International Journal of Computer Science and Informatics

Volume 2 | Issue 3 Article 2

January 2013

Combining Server And Storage Virtualization: A New Dimension for SMB's

Mohammad Arif Baig
Advance Networking and Telecom Dept (SOIT), International Institute of Information Technology, Pune, India, mohammadb_aug10@net.isquareit.ac.in

Combining Server And Storage Virtualization: A New Dimension for SMB's

Mohammad Arif Baig

Advance Networking and Telecom Dept (SOIT), International Institute of Information Technology, Pune, India
E-mail : mohammadb_aug10@net.isquareit.ac.in

Abstract **-** The virtualization wave is quickly reaching its way down into the small-to-medium-sized business. Virtualization provides unmatched flexibility, performance, and utilization by allowing you to move server workloads from one virtual workspace to the next, maximizing server resources on the fly based on your business needs. Server virtualization eliminates the conventional, one application per server model and allows businesses to run multiple, virtual servers on a single physical machine. Storage virtualization helps the storage administrator perform the tasks of backup, archiving, and recovery more easily, and in less time, by disguising the actual complexity of the SAN. Storage and server virtualization are complementary technologies that helps to build a completely virtualized infra-structure. When used together, server and storage virtualization are intended to derive greater benefit from each technology than deployed alone.

Keywords **-** *SAN, NAS, SMB, ILM, HBA, iSCSI.*

I. INTRODUCTION

A. What is Virtualization?

Virtualization is a method of running multiple independent virtual operating systems on a single physical computer. It is a way of maximizing physical resources to maximize the investment in hardware. Since Moore's law has accurately predicted the exponential growth of computing power and hardware requirements for the most part have not changed to accomplish the same computing tasks, it is now feasible to turn a very inexpensive 1U dual-socket dual-core commodity server into eight or even 16 virtual servers that run 16 virtual operating systems. Virtualization technology is a way of achieving higher server density. However, it does not actually increase total computing power; it decreases it slightly because of overhead.

Virtualization is being used by a growing number of organizations to reduce power consumption and air conditioning needs and trim the building space and land requirements that have always been associated with server farm growth. Virtualization also provides high availability for critical applications, and streamlines application deployment and migrations. Virtualization can simplify IT operations and allow IT organizations to respond faster to changing business demands.

B. When to use virtualization

Virtualization is the perfect solution for applications that are meant for small- to medium-scale usage. Virtualization should not be used for high-performance applications where one or more servers need to be clustered together to meet performance requirements of a single application because the added overhead and complexity would only reduce performance, e.g. We're essentially taking a 12 GHz server (four cores times three GHz) and chopping it up into 16 750 MHz servers. But if eight of those servers are in off-peak or idle mode, the remaining eight servers will have nearly 1.5 GHz available to them [1].

C. Advantages of Using Virtualization

Today's IT intensive enterprise must always be on the lookout for the latest technologies that allow businesses to run with fewer resources while providing the infrastructure to meet today and future customer needs. Virtualization utilizing Intel Virtualization Technology is the cutting edge of enterprise information technology. Intel is closely working with VMware, XENSource, Jaluna, Parallels, tenAsys, VirtualIron, RedHat, Novell and other VMM developers.

C.1. Server Consolidation

It is not unusual to achieve 10:1 virtual to physical machine consolidation. This means that ten server applications can be run on a single machine that had required as many physical computers to provide the unique operating system and technical specification environments in order to operate. Server utilization is optimized and legacy software can maintain old OS

configurations while new applications are running in VMs with updated platforms.

C.2. Testing and development

Use of a VM enables rapid deployment by isolating the application in a known and controlled environment. Unknown factors such as mixed libraries caused by numerous installs can be eliminated. Severe crashes that required hours of reinstallation now take moments by simply copying a virtual image.

C.3. Dynamic Load Balancing and Disaster Recovery

As server workloads vary, virtualization provides the ability for virtual machines that are over utilizing the resources of a server to be moved to underutilized servers. This dynamic load balancing creates efficient utilization of server resources. Disaster recovery is a critical component for IT, as system crashes can create huge economic losses. Virtualization technology enables a virtual image on a machine to be instantly re-imaged on another server if a machine failure occurs.

C.4. Virtual Desktops

Multinational flexibility provides seamless transitions between different operating systems on a single machine reducing desktop footprint and hardware expenditure.

C.5. Improved System Reliability and Security

Virtualization of systems helps prevent system crashes due to memory corruption caused by software like device drivers. VT-d for Directed I/O Architecture provides methods to better control system devices by defining the architecture for DMA and interrupt remapping to ensure improved isolation of I/O resources for greater reliability, security, and availability [2].

