



A CMI Business Brief White Paper:

Why Storage Virtualization is Critical

By Jeff Guenthner, CMI Storage Practice Executive





At CMI, we understand that companies today must balance conflicting priorities. In this White Paper, we address the conflict between lowering business costs and increasing operational efficiencies in mass data storage, while embracing innovation and new business models in the enterprise and in IT. CMI believes quite strongly that storage virtualization can and should play a pivotal role in all those areas.

That said, the goal of this White Paper is not to write a definitive work on 'everything storage virtualization'. Rather, our goal is to use a few real world technological and business examples to quickly define and qualify a handful of key aspects of this topic while raising the reader's interest in this fascinating and critical technology.

Let's start-out with a quick definition of storage virtualization at work: storage virtualization adds a layer of software and/or hardware between storage systems (disk drives, network attached storage (NAS), or storage area networks (SANs) and servers, so that applications no longer need to know on which specific drives, partitions, or storage subsystems their data resides. Using storage virtualization tools, system administration staff can identify, provision, and manage local and/or distributed storage as if it were a single, consolidated resource.

Storage virtualization has many benefits, as we will see, including the ability to increase application availability, since applications aren't restricted to a particular storage resource(s) and are insulated from most storage-based disruptions. In addition, storage virtualization eliminates vendor lock-in, preserving your firm's flexibility to choose storage from multiple vendors (more on this later). Moreover, storage virtualization software also permits storage resources to be updated "on the fly" -- without affecting application performance or impacting availability, resulting in savings of time and money.

Put another way, storage virtualization means storage assets are handled in a federated manner via a single, centralized storage console, allowing management of storage, data migrations, and legacy storage management in a simplified fashion. In addition, storage virtualization allows disks and logical drives (LUNs) from heterogeneous systems to be aggregated into "storage pools." This means that different classes (as opposed to simply different brands) of storage from multiple vendors can be tiered and managed as a whole, resulting in reduced hardware costs, increased efficiency, better flexibility, and more scalability of your data center.

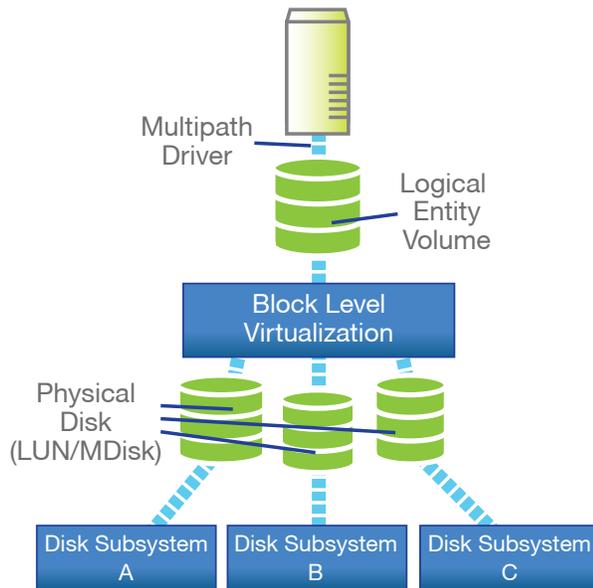


Figure describing how storage virtualization works

Storage Virtualization at Work

The Louisiana Supreme Court manages its redundancy and disaster recovery needs through virtual storage. When attorneys switched to electronic document management and video and audio notes, their IT Director switched to a virtual storage and disaster recovery system, enabling them to manage multiple instances of data more easily.²

As to what kind of savings people are seeing with virtualization, according to independent research firm IDC, an "advanced virtualization" infrastructure, consisting of virtualized servers, storage virtualization, and systems management tools, can produce an aggregated cost reduction of up to 52% per user per year as compared with an un-virtualized static Intel-based server configuration.¹ True, this example is a broader solution than just storage virtualization -- but the point is manifest and worthy of any executive's attention.

Now that we have a basic working definition of what storage virtualization is and what it does, let's look next at storage virtualization and a key issue of our time: data growth.

