Realizing the Agile Data Center Vision – Next Step: Hardware-Defined Storage

A Perspective by Jon Toigo

Managing Principal Partner,
Toigo Partners International

Chairman, Data Management Institute

TABLE OF CONTENTS

Synopsis

As a term, hardware-defined storage may have a "retro" sound to it, especially in the current technical lexicon that seems to relish software-defined everything. However, a further examination of the concept provides an important understanding of the foundations of current software-defined storage and the problems it seeks to address, while at the same time setting the stage for storage management requirements that have as yet been unaddressed by software-defined storage architecture.

Introduction

Today, the quest is on for the creation of a "virtualized infrastructure" that can be orchestrated and allocated "on a dime" to provide business end users with the data and application capabilities they need to do more productive work. The problem is that, while we hear a lot about the underlying technologies for enabling agility – such as server virtualization, hypervisor computing, and the like, as well as software defined networks, and even software defined storage – these are all actually just abstractions of the physical kit.

Today, the quest is on for the creation of a "virtualized infrastructure" that can be orchestrated and allocated "on a dime."

Truth be told, it is physical storage arrays, physical networks and physical servers actually host, process and distribute data. Virtualization wares are software, but the hardware beneath the software layer remains extremely important. Yet, for all the discussion of software-defined data centers, management of hardware and physical interconnects – always a challenge in contemporary computing – is rarely discussed.

It is also worth noting that in most discussions of agility and agile data centers, much more ink in the trade press dedicated to talking about

Server administrators are being rebranded as "virtualization administrators" in many shops.

ways that new virtualization technologies (when fully baked) will improve the efficiency and resiliency of data center processes. There is nothing wrong with such evangelism, of course, but we are seeing virtually no time dedicated to discussing the risks associated with deploying new agile technology or with operating new agile processes themselves.

Capping off the top concerns associated with the current craze around software-defined technology is the "dumbing down" it is fostering in terms of the practical skills and knowledge of IT practitioners especially with respect to the physical universe of computing. Server administrators are being rebranded as "virtualization administrators" in many shops, tasked to handle the configuration and maintenance of a software layer that provides the virtual to physical connection. That's important, of course, but myopic attention to the management of a proprietary hypervisor software package doesn't necessarily equip admins with what they need to know to configure and maintain the hardware infrastructure – the physical universe of things that must work together perfectly to carry an estimated 4 trillion transactions per day that traverse network cables, server backplanes and storage interconnects in medium-to-large data centers.

The simple fact is that virtualization doesn't fix any problems in the physical world – any more than plastic surgery changes genetics in human beings. Virtualization just masks physical infrastructure from view, sometimes making that infrastructure more difficult to monitor and manage.

So, the emphasis on virtualization and software-defined to the exclusion of any attention to the real world of hardware – disk drives, fans, cables, power, HVAC – can and has introduced a significant risk of failure to our quest for agility. Addressing this risk will require a broadening of the discussion of software-defined data centers to include practical physical layer concerns.

The Challenge to Agility of Software-Defined Technology

Recent survey data suggests that the narrow focus on virtualization technology as a building block for software-defined data centers has already taken a toll on the delivery of promised agile service levels. Failure to cope effectively with the demands of virtualization on storage infrastructure, for example, is cited as a leading cause for the “stalling out” of sever virtualization initiatives.

One survey of about 475 companies last year found 339 responses citing “storage expense” as a significant obstacle to realizing value from server virtualization initiatives. 308 respondents said that storage

Recent survey data suggests that the narrow focus on virtualization technology as a building block for software-defined data centers has already taken a toll on the delivery of promised agile service levels.

