



Storage Virtualization Software: A Cornerstone of Affordable Business Continuity

Internet Research Group

February, 2011

About The Internet Research Group

www.irg-intl.com

The Internet Research Group (IRG) provides market research and market strategy services to product and service vendors. IRG services combine the formidable and unique experience and perspective of the two principals: John Katsaros and Peter Christy, each an experienced industry veteran. The overarching mission of IRG is to help clients make faster and better decisions about product strategy, market entry, and market development. Katsaros and Christy published a book on high tech business strategy *Getting It Right the First Time* – Praeger, 2005 www.gettingitrightthefirsttime.com

This whitepaper prepared for DataCore by the Internet Research Group

Storage Virtualization Software: A Cornerstone of Affordable Business Continuity

Late 2008 saw an event that was a major inflection point in the development of business continuity solutions. At the September VMware User Group meeting that year, VMware co-founder and chief scientist Mendel Rosenblum demonstrated a remarkable capability that he called "Continuous Availability," in which a single virtual machine runs in multiple places at the same time. The setup was simple: two instances of a virtual Microsoft Exchange server running in parallel, a primary one supporting users and a secondary one shadowing it. When he unplugged the primary, the system switched over to the secondary instance without missing an e-mail message.

What Mendel's demonstration illuminated was something many had overlooked: that server virtualization not only saved money and simplified operations, but could also support business continuity and, perhaps even more important, make it more affordable. Two years later, the full implications of that demonstration are still being worked out, and one area where its biggest impact is yet to come is storage management, or to be more specific storage virtualization..

We say "yet" because it might fairly be said that most storage hardware vendors don't seem to have realized—or are resisting—those implications. Consider for a moment what storage virtualization comprises: a complex set of functions to manage the three key jobs of storage management – 1. Capacity, 2. Performance and 3. Data Protection at the primary and secondary storage management levels. The last of these includes data protection at the file or logical level, at the device or machine failure level (either locally or at remote sites). This also includes support for backup and restore, archival storage for compliance, the ability to withstand server loss (high availability), and local or remote site replication and recovery (disaster recovery). In other words, a lot of storage virtualization—especially its data protection and secondary storage management functions— is concerned with business continuity.

So here's the question raised by Mendel's demonstration that is disrupting the storage market. If business continuity can be supplied as part of a software virtualization platform (with all the convenience and flexibility that implies), what future do hardware-specific storage virtualization solutions have? After all, just as server virtualization eliminates the need for specialized server hardware, shouldn't storage virtualization eliminate the need for specialized storage hardware or appliances? In fact, doesn't the fundamental logic of virtualization eventually demand a hardware-agnostic storage management solution that can itself run in a virtual machine in one of the general-purpose x86 servers that make up today's ideal data center?

But what may be worrisome for some storage vendors is very good news indeed for customers, for following this logic to its inevitable end will deliver business continuity solutions affordable by a much wider range of companies. In this brief we'll review what's going on with the "business continuity journey" every IT organization faces, and discuss how device-independent storage virtualization software not only makes business continuity more affordable, but offers other cost savings and can have performance advantages as well.

Driving with the Brakes On

The ability of the internet to reduce friction in business-to-business and business-to-consumer communications and transactions has a dark side: it has steadily made businesses of all sizes much more dependent on the availability and resiliency of their computer systems. A decade ago, for many businesses, having a system go offline temporarily might not seriously impact productivity or customer activities. But today, business continuity—defined as continuous operation or recovery quickly enough to avoid serious business consequences—is a requirement for organizations of all sizes.

In the past, providing business continuity was an expensive process fraught with difficulty. For instance, consider what it took to create a viable disaster-recovery (DR) strategy: building a fully-configured replica of the data center at a completely separate location whose sole purpose was to stand by in case of a disruption. Not only was it necessary to replicate the servers and the network equipment, but also the storage subsystems with their data safely mirrored. Adding to the expense, of course, was the time it took to frequently test the DR system to ensure that it could perform its duties.

