White Paper

Storage Virtualization for Efficient Operations

With a Focus on IBM's Capabilities

By Mark Peters, Senior Analyst

March, 2012

This ESG White Paper was commissioned by IBM and is distributed under license from ESG.

Contents

Efficiency is Not a Fad

The Basic Need

Creating a more efficient IT infrastructure is not some passing fad—it is an essential ingredient in today's business environment. Moreover, focus on efficiency has increased during the past few years of difficult economic conditions. With flat or declining budgets, all organizations have been asked to "do more with less," and while budgets are beginning to recover, the need for efficiency does not dissipate.

Now that the possibilities for IT efficiency are known, there is no way to return to a care-free past. Even when times are flush, an inherently efficient infrastructure enables organizations to maximize their profits and make better use of resources and staff. It is simple logic, therefore, that in the best of all worlds, optimum efficiency should become part of the fabric of both business processes and IT infrastructures.

The Move to Greater Server Efficiency

To see the value of greater efficiency in living color, one need look no further than the server virtualization revolution. IT organizations were struggling under the weight of proliferating servers, each typically housing a single application. Server farms of hundreds, thousands, even tens of thousands of individual machines became nearly impossible to manage. Server virtualization turned that idea on its head. Consolidating multiple application workloads on far fewer hardware platforms has revolutionized the ability for users to be efficient in terms of "compute-ability"[1]. For many organizations, server hardware has been reduced by 50%, 60%, even 80%, enabling both capital and operational cost savings. As a result, server virtualization–which only a few years ago seemed thrilling, unusual, and even somewhat geeky –is now a staple of the data center. Excitement over its value as an IT strategy has not diminished, a fact that is borne out in ESG's most recent IT Spending Intentions research. As Figure 1 shows, the increased use of server virtualization is once again (for the third year in a row) the top IT priority according to organizations we surveyed in late 2011.[2]

A Stark Contrast to Storage

Server virtualization surely demonstrates the benefit of data center efficiency and renders the unyielding growth of data volumes—a reality from which IT cannot seem to escape—even more conspicuous. Not only is there a natural application growth occurring, but social media, web 2.0 applications, smartphones, and tablets enable an onslaught of both content creation and consumption. This has resulted in pressure for performance as well as capacity. Not only are organizations using the massive amounts of data collected constantly by websites and barcode readers, but that data is also then fed into analytical engines that spit out the right products targeted to the right buyers. Big data sets, as well as requirements for high application availability, data protection, and regulatory compliance have a tremendous impact on data growth.

For the most part, organizations have been addressing this particular challenge simply by adding storage capacity. If you need data copies for analytics, you get a server (physical or virtual) and some storage. Adding infrastructure stacks for testing and development, remote replication, backup, etc., seemed like a good idea at the time, but has resulted in parallel silos of infrastructure that keep growing yet are mostly poorly utilized. This response to growth has created tremendous inefficiencies in the storage domain, resulting in not only underutilized resources but also often with un-sharable storage capacity stranded behind specific uses. In some cases, upwards of 70% of storage capacity is sitting idle, locked behind a particular server or purpose, while another silo might be running out of capacity. Buying more storage is such a waste when unused capacity is sitting around! In some cases, storage has become the poster child for inefficiency, driving up acquisition costs, consuming data center floor space, energy, and increased management time.

[1] This constructed word is being used to convey not just 'regular' efficiency in terms of reducing the number of servers (and therefore also space, energy, and cooling demands), but also other elements such as greater flexibility, better responsiveness, and reduced complexity.
[2] Source: ESG Research Report: *2012 IT Spending Intentions Survey*, January 2012.

Server virtualization and traditional storage architectures seem to be a gigantic mismatch.