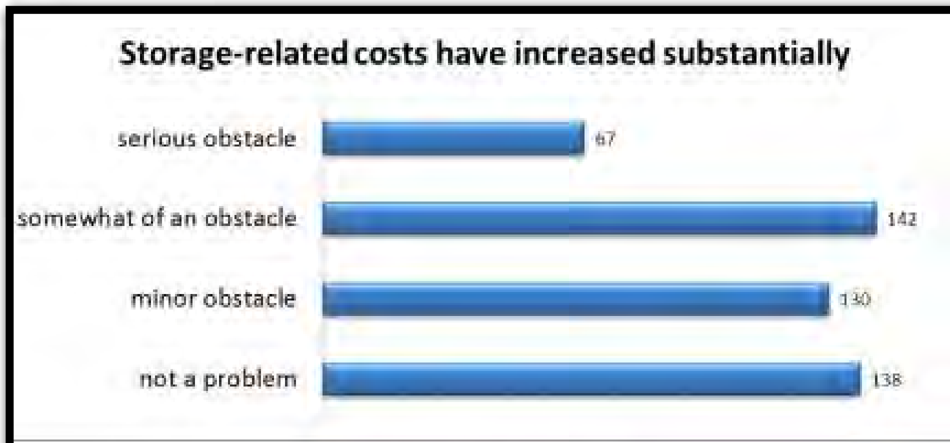

Figure 6- Impact of storage-related costs

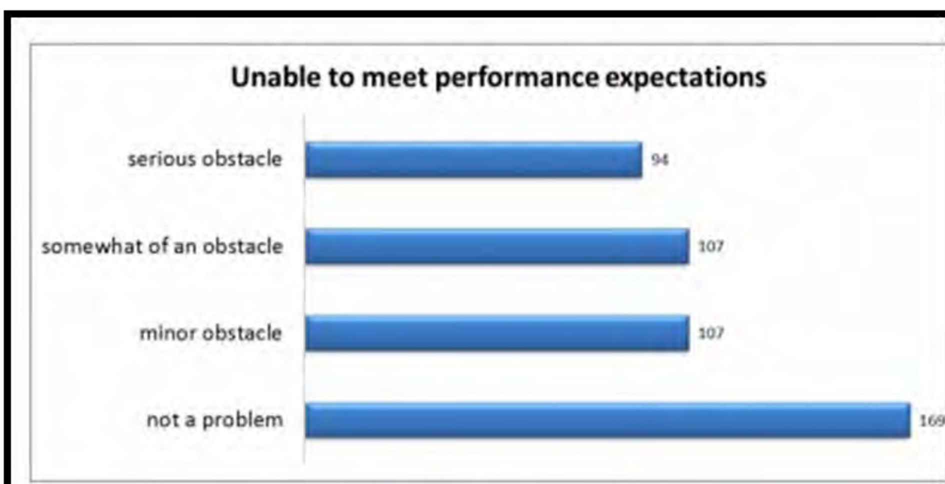

Figure 7- Performance considerations

Source: DataCore Software Third Annual State of Virtualization Survey
http://pages.datacore.com/StateofvirtualisationSurvey.html

inefficiencies were limiting their ability to realize the performance levels promised by their hypervisor vendors.

Recently, server hypervisor vendors have advanced a solution to the issue of storage provisioning delays that requires customers to "rip and replace" the storage we have behind virtual servers. The argument, which was not advanced during discussions of the original value case for server virtualization holds that, with the adoption of virtualization, comes the need to abandon "legacy storage" infrastructure.

In fact, the same vendors that sold companies storage area networks (SANs) only two or three years ago are using the term, legacy SAN, to describe those investments today – and often without so much as a blink. The new architectural model required by server virtualization, they say, is a Virtual SAN.

Virtual SAN is really a return to direct-attached storage, often with some mixed internal server storage thrown in.

Digging into the architectural model, where such a model has been formalized and articulated clearly, Virtual SAN is really a return to direct-attached storage, often with some mixed internal server storage thrown in. It could be argued that Virtual SANs are not an evolution of storage architecture, but instead a de-evolution – a turn of the evolutionary path back toward server-side or direct attached storage and isolated islands of storage technology rather than improved storage capacity and access sharing.