In reality, DR was so costly that only the largest organizations could afford it, and most businesses had to accept the considerable risk that they would not be able to recover fully from a disaster. The same cost/risk disadvantage obtained all the way down to backup systems, for much the same reason: all levels of the business continuity continuum required some degree of hardware and software duplication in systems whose only function was to standby in case they were needed.

Fortunately, the convergence of three technology trends are transforming the business continuity landscape: virtualization, the cloud, and storage management.

Server virtualization—the insertion of a thin shim of software that abstracts an application implementation from the server it actually uses—has had a profound impact on how data center computing is organized. As alluded to above, the ideal data center today is basically a collection of general-purpose x86 servers running virtual machines under the control of hypervisors, which delivers cost savings and agility too well known to repeat here.

Perhaps the biggest impact of server virtualization has been the development of cloud computing, which might be regarded as a meta- virtualization in the sense that it abstracts much or all of the entire IT infrastructure from the business that it serves. The result is on-demand computing, either within the organization (Private Cloud computing), or using third-party IT resources (Public Cloud computing).

The third part of the equation is storage management, which like server virtualization, is basically a thin shim of software that in this case abstracts the details of storage systems from the applications that use them. The demands of server virtualization and the cloud have driven the evolution of storage management towards a quite complex set of functions, as illustrated in Table 1.

Table 1: Storage Management Functions

I/O performance and disk capacity optimization	Match workloads with the appropriate storage resources and caching algorithms to improve application response and maximize disk space utilization
Shared storage access controls	Set up the appropriate routes and permissions to facilitate access from multiple servers
Business continuity, data protection and secondary storage management	Maintain mirrored copies across separate locations accessible over redundant paths. Replicate selected volumes offsite for disaster recovery. Create snapshots for online backup /restore and migrate older contents to archival storage for compliance

Server virtualization and cloud computing are putting a great deal of pressure on storage management solutions. The problem is that the dominant model for storage management solutions is still hardware-oriented (whether hardware storage systems or appliances front-ending commodity storage), which makes them less agile than the virtualized-server data centers they are intended to serve. In effect, hardware-based storage management makes running a data center like driving with the brakes on. The necessity for specialized hardware storage systems or appliances also drives up cost considerably leaving smaller businesses in the same bind as before: the business continuity functionality they need is still too expensive.

It's time to take the next step in the business continuity journey.

The Business Continuity Journey

In our conversations with customers we have encountered many IT organizations who have undertaken the “journey to business continuity,” a four-stage process illustrated in Table 2 that's a microcosm of the journey the IT industry as a whole has taken.

Table 2: The Business Continuity Journey

Stage 1	Stage 2	Stage 3	Stage 4
Server Virtualization	Beginning Business Continuity	Intermediate Business Continuity	Advanced Business Continuity
Sets the stage for business continuity	Backup & Archival Operations	High Availability	Disaster Recovery
Virtualizes the majority of servers – eventually including business-critical servers	Implements server backup strategy to recover lost data and preserve data required for compliance	Implements local strategy to recover from server and storage failure(s) including planned outages.	Implements offsite strategy to recover from temporary or permanent loss of data center

Stage 1

The first stage of this journey is simply server virtualization, which generally starts with simple, non-critical functions that are characterized by low server utilization (e.g., print servers). As these projects prove the benefits of server virtualization, the organization begins to virtualize more complicated systems such as Microsoft Exchange and SQL servers, and eventually the full range of business-critical applications.

The first stage is the most critical, for it's here that organizations usually make a crucial misstep by leaving out business continuity planning. Instead, one day they wake up and realize that they've virtualized just about every business system, but their ad hoc approach to storage management has left their IT infrastructure dependent on a collection of business continuity solutions: backup, archive, replication, etc. Management costs are growing out of control, they're locked into various vendor solutions, their storage operations are limiting their agility, and they still don't have the availability and resilience they need.

Instead, it's in the first stage of the business continuity journey that IT ought to step back and think seriously about where the logic of virtualization might take them: to a software-based storage virtualization solution that can help them avoid these problems. Let's look at how the following stages of the journey can work out if IT makes that choice.