Growth Drives Complexity -- and Waste

Businesses worldwide generate more data per day today than ever before. Take just one company as an example: according to Niall Kennedy's weblog of July 5, 2010, "Google processes over 20 petabytes of data per day..." (A petabyte is 1,000 trillion bytes of data.)

¹ http://h18004.www1.hp.com/products/servers/management/vse/Biz_Virtualization_White_Paper.pdf

² Source: http://www.cio.com/article/593296/Storage_Virtualization_When_to_Invest



Google may be somewhat obvious, but in organizations large and small, there is growing demand for storage environments that are as flexible (in terms of sizing, elasticity, and operational capabilities) as the applications and underlying server virtualization that supports them.

According to a 2012 report by IDC, the amount of data generated by enterprises is expected to grow by 48% in 2012 alone -- and 90% of that growth will be in what is called “unstructured data.”³ According to Wikipedia: “Unstructured Data (or unstructured information) refers to information that either does not have a pre-defined data model and/or does not fit well into relational tables. Unstructured information is typically text-heavy, but may contain data such as dates, numbers, and facts as well.”

The traditional approach to addressing the issue of huge data growth has been to purchase or lease vast amounts of storage media and fill it up over time. Leaving the obvious monetary issues aside for a moment, this approach makes it difficult to allocate additional storage or make changes to the underlying infrastructure. However, there are other utilization issues at work here as well. CMI has seen data from IBM that paints a scenario in which many organizations are only utilizing storage assets at about 20 to 50% of capacity. That is wasteful storage utilization on an industrial scale. Storage virtualization addresses these issues by allowing administrators to intelligently manage data storage infrastructure more efficiently and on an “as needed, as we grow” basis rather than the “all at once, up front” or “poorly utilized” models.

As an example of that intelligent management, let’s look very briefly at a storage virtualization process known as thin provisioning. With this technology, a total storage requirement of 15 TBs, projected across an entire data storage pool, might only need the purchase of 5 TBs of physical storage space. Net-net: savings on both initial storage purchases as well as savings in HVAC, electricity, and data-center costs.

³ <http://www.idc.com/getdoc.jsp?containerId=prUS23177411>

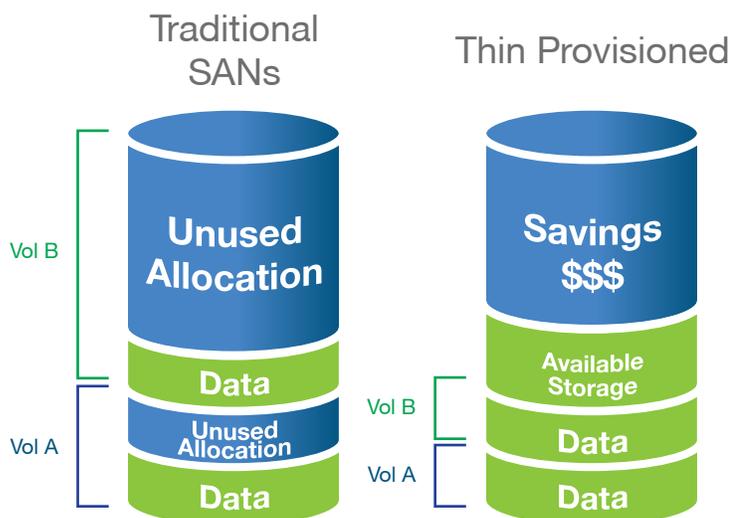


Figure: Thin Provisioning at Work

Storage Virtualization at Work

The city government of Richmond, VA uses IBM's SVC platform to manage its EMC and IBM storage devices, increasing storage utilization by 45%.⁴

Brass Tacks

At CMI, we see additional client concerns driving interest in storage virtualization, including (in no particular order):

- Data center costs
- Ensuring High Availability of data under storage management
- High variability of data acquisition costs
- Inability to respond quickly to business demands within IT
- Lack of awareness/experience of/in Best Practices in storage strategy
- Lack of skilled staff for storage administration
- Poor availability or service levels (within the storage environment)
- Poor storage asset utilization

Marrying the demand we just mentioned above with these concerns and you can quickly see the added importance of addressing complexity by applying storage virtualization solutions.