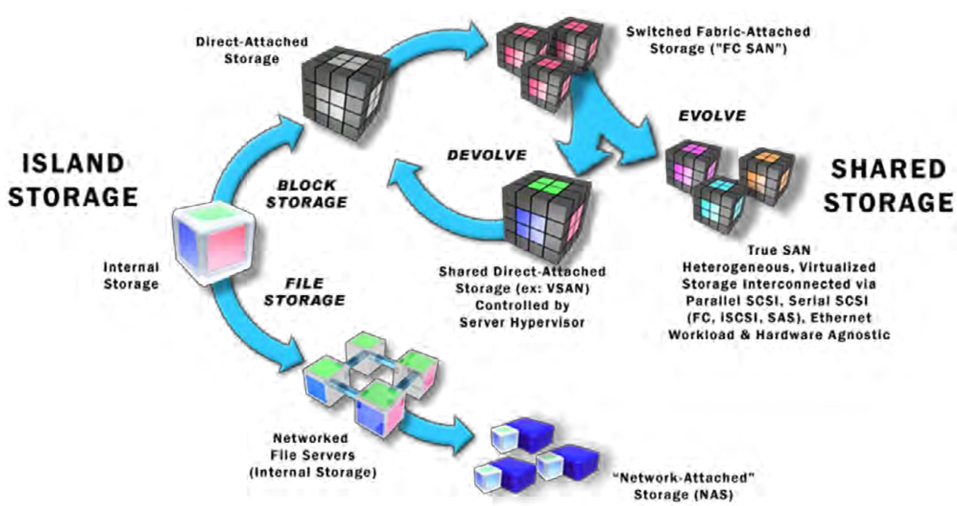

Hypervisor vendors are proposing that customers deconstruct their legacy storage and move to clustered virtual hardware/ software stacks.

There are incentives to return to direct attached, of course. Fabric interconnects are challenging to set up, manage and maintain. Moreover, storage shares – allocations of storage to workload are "hard wired" in anticipation of static hosting (e.g., the application never changes hosting locations under normal conditions, so a fairly permanent route is typically established between the workload and its data storage). With server virtualization, the much touted ability to shift virtualized workload from one physical host to another for purposes of load balancing or "high availability" comes at a significant price: when workload shifts hosts, administrators must intervene and modify application configuration parameters to re-establish routes to workload data, which is still located on hardwired storage locations.

To meet this challenge, hypervisor vendors are proposing that customers deconstruct their legacy storage and move to clustered virtual hardware/ software stacks in which storage is direct-attached to each node and workload data must be mirrored between all nodes in the cluster – and between clusters that might be potential hosts for the workload. The impact is to drive up capacity demand and cost. This is why leading analysts have moved off of their 2011 estimates of storage capacity demand growth (20 – 40% growth per year through 2016) and have begun talking about growth rates of 300% to 650% per year in highly virtualized shops.

Moreover, adopting server-side architectures means that multiple data replication processes will need to be supported by infrastructure and managed by administrators. Mirroring, if controlled by server hypervisors, can be problematic owing to poor I/O handling capabilities even in market leading hypervisor software. Offloading mirroring to the controller on each server-side array might provide an alternative, but invites vendor lock-in (most vendors will mirror data only to kit that has their own moniker on the bezel plate) and cost. Moreover, mirroring consumes bandwidth of either networks or storage interconnects, with potential impact on the overall performance of storage and of the applications using the data.

And all of the mirroring will require administration. While it could be argued that server-side storage architecture may be a better fit with the lesser degree of storage skills and knowledge of the virtual server administrator (they often lack understanding of FC, SAS or iSCSI SANs, or even traditional array controller functionality such as RAID, not to mention value add services sold on array controllers to increase the price of the kit), the truth is that administering server- side storage architectures and myriad mirroring processes might prove even more problematic.

Moreover, if current trends hold, server side architectures proposed by server hypervisor vendors will, in part, lock the consumer into the hypervisor vendor 's proprietary functional stack while locking out the technologies of competitors. That means a hypervisor-storage stack from vendor A will only be able to be used with vendor A's workload, creating an "island" of storage. Given the likelihood that larger data centers will have a mix of hypervisors supporting different types of workload, plus some applications running without a hypervisor at all, the very real prospect exists that software-defined data centers will present a nightmare scenario of several storage infrastructures to manage.

If current trends hold, server side architectures proposed by server hypervisor vendors will, in part, lock the consumer into the hypervisor vendor's proprietary functional stack while locking out the technologies of competitors.

The hypervisor vendor's thinking is clear, of course: Younger "virtualization administrators" probably do not remember the problems of data isolation or the challenges of sharing data between different islands of storage locked up behind different servers. These problems had reached epic proportions by the late 1990s, helping to create a push to build a more centralized storage infrastructure – something vendors called a Storage Area Network – to enable greater data sharing.

Of course, these SANs (FC fabrics, in fact) weren't perfect, or even true networks by standard definition of networks as embraced by IEEE, IETF, ISO, and the other standards groups. An important missing component from the SAN protocol stack was a dedicated management layer, present in general network definitions for nearly 30 years. Fibre Channel protocols didn't provide in-band management, and that fact called into question vendor representations of FC SANs as true networks. Still, Fibre Channel provided a fairly robust fabric, a unified interconnect leveraging serialized SCSI that enabled large numbers of storage devices to be aggregated into a common infrastructure.

Both capacity and performance must be provisioned elastically.

