CLOUD COMPUTING AND VIRTUALIZATION

White Paper

JAN KREMER
CONSULTING SERVICES

TABLE OF CONTENTS

1. INTRODUCTION

This white paper provides an overview of Cloud Computing and the related Virtualization aspect of Cloud Computing.

For a detailed white paper on Cloud Computing see my white paper "Cloud Computing Private and Public"

1.1. Document Outline

Chapter 1 provides an introduction and outline of this document.

Chapter 2 provides an overview of "What is Cloud Computing and Virtualization".

Chapter 3 provides an overview of the relationship between Cloud Computing and Virtualization

Chapter 4 provides a more detailed overview of all aspects of Virtualization

2. What is Cloud Computing and Virtualization

2.1. Introduction

This Sections provides a high level overview of Cloud Computing and Virtualization

2.2. What is Cloud Computing?

Cloud computing is being heralded as the Next Big Thing in IT infrastructure, promising new levels of efficiency, flexibility and cost savings—particularly in the area of outsourced hosting, also known as infrastructure-as-a service (IaaS).

But because cloud computing is at the early stages of what Gartner calls the "hype cycle," there is widespread confusion about what the term actually means, not to mention questions about how this new technology can deliver practical business benefits.

One IT Web site recently collected 21 different definitions for cloud computing from various experts, but in the broadest sense the definition is simple. Cloud computing provides a remote service that users can access via the Internet. Under this broad definition, we are all familiar with cloud computing.

Facebook, Flickr and the various Internet-based e-mail offerings such as Yahoo! Mail and Google's Gmail all store and process data remotely, and we take this for granted. Because of the close association of these familiar services with cloud computing, some analysts and trade-press editors are beginning to refer to all Internet-accessible services as cloud computing. This is both confusing and inaccurate. Such services more properly fall into the category of software-as-a-service (SaaS).

It is more accurate to look at cloud computing as a new approach to infrastructure one that is a logical next step on a path to more efficient use of computing resources.

The search for economic efficiency in IT infrastructure is a relatively new concern. Until recently, the only energy-related question that data-center managers worried about was whether enough of it would be available on a reliable basis.

Three factors have changed this mindset:

❖ Rising energy costs

❖ A trend toward IT and Facilities Management sharing responsibility—and budget—for data-center energy usage

❖ Concern about the general ecological impact of IT, and its contribution to companies' carbon footprint in particular

Increasingly, data-center managers must meet not only SLAs but "performance per watt" goals as well. Two technologies have already emerged to meet this goal: virtualization and grid computing (also known as utility computing). Both embody principles that are central to cloud computing.