In addition, storage virtualization also saves money in the long term for older

⁴ Source: IBM Storage Virtualization Value Report 2005



hardware. According to a 2011 Gartner Group report, virtual storage appliances can extend the life of legacy storage environment, too, reducing long-term maintenance costs and simplifying migration to new devices.⁵

In addition, storage virtualization savings can also be realized in dealing with the everyday problem of upgrades or changes to a traditional storage infrastructure environment. Here, organizations often have to take critical applications and storage offline when an IT systems administrator performs the upgrade or change. However, doing so interrupts the availability of applications and prevents users from accessing data. IBM analysts estimate that the costs of this downtime can be up to \$2 million per hour for large companies that rely on mission-critical data.

Contemporary virtualization platforms are designed to support continuous, uninterrupted access to data by applications and users while IT staff migrate, share, reconfigure, expand, or update storage infrastructure. Utilizing a storage virtualization solution in this situation means an IT manager can reduce or eliminate costly downtime and increase (or maintain) application availability.

However, it would be invalid to presume that only large enterprises benefit from storage virtualization. CMI has seen presentation after presentation from firms including IBM, which show conclusively that storage virtualization can help mid-market-sized companies take advantage of the cost savings of midrange-sized storage solutions -- while still benefiting from high-performance enterprise storage.

Moreover, IBM points-out that since cost savings on midrange-class storage can be up to as much as 50% over enterprise-class gear, the flexibility to choose midrange hardware can be a significant cost savings to any sized organization. If an organization doesn't need the performance that comes with enterprise-level hardware, it can still enjoy the business benefits of storage virtualization on more modest midrange-class hardware while providing the robustness of federated storage.

The following graphic from IBM illustrates just one example of the tremendous possibilities inherent in this technology. This is what we at CMI call "Brass Tacks Savings."

⁵ Gartner Technical Professional Advice: Storage Virtualization: Steppingstone to a Better Environment. 12 May 2011. ID:G00211854

Potential Savings on \$1 Million Data Storage Budget



Source: IBM Enterprise Strategy Group 2011

Finally, let us be clear that these are not isolated companies in ones-and-twos who are deploying storage virtualization. According to a 2010 survey by CIO magazine, 38% of responding IT professionals had already deployed virtual storage in their companies and another 31% were actively researching virtualization.⁶

Making It Work

Now that you have a survey of the advantages of storage virtualization, you might be wondering how to make this technology work in your environment? As CMI sees it, in order to best optimize your storage virtualization strategy we recommend taking the following steps:

1: Start With an Assessment of What You Have Now

You need to know where you are to know where you are going. As such, it is critical for you to take an inventory of your existing storage infrastructure to determine what is currently in place and where it is located. Follow that by a comprehensive inventory of your storage software, too. Combined, these will enable your IT department to determine what hardware, software, or application upgrades might need to be made, and what data (if any) will need to be migrated to your new environment. You'd be amazed at how often organizations who undertake this

⁶ <http://www.cio.com/documents/pdfs/CIOTechnologyPrioritiesJan2010-FINAL.pdf>

⁷ Source: http://www-01.ibm.com/software/success/cssdb.nsf/CS/DLAS-8SBMEE?OpenDocument&Site=default&cty=en_us

Storage Virtualization at Work

Ricoh, a 75-year-old leading provider of advanced office technology and innovative document imaging products, services and software, implemented storage virtualization and saw benefits immediately. Marty Everett, Director of Architecture and Solutions Design says, "The amount of storage Ricoh manages has jumped from 500 terabytes to almost two petabytes of physical storage, plus 1.2 petabytes of virtual storage." Storage virtualization has been key to providing efficient and cost-effective storage solutions for our data center," says Everett. "We have maintained a lower price point and effectively managed the explosive storage environment within our data center."⁷



exercise find storage sitting around they didn't know they had.