II. SERVER VIRTUALIZATION

Server virtualization is the partitioning of a physical server into smaller virtual servers. In server virtualization the resources of the server itself are hidden, or masked, from users, and software is used to divide the physical server into multiple virtual environments, called virtual or private servers.

One common usage of this technology is in Web servers. Virtual Web servers are a very popular way of providing low-cost web hosting services. Instead of requiring a separate computer for each server, dozens of virtual servers can co-reside on the same computer.

Virtualization can drastically reduce the number of servers in a data center, thus decreasing electricity consumption and waste heat, and consequently the size of the necessary cooling equipment. Some investment in software and hardware may be required to implement virtualization, but it is usually modest compared to the savings achieved. There are three popular approaches to server virtualization: the virtual machine model, the paravirtual machine model, and virtualization at the operating system (OS) layer.

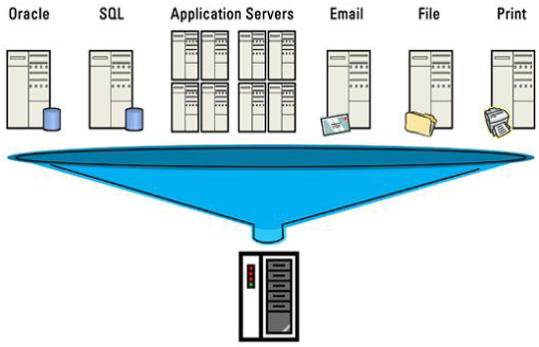

Fig. 1 : Virtualization of Servers

Benefits of Server Virtualization

Advances in Server Virtualization technology continue to highlight the many benefits of using server virtualization to efficiently consolidate servers, save physical space, provide more flexibility, and manage disaster recovery.

A. Consolidate Servers

The most benefit of server virtualization is the capability to consolidate applications on existing servers. Business applications usually don't require anywhere close to the resource capacity of their computers—perhaps an average of 20%. Several applications can be run on shared servers using virtual environments, allowing up to a 60-80% [3] resource utilization of a company's servers. The need for fewer servers results in significant savings in new server hardware spending, and additional savings are realized in lower maintenance costs for servers.

B. Conserve Physical Space

Using fewer servers also reduces the need for additional physical data center space. Limiting the size of the physical footprint means cost savings in heating, cooling, and electricity, as well as other maintenance needs for a facility.

C. Provide Flexibility

In the constantly changing IT environment, flexibility is crucial to maintaining older applications as well as developing and testing new applications. Different operating systems may also be required on the same hardware platform. Server Virtualization can allow a standard virtual server that can be easily duplicated to speed up server deployment, and can provide the

environment to run legacy systems along with the newer versions of applications. Programs in development can be easily tested in virtual environments, and the migration of applications can be accomplished without interrupting business.

D. Manage Disaster Recovery

Planning for better disaster recovery management is another benefit of server virtualization technology. Traditional tape backup systems have typically one or two days for a complete restoration. Virtualization could allow a restoration in less than a half a day because the systems and applications already exist in other untouched virtual environments and just need to be brought online if proper planning, updating, and testing have been done. If a system fails, it can be automatically switched to a standby server or brought back to its normal state with a virtual copy of the original image. Virtualization strategies vary, but the benefits continue to be impressive in reducing hardware requirements and physical space allotments, as well as allowing the maximum strategic flexibility for allocating resource and planning for efficient disaster recovery [4].

III. STORAGE VIRLUALIZATION

In its simplest form, storage virtualization pools all of your storage resources into a single logical entity; this should sound familiar since it's the same process you went through when you started virtualizing your servers.

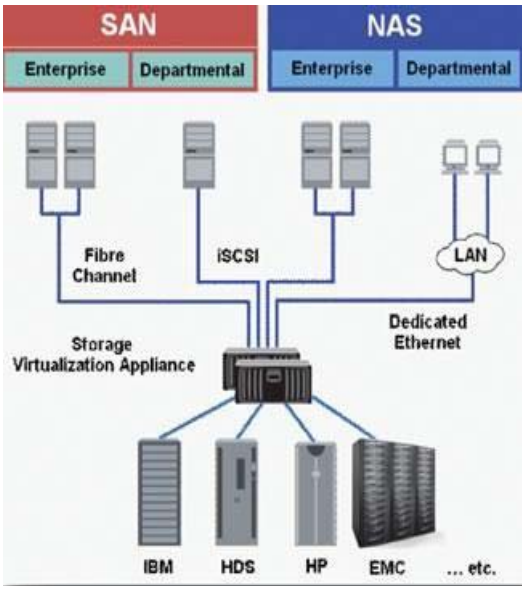

Fig. 2 : Storage Virtualization

Storage virtualization is an abstraction that presents servers and applications a view of storage that is different from that of actual physical storage, typically by aggregating multiple storage devices and allowing them to be managed in one administrative console.