Figure 1. Top Ten IT Priorities for 2012

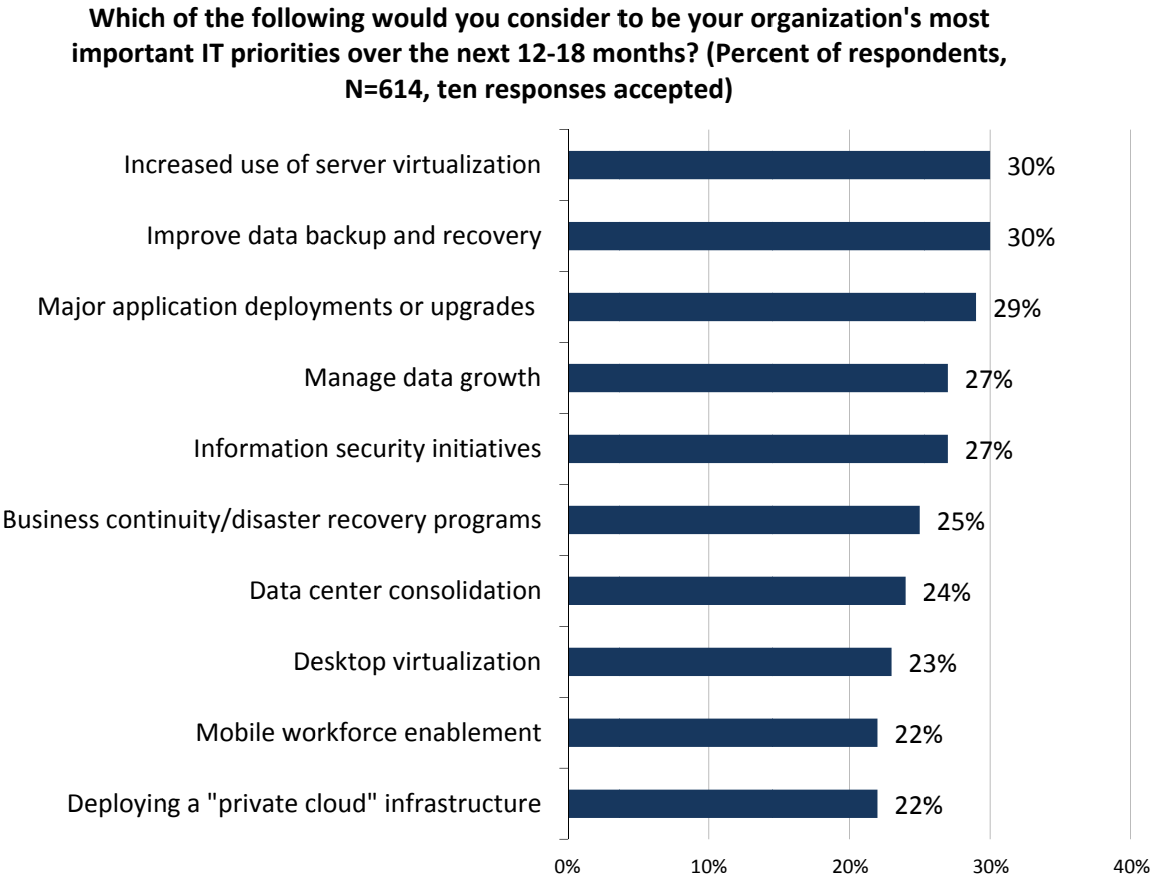

Source: Enterprise Strategy Group, 2012.

Server Virtualization Can Reflect Poorly on Storage

Server virtualization lets physical servers host multiple application workloads. Virtual machines (VM) can be moved between physical machines without interrupting operations, VMs can be spun up in minutes to address a particular need, and organizations can save tremendous amounts of money by having to buy, manage and software-license-provision fewer servers. This makes the IT organization much more flexible and responsive. Dramatic improvements in utilization rates and more efficient energy consumption can create significant savings, as can reductions in both maintenance and software licensing.

Server virtualization truly shifted the IT world on its axis, enabling organizations to improve application availability and scale quickly, as well as bringing remote replication into the realm of possibility for many who could not afford the like-for-like infrastructure that was previously required. The advances that server virtualization has generated highlight the inadequacies of other infrastructure components in this type of environment. Aggregated server workloads put tremendous stress on both networking and storage. In addition, while you may be able to move a running application between servers, the storage doesn't automatically go with it. Virtualizing the compute side but not the storage side prevents the complete agility and truly utility-like resource provisioning that users are looking for. Unfortunately for IT, users are becoming accustomed to accessing an application server pretty much instantly, which results in ever-higher service-level expectations. If the storage cannot match what the servers can do, those service levels cannot be met.

Users of Server Virtualization Speak

When 21st-century server technology meets 40-year-old storage design, it is—to put it mildly—a bit of a cultural and technical clash. As a result, storage implementations can prevent IT organizations from expanding their virtual server deployments in the ways that they would like to. That is an easy, even glib, assertion, but it is proven by what IT managers revealed in ESG research focused on server virtualization[3]. First, it is clear that the more a user virtualizes their environment, the more benefits that user will invariably gain. ESG's research segmented server virtualization users into three categories based on the maturity of their implementations—this was defined by the scope of deployments, virtual-to-physical consolidation ratio, percentage of virtual machines running in a production environment (as opposed to test and development), and penetration of server virtualization into production and mission-critical applications. It quickly became clear that basic benefits such as reducing acquisition costs and floor space requirements were benefits shared by organizations at all levels of the maturity model. However, only those organizations that had been able to expand their virtualization deployments were experiencing higher level benefits such as improving disaster recovery, increasing application availability, and speeding application provisioning.