2: Plan for Growth

Next, determining the areas in your data/storage environment that have the greatest potential for growth is essential, for reasons we hope we've made clear. In addition, planning for growth helps identify where an investment in storage and storage virtualization will yield the greatest benefit(s). In addition, knowing what the performance, availability, and storage needs of your current (and any planned/future) applications will also help you determine the mix of storage infrastructure that will be needed.

By the way: at CMI, we have seen many organizations try to skimp on this step – do not do it. The old adage “Prior Planning Prevents Poor Performance” is as true in storage planning as anywhere, so take the time to do this step in detail.

3: Choose a Manager

Next, you need to continue with the storage virtualization process by choosing a virtualization manager. To reduce interoperability headaches, consider asking storage virtualization vendors for a compatibility matrix that outlines not only the storage products each supports but the versions and configurations. Keep this data with your inventory.

One Example is IBM's SAN Volume Controller (SVC). As you recall, this software's job is to take over the management of existing storage hardware, thereby easing the transition to a fully virtualized platform.

In addition, IBM's SVC is:

- A mature, sixth-generation virtualization solution so it has plenty of real world experience behind it
- Is based on the “open standards” you hear so much about from IT people
- Is consistent with the Storage Networking Industry Association (SNIA) storage model (basically, the industry's standard model).



Those are three key reasons why at CMI we like IBM's SVC. Additionally, we also appreciate that IBM's SVC has been engineered from the ground-up to address several other key requirements, including:

- Solid State Drive Storage (SSDs). With no moving parts, SSDs are the future of high performance data storage. IBM's SVC is designed to deliver ultra-high performance capability (up to 800,000 read I/Os per second), yet it has a small entry configuration, which helps make SSD more affordable. We see that as a win-win.
- Next, we appreciate the fact that SVC will allow your IT folks to migrate data among storage devices without taking applications using that data offline! This means improved operations, less impact to your users, and the ability to manage and scale storage capacity without disrupting running applications.
- SVC includes thin-provisioning capabilities (as described earlier) meaning organizations can save money by dramatically lowering their physical storage requirements.

That's just the tip of the iceberg. In our opinion, IBM's SVC is an excellent example of storage virtualization at work, and we could wax-on about it all day. Suffice it to say it is an offering we point to that demonstrates, to CFO's and IT types alike, the benefits of contemporary storage virtualization at work.

4: Know When to Ask for Help

Finally, know when to ask for help. Storage virtualization is a complex, highly intricate science. There are many 'tricks-of-the-trade' and common sense, experiential aspects to this subject that simply can't be picked-up casually, or from a book, or from a fist full of brochures. Knowing when to call on someone with real world experience, such as CMI or IBM, can make the difference between realizing storage virtualization and cursing it.

It's a Wrap

As you've seen in this White Paper, making use of storage virtualization as the foundation for a flexible and reliable storage solution helps enterprises of all shapes and sizes do a better and most cost-efficient job of aligning business and IT (and



enhancing/optimizing IT operations as well), by optimizing their storage infrastructure and storage management capabilities to meet business demands and the growing tide of business data.

At CMI we have lots of experience in storage virtualization. Give us a call and let's discuss this exciting technology today.



Status Check -- by Kris Neely, CMI's Chief Technology Officer

Storage virtualization has quite simply revolutionized the business of data storage. I'm frankly appalled at the number of clients I encounter who have not virtualized their storage environment and are simply drowning in data – and hemorrhaging money to try and cope.

While no one technology can be said to be a “one-size-fits-all” solution, the business case for storage virtualization is so overwhelming that it's fast approaching “Must Have” status.