Stage 2

The second stage of the journey to business continuity generally involves putting in place backup systems which take advantage of virtualization. Prior to server virtualization, backup systems were an assortment of equipment —over-engineered systems which performed simple tasks in a complex way. The evolution of storage management has made backup much simpler, but we can take a further step. With server virtualization, backup can be accomplished by taking a “snapshot” of the server's image and mirroring that data. This yields a huge advantage in that it not only preserves the data, but also the state of the system when the snapshot was taken, thus eliminating much of the difficulty involved in rollback. The restore function is simply a matter of loading the image into another virtual machine and restarting it.

Newer technologies like Continuous Data Protection CDP can now be utilized to recover from file or logical level problems such as those caused by virus attacks. Some virtualization systems include 'rollback' capabilities that allow users to go back in time not just to a particular snapshot but to any point in time prior to the attack and undo undesirable changes to files or systems.

Of course, there will always be a need for “simple” data backup, especially as regards the archival demands of various regulatory regimes and just basic good governance. But again, why not virtualize this functionality as well? What is the business case for specialized hardware systems or front-end appliances to supply such functionality, especially given the limited agility and vendor lock-in they involve?

Stage 3 & 4

Lastly, the third and later stages of building business continuity systems involve developing high-availability and disaster-recovery scenarios which will allow continuous access to systems using redundant facilities. These extend to the recovery of data center functionality from unplanned system outages. Virtualization actually blurs the line between this and the previous backup stage, for as Mendel's demonstration illustrated, it's a short step from making a static copy of a virtualized server image to actually running the redundant system as a virtualized duplicate of the application, whether locally or remotely. And since there's no need to match hardware configurations, resources for such redundancy can easily be found, whether inside the data center, at a co-location site, or even in a public cloud such as Amazon AWS. The end-state of this process is the provision of business continuity as a utility, at a cost that almost any business can afford.

The Bottom Line

Naturally, the devil is in the details, and the business continuity journey is not an effortless one. But by basing a business continuity strategy on a storage virtualization software solution that layers across the wide variety of desirable hardware choices, each step becomes much easier and less expensive.

To begin with, a totally software-based storage virtualization solution is hardware-neutral, which yields a lower overall TCO for storage. IT enjoys the freedom to choose the optimal balance of price and performance for the storage hardware needed to fulfill a given function (e.g., archival vs. transaction processing), without any worries about vendor lock-in. The solution preserves an organization's present storage hardware investment and permits upgrades as the IT budget allows. It also reduces storage costs by eliminating the need to reserve large amounts of storage in anticipation of future use—what's called "thin provisioning."

In addition, look for those products that can be implemented directly on host systems via virtual machines or natively above the OS, eliminating yet another hardware dependency and enabling them to supply storage management services and features across a heterogeneous storage infrastructure. Equally important, more sophisticated storage virtualization software can inexpensively leverage server RAM and spare processors for high-speed caching, making it even more cost-effective at improving I/O performance. After all, it does no good to virtualize if users take a big hit on application response.

With a software-based solution, all of an organization's storage resources are manageable from a single console. This makes administrators far more productive, helps reduce storage-related downtime and disruptions, and provides a high-level view across a variety of hardware and vendor systems that can only improve storage and business continuity planning. When choosing storage virtualization software, pay special attention to how they automate storage management to minimize dependencies on storage specialists. Also consider the software supplier's track record with regards to projects involving server and desktop virtualization.

Finally, by increasing the overall agility of the storage system and decoupling it from the underlying storage hardware, software-based storage virtualization makes it much easier to provide whatever

business continuity functions are called for, whether they are currently-understood functions such as automatic mirroring of data between separate storage devices and the movement of the resulting data files both locally and remotely, or as yet undreamed-of techniques for ensuring non-stop business operations. Thanks to virtualization, evolution to new business continuity capabilities is merely a matter of a software upgrade.