Server-side direct-attached storage is being posited as a more scalable and flexible storage architecture that eliminates the problem having to perform time consuming and complicated gyrations to revise the routes through the storage fabric between workload and data when workload moves to different servers. However, for that benefit to be realized, considerable resource management and orchestration technology is required – management that is both hardware and hypervisor agnostic.

The management challenge goes beyond data replication and mirroring. Storage must also be allocated in an elastic way. Administrators need a simple way to tune storage resources allocated to a workload so that it is not starved of storage capacity or performance required even under peak load — or wastefully overprovisioned with too many resources.

Both capacity and performance must be provisioned elastically. Depending on the application workload, performance can be as critical to elasticity as capacity. Certainly this is a concern of many IT planners who are considering server-side flash and disk storage architecture to support performance-sensitive applications like virtual desktop infrastructure (VDI). Anyone working on a desktop virtualization solution has likely encountered (or read about) problems like "boot storms" (the sudden request for virtual desktops by many users at the same time — say at the beginning of a shift). Such demand puts the elasticity of storage solution to the test, to be sure.

Another performance challenge may have to do with how well server-side storage architectures cope with random I/O processing, a common issue in a virtualized desktop environment. In the experience of many firms, VDI challenges all storage infrastructure types with random writes. Both memory caching and storage itself need to be engineered specifically to support random read/write workload – whether using solid state components (DRAM or flash memory) or the disk components or a hybrid of both. None of the leading hypervisor vendors currently provide the level of sophistication or configuration control that will enable virtualization administrators to obtain the performance levels needed by demanding applications. This issue does not go away because vendors speak about elasticity only in terms of capacity management.

Achieving resiliency in storage resource allocation is where many agile IT strategies are encountering problems.

In addition to agility, availability, and elasticity, software-defined storage evangelists also posit the greater resiliency of their server-side storage architecture. Resiliency refers to the predictability or reliability of service delivery and is, in fact, a foundational attribute of agility. An agile data center has a set of processes that must work in a relatively error free manner to provision pooled resources to specific application workload. Resiliency, according to IBM, is a measurement of the reliability or predictability of those processes and the services they provide.

Achieving resiliency in storage resource allocation is where many agile IT strategies are encountering problems. Using metrics introduced by IBM for measuring resiliency, we find that delays in provisioning resources (storage, in this case) to workload may well be more protracted in Virtual SAN architecture than in legacy storage because of hardware challenges – adding new storage requires scaling capacity across multiple nodes in a cluster, for example, and old issues such as locating and installing drivers on all server hardware continue to add steps to provisioning workflow. Another metric suggested by IBM researchers, rates of job rejection (a measure of frequency of rejected resource allocation requests, usually because of inventory management issues or a shortage of needed resources) have produced data in many shops suggesting that, virtualized or not, we have not yet developed the technology or tools for forecasting resource requirements or for inventorying hardware assets in a cost-effective manner.

Conclusion

Resiliency metrics go to both the actual and the perceived agility of IT by those who pay the bills. Until management and orchestration tools are developed, still very much a work in progress, achieving resiliency will likely require smarter hardware.

Simply put, storage gear needs to be smarter, both in terms of how it physically performs and scales and how it avails itself physical monitoring and management. Ultimately, the solution – at the storage layer at least – to the challenges of agile and resilient data center operations is hardware-defined, rather than software-defined, storage.

Only by including in the hardware/ software stack hardware components that can deliver functions still absent from hypervisor software can we begin to make progress toward the agile vision.

X-IO provides an example of a building block for hardware-defined storage worthy of consideration by IT planners pursuing an agile software-defined infrastructure. The implementation of RESTful management enables any device supporting HTTP to be used to manage and administer the provisioning and scaling of X-IO Intelligent Storage Elements (small high performance building block arrays). Moreover, the capabilities of the ISE array to provision solid state and magnetic disk components automatically enables the product to adapt to demanding I/O workloads in an elastic manner – including the workloads hosted under different hypervisors and on bare metal servers.

While selecting hardware for its on-board functionality may seem contradictory to software- defined storage mantra, only by including in the hardware/software stack hardware components that can deliver functions still absent from hypervisor software can we begin to make progress toward the agile vision. Agility is the practical objective of software-defined data centers and may require hardware-defined storage to help to bring it to fruition.

A Perspective By Jon Toigo

Managing Principal Partner,
Toigo Partners International

Chairman,
Data Management Institute

Data Management Institute