Benefits of Storage Virtualization

A. Single point of administration:

Customers understand that a little friendly competition between storage vendors can help reduce the price of storage; SAN virtualization can be implemented in multi-vendor storage environments, so if a potential customer doesn't run the brand of storage equipment that you offer, converting them to a SAN virtualization setup could open the door to storage hardware sales. SAN-based storage virtualization provides this benefit by virtue of the fact that all of the SAN storage is provisioned to the virtualization device, so from that point on, all of the storage administration occurs at the virtualization layer.

B. Non-disruptive data migration:

Many customers do not replace their storage arrays when the lease or support expires or when the products are fully depreciated -- not because it is cheaper to keep the storage (the manufacturers make sure it's not), but because it is very difficult to migrate to the next storage platform. Without storage virtualization, migrations often require application outages and lots of sweat equity. These efforts are extremely difficult to coordinate across the various teams and business units involved. With SAN virtualization, the storage team can execute disk array swap-outs without impacting anyone else. This capability brings benefits to storage managers and VARs alike; migrations suddenly become not only possible but easy enough to make them worthwhile, allowing storage managers to take advantage of the declining cost of storage and bringing VARs an opportunity to sell new storage.

C. Information lifecycle management (ILM):

Customers want to know that they are putting their application data on the most appropriate tier of storage. Data access patterns are a key criterion in determining where to put the data. Often, a single business application or database has certain regions of data that are frequently accessed and require high-performance storage, while other regions are rarely accessed and could exist on more cost-effective storage. Because SAN virtualization tools sit between the server and the storage hardware, they have awareness of the access patterns. If the virtualization engine can use this access pattern information and leverage its online migration features, it may be possible to transparently relocate frequently accessed data to more expensive, high-performance storage and move less frequently accessed data to less expensive storage, bringing true ILM within reach.

D. Improved allocation efficiencies

Storage managers know that improving asset utilization is a quick way to lower the total cost of ownership (TCO) for their department. One of the common causes of low utilization is that the application teams demand more storage than they need. This may be because the process for requesting more storage is too slow; on the other hand, it could be because the application is new and there's not enough history to properly plan for growth. Storage virtualization promises to solve both problems. In the first case, the pace of deployment can be improved when all storage, regardless of brand or type, has a single administrative interface for allocation. And capacity planning challenges could be alleviated with thin provisioning services in the virtualization layer, which allow pre-allocation of storage and shared free space across applications optimizing unused disk, which is the most expensive storage asset.

E. *Heterogeneous replication*

One of the huge challenges associated with maintaining agnosticism among disk array vendors is disaster recovery replication. Most array-based storage replication is not heterogeneous, meaning that the production and disaster recovery frames must be of the same brand and often of the same type. Host-based replication options are heterogeneous, but management is cumbersome when a large number of hosts have replicated data. SAN virtualization can split the difference, providing a single method of replication for multiple types of storage arrays and a limited number of management points [5].

IV. INTEGRATING STORAGE AND SERVER VIRTUALIZATION FOR SMB's

Server virtualization projects are no longer being implemented in the largest data centers. The virtualization wave is quickly reaching its way down into the small-to-medium-sized business. In fact, the payoff for the SMB may be even greater than in the larger enterprise. SMBs, however, have typically had one significant disadvantage to the larger enterprise, accessibility to shared storage. Shared storage is the key in how the SMB can achieve server virtualization's full potential.

One of the prime focuses in any virtualization project, whether it's storage or server and especially both, is to conduct an inventory of the servers, storage devices and such that will be involved. This includes things such as the host bus adapters (HBA) and storage area network (SAN) switches, and the software and firmware revisions.

Check the hardware compatibility lists (HCL) for both virtualization products and make sure your configuration conforms. This is getting easier as virtualization vendors work to make their products interoperable. For example, VMware Inc., now owned by EMC Corp., is aggressively promoting its VMware Infrastructure 3, which ties VMware's ESX Server 3 and related products with storage virtualization, and associated hardware and software. Recently, both Emulex Corp. and QLogic Corp. announced that they now have HBAs that are supported by VMware's architecture [6].