To no one's great surprise, therefore, expanding server virtualization is a high priority for many organizations, and most of the research respondents had immediate plans to virtualize more applications, consolidate more VMs on each physical machine, and move server virtualization increasingly into production operations. Unfortunately, many of them were also experiencing roadblocks to virtualization expansion, and storage implementations were often the culprit. The storage challenges associated with server virtualization cited by respondents included the capital cost of new storage infrastructure (36%), scalability problems (25%), operational costs of new storage infrastructure (24%), and the impact on overall volume of storage capacity (21%). When asked what would help them to expand their virtual server implementations, the most-mentioned needs were not just more training, but better storage capabilities—faster provisioning, better scalability, more use of virtualized storage, and better storage migration tools (see Figure 2).

Figure 2. Storage Developments That Would Enable Wider Server Virtualization Usage

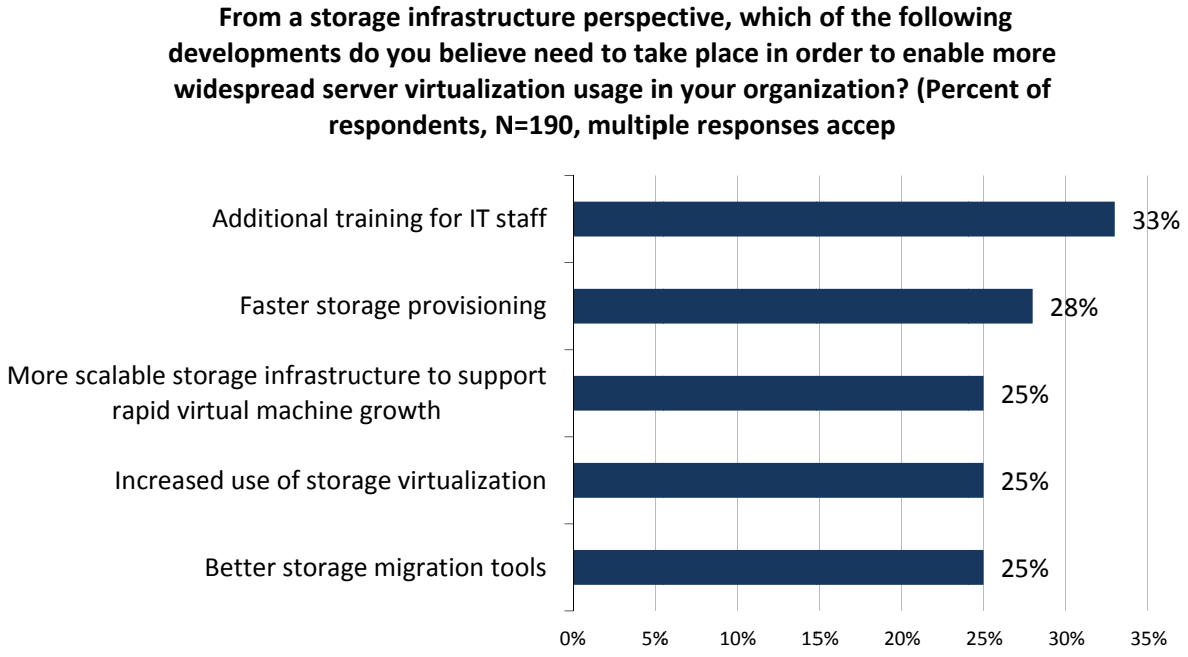

Source: Enterprise Strategy Group, 2011.

[3] Source: ESG Research Report, *The Evolution of Server Virtualization*, November 2010. All server virtualization research references come from this report.

Storage Virtualization Can Help

Given this vexing situation, many organizations are looking for more flexible and scalable storage implementations. What server virtualization deployments need is storage that:

- Helps to increase application availability instead of diminishing it.
- Promotes flexibility instead of hindering it.
- Speeds provisioning instead of delaying it.

Virtualizing storage is essentially accomplished in the same way that servers are virtualized: the physical hardware and disk resources are logically separated from the data services. With virtualized storage, a pool of flexible, consolidated resources with a common management interface is built using capacity from multiple arrays—arrays that can have different specifications, software features, and sometimes can even be from different vendors. Storage virtualization enables organizations to improve performance and increase capacity without disrupting operations. Because storage is not dedicated to any specific server or application, it creates a much more flexible and responsive environment while also increasing utilization rates. Virtualization can also simplify storage tiering, another good efficiency strategy: its benefits are twofold: firstly, more expensive, tier 1 storage can be used for frequently accessed "hot" data while less expensive, capacity-optimized storage can be used for "cold" data. Secondly, users can get away from a "default" or "one-size-fits-all" storage approach (which is all-too-often tier 1) and move to a storage mix that better meets their application needs.

