistakes that we made this last go around.

Creating Truly Open Systems

Specifically, despite the clear benefits of the open systems architecture(s) now fully achievable in a virtualized environment, we still tend to acquire and deploy "ecosystem" data storage solutions.

IDC reports that 2009 is the last year in which physical servers will outnumber virtual machines with the cross-over coming at around 6.3 of each type

This is understandable because the industry heavyweights, particularly in storage, pull us in by touting their flexibility on the one hand, all the while “inferring” that the shortest path to success is, of course, their single vendor solution.

This is an all too familiar cycle that compute and storage vendors put each successive generation of their customers through. There is even an acronym to describe this method, FUD, or casting fear, uncertainty, and doubt on open systems.

To date these seller and buyer behaviors have simply produced overspending for underutilized assets.

Moving forward in this virtualized world however, the challenges are much more complex and multifaceted; with a corresponding decrease in the margin of error for sub-par investments.

Budgets are shrinking not only for the usual suspects like physical assets, labor and energy, but also for “starting over” in the event of a mistake.

Worse yet, many of the under-utilized legacy computing and storage assets already in the field are inadequate to support full virtualization. In the case of SATA drives, they are downright unreliable.

Sadly, in many enterprises, these assets are embedded in one of the aforementioned “ecosystems” which compounds the problem:

- For starters, existing data storage installations are cash cows for support and maintenance so an ecosystem vendor will tend to “roll out” enhancements to meet their own internal fiscal revenue cycles.
- Another critical cornerstone in perpetuating storage ecosystems is carrying forward legacy data paths and software layers; many of which are past their prime and becoming bottlenecks for virtualization.
- Finally, every generation ushers in a new “must have” technology designed in collusion with the marketing department to lock-in customers. De-duplication is the latest such feature; offering dubious benefits, and in the case of virtualization, yet another layer of abstraction to deal with.

The key point here is that the major hypervisors are “ecosystems” unto themselves and we will need to guard against overlapping functions that waste compute cycles and deteriorate data storage access feeds and speeds.

RELDATA

What Your Hypervisor Wants

VMWare, Hyper-V, and Xenserver use sophisticated virtual networks to route information within themselves and to interface with the external world.

This virtual network is the key to abstracting computing and storage assets so they can be provisioned on an as needed basis. It is important to understand that the internal virtualization adds an additional layer of switching and translation on top of those already present on the storage fabric.

These added layers put a premium on ensuring the storage path is as clean and direct as possible.

RELDATA defines this process as “Pipelined I/O”, or IØPS™ which represents “near-zero” latency.

When you look inside VMWare for example you can see (as shown in the graphic below) that ESX wants to talk directly to a LUN or file system.



Each of these storage types plays a specific role in the virtualized realm and each of them benefits from the RELDATA IØPS™ architecture as follows:

- For block based LUNs, IØPS™ mitigates the latency and translation delays associated with typical storage ecosystems by reading and writing data using the fewest possible translations between the disk and the virtual interface.
This benefit is compounded as the number of concurrent multiple virtual machines rises.
- For the Network File System datastores, IØPS™ means that native NFS is provisioned to volumes within the same unified frame and does NOT traverse additional external interfaces.

Now that virtualization is here, it is a good time to revisit some of our assumptions about data storage because methods that were good enough in the past may not be up to these new challenges.

2009 RELDATA, Inc. All rights reserved.

All trademarks and product names are the property of their respective owners.