The first step is to drive out the upfront costs. Shared storage in the enterprise has typically had a high cost because it meant purchasing proprietary storage from a single vendor and implementing that shared storage on a new network infrastructure like Fiber Channel. While the costs of these systems and their infrastructures have come down, the purchase of something new will typically raise the cost of a project. Second, because the storage is from a single vendor there is often, (always?), a premium price associated with it.

For the SMB it would be more cost-effective if they could leverage the investment they already made in servers, storage and network. A way to accomplish this is to use separate storage software from the storage hardware. All storage systems, whether they are designed for small business or large enterprise, have essentially two components; hardware and software. The software is the intelligence that manages sharing of the storage and other capabilities like replication or snapshots.

Another key aspect in making shared storage affordable for the SMB Virtual Server project is to similarly leverage the existing network infrastructure instead of implementing a new one. Most SMBs will have some form of IP network in place long before they begin a virtualized server project. As a result, a storage system that can be shared with the existing IP networks would be able to further keep storage costs down even if the storage network can use its own switch or networking gear to ensure high performance of the storage traffic.

There are two storage protocols that leverage an existing IP infrastructure. The first is the file sharing protocol NFS, which is commonly found in NAS environments (file-level storage). NFS, while supported by some server virtualization platforms, is not supported by all. If it were to be used, for most SMBs, the purchase of a specific NAS based appliance would be required. For hosting virtual machine images a fairly powerful, and even more expensive, NAS is also required.

The second protocol that leverages the existing IP network and is used for block-level storage (i.e. SAN) is iSCSI. iSCSI encapsulates standard SCSI commands and sends them across an IP network. This means that almost anything that can be done with a local SCSI hard disk can also be done shared, across a network with iSCSI, including clustering and booting from the environment. For the small-to-medium-sized business this may be the ideal situation. iSCSI allows the SMB to leverage the existing IP network that they already have.

Once the storage platform is built, the SMB has to be confident that they can implement it and operate the environment. Since this iSCSI environment is built on a networking protocol that they are likely to be very familiar with, it is as simple as continuing to manage that network. Other than centralizing the storage and making it shareable, nothing new has been added. The SMB simply has another component on their existing IP network [7].

V. HOW STORAGE VIRTUALIZATION CAN ENHANCE SERVER VIRTUALIZATION FOR SMB's

Server virtualization creates a dynamic environment where large numbers of virtual machines are applied to applications. Storage virtualization makes it very easy to allocate capacity to these servers and then reallocate capacity as server's needs change.

A. Easy Provisioning of Volumes to Virtual Machines

In general, the value of virtualization grows with the number of servers being managed. Typically, therefore, virtualized environments support tens or even hundreds of logical machines on a very few, large physical boxes. Therefore, the ability to virtualize the storage to these virtual servers becomes very attractive. If a user has hundreds of virtual machines, where each server requires about ten volumes, the number of volumes required would be in the thousands. Having a SAN-based Volume Manager (SVM) on the back-end of the server virtualization, allows users to quickly and efficiently create volumes for each of the virtual machines. It eliminates the need to deal with LUN management at the array level. Due to the fact that each volume is an independent volume, those volumes can be mirrored, replicated be the source of snapshots, and even mounted into standard servers without the VMware operating system if needed.

Storage virtualization provides the flexibility to rapidly allocate capacity, and the ability to allocate thousands of volumes – as needed. Due to the fact that in a virtual environment it is very easy to add and remove virtual machines and applications, the environment becomes very dynamic. Virtual machines are created, used and then reallocated or removed. This provides tremendous flexibility. In this highly dynamic environment, it is also important that the storage be provisioned and reallocated after use with the same simplicity as the virtual servers. Storage virtualization enables this flexibility.

B. Test environments

Test Environments are one of the killer applications for virtual servers. A storage virtualization technology that supports low-capacity, point-in-time snapshots can increase this advantage by enabling the rapid creation of multiple snapshot copies of production volumes and their assignment to virtual test environments. Additionally, snapshots can reduce data preparation time before each test. So, testers can be assigned real ―live‖ data within seconds and then take snapshots of the data throughout the testing process. Should a multi-stage test fail at say, stage 13, the tester could go back to the snapshot taken at the beginning of the stage and run the test again, eliminating the need to repeat the 12 previous tests. Additionally, since the real failure may have occurred earlier in the testing process, the user could go back to previous snapshots taken at each stage and view the data to determine the root cause of the failure. All of these features significantly reduce the time needed for testing and increase the productivity of the testing team. Bringing a product to market quicker or isolating a software bug quicker can improve the profitability of a company.