Advantages of a Virtualized Storage Pool

With storage virtualization, the data services provided to users are no longer limited by hardware characteristics, and with tiering, IT can match the performance and availability characteristics with data value to streamline costs. Instead of silos of storage capacity that cannot be shared, a virtualized storage pool can provide a number of benefits:

- **Reclaiming stranded storage capacity**. Unused capacity now becomes part of the virtual storage pool and is available to any servers, users, or applications accessing the pool. In addition, when storage is no longer needed it becomes part of the pool again to be used elsewhere.
- **Consolidating storage management**. Storage virtualization centralizes management of all storage resources, eliminating multiple tools and processes for different arrays. Pulling storage management all together under one umbrella makes the environment more reliable, providing less opportunity for administrative errors.
- **Advanced software features**. With storage virtualization, high-end features (which are also very often economically-efficient too) such as thin provisioning, snapshots, and remote replication that come with the storage system can be shared with arrays that don't have those features. This can extend the useful life of storage devices as well as maximize the return on investment. In many cases, it can prevent or at least delay additional storage purchases.
- **Green operations**. By improving utilization and so reducing the number of physical machines, storage virtualization reduces data center floor space, power, and cooling requirements.
- **Non-disruptive migration**. Data movement is one of the most common and disruptive storage tasks. it is done for many reasons, perhaps the most common of which is to address performance problems. Storage virtualization enables data to be moved non-disruptively and so turns an irksome but necessary task into something to be welcomed.
- **Total cost of ownership**. Storage virtualization replaces a fragmented environment,- in which silos of SAN, NAS, virtual tape, etc., all operate with separate storage networks and management tools—with a centralized, flexible storage pool. Equipment, management, energy, maintenance, and licensing costs are reduced, while storage is used more efficiently.

IBM's Storage Virtualization Leadership

The old adage that "nobody ever got fired for buying IBM" is rooted in the underlying trust held by users that IBM's solutions do what they promise. The company does just about everything in IT, and it invariably does it well. The only downside of this scenario is that IBM doesn't necessarily put its storage solutions front and center—which is unfortunate, because what it has to offer is really quite special, and highly valuable for users.

Key IBM storage virtualization products are the SAN Volume Controller (SVC) storage virtualization system and the XIV and Storwize V7000 storage systems. SVC is a storage virtualization system. Storwize V7000 and XIV are disk systems with advanced virtual capabilities that help analyze, adapt to, and improve application performance while remaining easy to use. These intelligent storage solutions help users to maintain application availability and performance levels in dynamic virtual and cloud-based infrastructures, and are good examples of IBM's commitment to "Smarter Computing." Smarter Computing is about seeing the issues, and hence the infrastructure, holistically; designing strategies and solutions that enable breakthrough service delivery - optimizing systems for the workloads they support; and making infrastructure flexible, intelligent, and interconnected. All of these approaches and technologies improve efficiency and flexibility and provide a responsive IT infrastructure that is based on virtualization, consolidation, and automation—essential for many environments, including contemporary cloud deployments.

SAN Volume Controller

The goal of SVC is to provide a virtual storage pool with a single point of control that manages all storage resources to the best advantage of the user's business—quickly, efficiently, in real-time, and without increasing administrative complexity or cost. SVC is a modular, highly scalable system of redundant pairs of storage engines. Along with cache, and both Fibre Channel and iSCSI ports, the system scales to support up to 1,024 host servers and 32PB of virtualized storage capacity. In addition to innovative support for solid-state drives that can scale out for performance, SVC includes IBM's System Storage Easy Tier functionality which makes the most efficient use of the solid-state storage by automatically identifying and moving highly accessed data to that storage. Easy Tier also supports solid-state disk in the virtualized arrays, and because it is tightly integrated can speed data migration and replication functionality. SVC capacity pooling increases utilization and helps insulate host applications from storage changes to improve application availability. The consistent management interface minimizes administrative effort. Instead of having to operate multiple interfaces for storage configuration and management, SVC enables users to dramatically increase the productivity of storage administrators. SVC also enables a stretched configuration with clusters supporting storage and servers in different data centers.