That said, there are some common-sense questions I always asked myself when buying this technology and I recommend you ask yourself before you head down the storage virtualization path:

- What problem are we trying to solve?
 - Drowning in data? Operational inefficiencies? Vendor lock-in? Know what you're trying to do – it'll keep you focused on the steak and not chasing the vendor-supplied peas around the plate.
- How much complexity can my IT shop handle?
 - Storage virtualization involves hardware, software (very sophisticated and complex software, to boot), and often large changes to how IT operations are done. For some environments, the extra burden just isn't worth the candle. The financial promises of storage virtualization can be appealing – but they won't mean a thing if your IT Department can't deliver professionally managed storage (and data, and applications) on a consistent daily basis.
- What's my budget?
 - Storage virtualization adds to your budget, to be sure. So let's make sure what we're buying, what we need and why we need it, and let's get some hard numbers on paper. Prices for storage virtualization vary widely. And, don't forget that “Annual Maintenance” bill that never seems to be mentioned during the product presentation!
- Do I have an exit strategy?
 - What happens if the storage virtualization doesn't live-up to what you had hoped? Or what happens if your IT Department can't keep up with the operational demands? It's not always a “baby-with-the-bath-water” situation, but having a well-thought out exit strategy is simply common sense.
- Do I have a technology partner who knows what's-what?
 - This is one technology I frankly don't like to take-on alone. The field, the technology, the models, and the Best Practices of storage virtualization change constantly as this stuff becomes more and more powerful – and complex. Don't wade into this water with one boot off: find a good technology partner who can tell you the truth about how, why, and where this technology might fit in your environment.

Storage virtualization is a very important technology that every CFO, CEO, and CIO needs to carefully consider while bearing in mind that it certainly is not a “Lowest Price Is Best” solution.

I'd be happy to talk with you or your staff about this topic.

*Best,
Kris Neely
Chief Technology Officer, CMI*



Reality Check -- A Veteran CFO Comments on Storage Virtualization

First off, I support both the assertions and the conclusions of this CMI White Paper. Let's establish that up-front.

Next, I have worked in both Visual Effects and Video Game companies and have personally seen how these industries often set the pace for IT technologies, and in particular the scalability of IT infrastructure. These industries are already seeing enormous cost benefits from storage virtualization. As just one example, not all data is created equally. For example, in the Visual Effects industry, it is not uncommon for a single data file to be accessed literally tens of thousands of times a day by dozens of sources. Other data may not be accessed on a daily basis at all, but is equally as important in terms of Intellectual property to the organization that owns it. This is an extreme example, but every company has a similar issue in its data layers – consider how many forms of data your company has and the usage profiles and importance of each.

In the past, this type of situation would be managed internally by the IT department in the types of storage devices purchased and the choices in how and where data is stored. Ultimately, you would make the best of the infrastructure you had and because internal IT infrastructure is usually a fixed cost, efficiency was hard to measure. These days, with storage virtualization, the cost of storage moves from being a fixed to a variable cost. As a result, the benefits of being able to manage data more efficiently become enormous.

As a CFO, one is constantly trying to balance the choices between the costs and ROI of providing IT infrastructure with the risks of being undersupplied in hardware. You never want to be the person who is responsible for an issue that arises due to a lack of IT capital equipment and so any time a solution is available that enables increased efficiency of resources, it is worth investigating. Storage virtualization is just such a solution, as this White Paper details.

Also to a CFO, protecting resources is an equally important responsibility, and storage virtualization allows for additional benefits in terms of the ability to seamlessly migrate and backup data. For many companies, this alone can justify further investigation.

Similarly, being able to be vendor agnostic in your IT purchases not only provides potential cost savings, but also removes a very real business risk of being dependent on a single supplier.

To put it simply, Storage Virtualization is, or should be, a cornerstone of any contemporary IT strategy. Based on my experience as a CFO, if you do not consider it you are likely wasting money directly or indirectly. Who can afford that?

*Good Luck,
Kevin*

*Mr. Kevin Weston, former CFO Digital Domain Productions, VP Finance and Operations of LucasArts
and CFO/SVP Finance and Operations at Eidos Interactive Inc.*