C. Enhanced Backup

Having a virtual environment with hundreds of virtual machines can create a complicated, expensive backup proposition. Snapshot functionality obviates the need to install backup agents on every virtual machine. The backup can be done by creating snapshot copies for every virtual server and then assigning the copies to a virtual machine with the dedicated role of backup server. In this manner, the backup server is the only virtual machine that needs the backup software. When dealing with hundreds of virtual servers, this can reduce the cost of backup licenses considerably.

With capacity growing at exponential rates and processing hours becoming more and more important, backup windows are becoming non-existent. Simply stated, there is too much data to backup during off-hours. By comparison, snapshots can be taken at any time without taking the application offline. This creates a zero-window for backup. For many users, this solves the "shrinking backup" window problem.

Snapshots can add another significant benefit to the overall backup strategy. It is possible to keep point-in-time snapshots online for extended periods of time. If the data needs to be restored, the restoration can be done

in seconds from a point-in-time snapshot rather than in hours off tape, reducing the recovery time objective (RTO) from hours to seconds.

One issue in using snapshots is ensuring point-in-time consistency between volumes. This often requires that the application be suspended for a time while (for example) database buffers are flushed and all I/O completed. More complex solutions such as the use of consistency groups should be evaluated. Snapshots will not usually eliminate the requirement to take a full backup of the systems on a regular basis.

D. Consolidation of Servers at the Disaster Recovery Site

Today, remote mirroring and disaster recovery is a requirement even for very small companies due to regulation, corporate policies, or simply common sense. Large enterprises typically have the resources to spend on necessary communication lines, equipment, software, and training for disaster recovery. However, small and medium-sized businesses do not always have these resources, leaving the company exposed to regional disasters. This is another area where the combination of storage and server virtualization software can enable an affordable solution for disaster recovery.

Over the life of a remote mirror implementation, the most expensive component is the communication lines between the source and target locations. A solution that can remotely mirror using a snapshot-based technique, where only the differences between the snapshots are transmitted, avoids the need for very expensive communication lines between the locations. It is possible to use T1 or T3 lines that often cost hundreds of dollars per month, rather than higher bandwidth lines that can cost tens of thousands of dollars per month.

A remote mirror is an insurance policy to make systems available should a regional disaster render your primary site unavailable. However, there is no need to dedicate resources until a failure occurs. A virtual server gives users the flexibility to assign just very small amount of resources required to accomplish the mirror to the remote site. With snapshot-based remote replication, those resources are minimal because of the reduction of the amount of data being transmitted to the remote site. Should a failure occur at the primary site, server virtualization could then be used to assign more virtual machines to support the production workload. Virtualization of the remote resources allows them to be used for other purposes such as testing while they are in "standby" waiting for a failure, while guaranteeing that they will be instantly available when that primary site failure happens. This is true even if the primary site uses "physical" servers.

Thus the powerful combination of storage and server virtualization, therefore, allows users to build a disaster recovery site at a fraction of the cost. While large enterprises use expensive communication lines (e.g., OC-3, OC-12), high-end arrays in both the local and remote sites, and dedicated servers in the remote site, virtualization allows the use of inexpensive communication lines, inexpensive storage arrays at the remote sites and the use of virtual servers instead of physical servers at the remote location [8].

VI. CONCLUSION

Virtualization is a great opportunity to lower cost and raise productivity while reducing risk for businesses of any size and with budgets as low as zero. Many technologies promise important improvements for businesses but most create questionable value while incurring real cost. Virtualization brings real, measurable value while often costing nothing and often reducing spending immediately.

Virtualization technology is becoming more and more prominent in IT environments. Both server virtualization and storage virtualization offer unique benefits. However, it is not until the two technologies are combined that users truly recognize the full benefits of the solution. The ability to scale both server and storage resources as needed is a tremendous benefit. Additionally, the combination of the two technologies, create new opportunities for value that previously did not exist. Storage and server virtualization, when combined provide full resource virtualization to capitalize on these new benefits.

REFERENCES

[1] George Ou, "Introduction to server virtualization", May 2006

[2] Thomas Burger, " Intel Technology Journal: Special issue on virtualization technology", Volume 10, Issue 03, Oct 2008

[3] Richard Talber,"Using Virtualization to Improve Datacenter Efficiency", July 2009

[4] [4] Anup Pal, ―Server Virtualization: Its Challenges and Benefits‖, Oct 2010

[5] Brian Peterson, "SAN-based storage virtualization: Five benefits for your customers", Feb 2008

[6] Rick Cook, "Integrating storage and server virtualization", July 2006

[7] George Crump, "Why SMBs Should Combine Server Virtualization and Storage Virtualization", November 2009

[8] Nelson Nahum, " Combining Storage and Server Virtualization", Jan 2009

◈◈◈