SVC offers plug-ins to both VMware vCenter and Microsoft System Center Operations Manager for even greater management consolidation. SVC features are built to complement the server virtualization paradigm. Provisioning is fast because it is a non-disruptive software operation. Non-disruptive data migration in the storage pool means that functions such as VMware vMotion are not stymied. SVC also supports the VMware vStorage APIs (VAAI) so that storage can offload some tasks from VMware to free up server resources. Like server virtualization, SVC enables different physical configurations at production and DR sites but creates the same virtual configuration at each. And, in the same way that virtual servers can be deployed on various types of server hardware, SVC enables a heterogeneous storage environment.

XIV

Enterprise-class XIV storage is grid-based, automatically spreading data across all disk drives to optimize application performance and eliminate hot spots. Its grid architecture is purpose-built for virtual and cloud infrastructures. Consolidated virtual workloads, by definition, create increasingly random I/O profiles because many different applications are sending reads and writes. The grid structure was designed to use massive parallelism to optimize random I/O, so it's faster by nature. In addition, performance can grow along with capacity, eliminating that trade-off. XIV's self-tuning ability enables high performance using a combination of low-cost, high-capacity drives to reduce TCO as well as a sophisticated SSD cache layer. This strategy differs from trying to improve performance

with a tier of high-speed, low capacity spinning disk drives that is constantly under-utilized, and it is the kind of new idea customers have come to expect from IBM.

XIV is built for high reliability and availability with redundancy and proactive self-healing. Built-in management tools are extremely simple to use, and arrays seamlessly scale in capacity, cache, processing power, and host connectivity. XIV even has free iPad and iPhone apps (XIV Mobile Dashboard) which can be used to monitor XIV systems from anywhere. Other built-in advanced features include space-efficient snapshots, global mirroring for replication and disaster recovery, consistency groups, thin provisioning, and automated data migration. In addition, the product is deeply integrated with VMware solutions and APIs, as IBM and VMware have a long-standing partnership. Not only does integration with vCenter Site Recovery Manager simplify site failover, but some VMware functions can be offloaded to the XIV array through VAAI for processing efficiency. These are features that customers should expect from their storage vendors because they help get more from the virtualized environment.

Storwize V7000

For the mid-range, IBM Storwize V7000 tackles performance and efficiency in a virtual world with built-in thin provisioning, Easy Tier automated storage tiering, and storage virtualization. The new Storwize V7000 Unified system can support file and block storage with fully integrated management to handle the massive growth of unstructured data without requiring a different storage silo. These systems scale up (and Storwize V7000 block systems scale out) easily to accommodate the rapidly changing needs of virtual and cloud-based environments, and Storwize V7000 Unified includes automatic, policy-driven tiering to optimize performance and cost while simplifying administration.

Integration with VAAI lets the Storwize V7000 and Storwize V7000 Unified take some of the load from servers to free up processing resources. Improved mirroring lets organizations balance network costs with recovery point objectives. External storage virtualization enables non-disruptive migration from legacy and non-IBM arrays or allows them to be retained and consolidated—another efficiency measure designed for investment protection. Snapshots and replication are easy to manage for efficient protection. Storwize V7000 Unified also supports IBM's new Active Cloud Engine, delivering policy-based file management for backup archiving, and migration. In today's efficiency-focused world, organizations can reduce CAPEX and OPEX with consolidation, virtualization, and automation. Storwize V7000 and Storwize V7000 Unified have been designed for exactly those needs.

The Bigger Truth

For the most part, organizations have had a fairly cavalier attitude towards storage. It's not that IT wasn't unaware of its rising costs—it's more that it didn't see any other ways to address the problems with which it was faced. But server virtualization changed all that, demonstrating a way to operate with greater efficiency and less waste.

The recent flooding in Thailand that created a disk drive shortage has been a new, rude awakening. Suddenly, it has become more expensive—even supply constrained—for additional storage silos to be the answer to every problem. This is a wake-up call for those who are willing to see it. As long as storage infrastructures remain inflexible and inefficient, any disruptive force can quickly overwhelm an IT infrastructure. Natural disasters are not the only types of interference–they may also come in the form of sudden financial problems, regulatory needs, and unforeseen competitive threats that will require a prompt response. Only if a storage environment is built on a foundation of efficiency and flexibility will IT organizations be able to handle these problems.

The efficiency and cost savings of the virtualization paradigm have been proven in stunning fashion on the server side of things. Storage can–and should–provide the same types of benefits. IBM's SVC, XIV, and Storwize V7000 offer benefits that match the efficiency needs of today's organizations, and are part of IBM's ongoing leadership in creating a Smarter Planet through Smarter Computing. The economic and operational value of virtualizing storage is clear. So, for IT users that want to spend more time thinking about the business services they deliver and less about the boxes underneath, storage virtualization is the way to go. It is , if you will, *"Smarter Storage*. "

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com