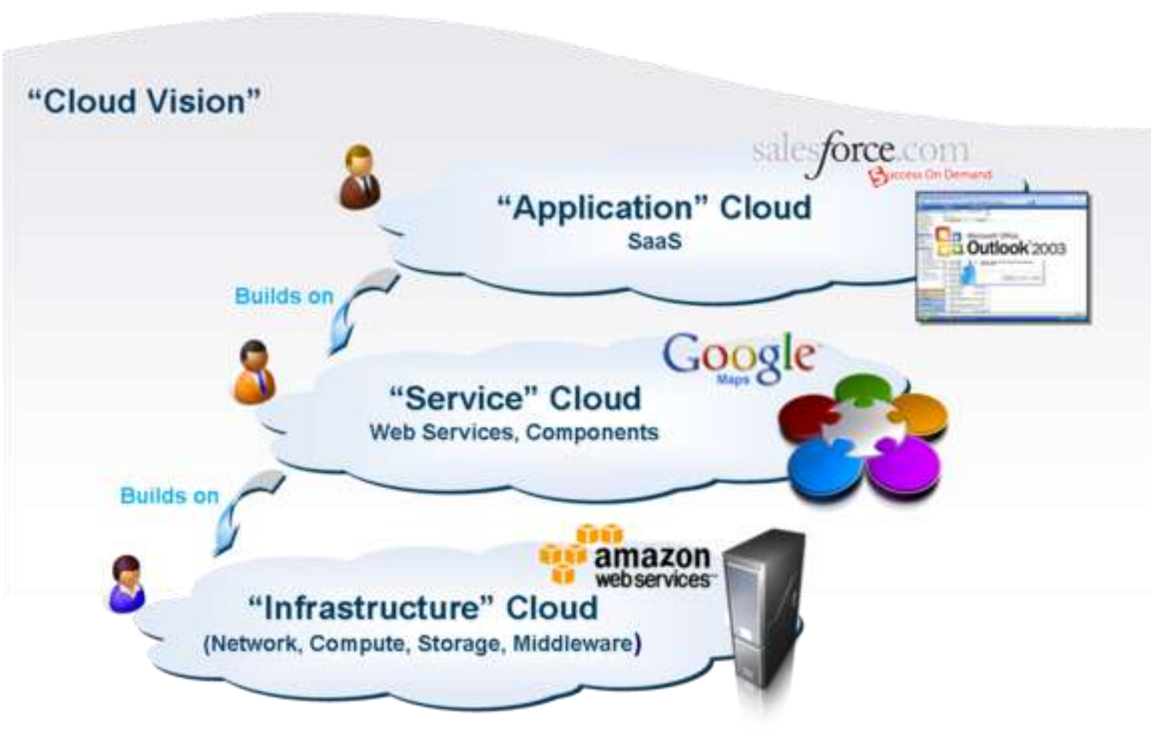

2.3. What is Virtualization?

Server virtualization was one of the first and most important responses to demands for IT energy efficiency. It is specifically a response to the one-application-per-server mentality that has been prevalent in data centers for many years and that result in a situation where only 15 or 20 percent of their total computing capacity is in use at any given time. By enabling IT departments to run multiple applications on the same server, virtualization provides dramatic gains in server utilization. It is also an approach that non-technical senior managers can understand and may even demand (often without realizing that virtualization can be more diff cult in practice than in theory).

Although virtualization clearly increases server resource utilization, it is not the ultimate answer to computing efficiency because it does not deal with the problem of usage spikes, such as the increased demand put on financial applications at the end of each financial quarter or the huge (but short-lived) spikes in e-commerce traffic following the airing of an infomercial. For example, if application X exceeds the processing capability of the initially provisioned and available hardware during a peak transaction period, the application suffers a performance hit because virtualization cannot dynamically add additional hardware components and additional virtual machines on demand.

This is where technologies like VMware DRS (Distributed Resource Scheduler) come into play. VMware DRS can aggregate resources from multiple servers into pools of resources. By doing so, DRS can intelligently allocate available resources among virtual machines according to business needs. Highly demanding, short-lived applications will have access to all the necessary server resources when needed without isolating those resources specifically to those applications. When those highly demanding, short-lived applications are not demanding excessive resources, those resources can be dynamically reallocated to other virtual machines as needed.

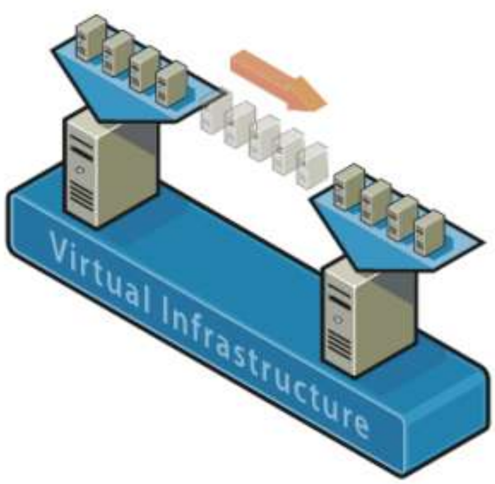

3. Cloud Computing and Virtualization

3.1. More about Cloud

Cloud computing is both a business delivery model and an infrastructure management methodology. The business delivery model provides a user experience by which hardware, software and network resources are optimally leveraged to provide innovative services over the Web, and servers are provisioned in accordance with the logical needs of the service using advanced, automated tools. The cloud then enables the service creators, program administrators and others to use these services via a Web-based interface that abstracts away the complexity of the underlying dynamic infrastructure.

The infrastructure management methodology enables IT organizations to manage large numbers of highly virtualized resources as a single large resource. It also allows IT organizations to massively increase their data center resources without significantly increasing the number of people traditionally required to maintain that increase.

For organizations currently using traditional infrastructures, a cloud will enable users to consume IT resources in the data center in ways that were never available before. Companies that employ traditional data center management practices know that making IT resources available to an end user can be time-intensive. It involves many steps, such as procuring hardware; finding raised floor space and sufficient power and cooling; allocating administrators to install operating systems, middleware and software; provisioning the network; and securing the environment. Most companies find that this process can take upwards of two to three months. Those IT organizations that are re-provisioning existing hardware resources find that it still takes several weeks to accomplish. A cloud dramatically alleviates this problem by implementing automation, business workflows and resource abstraction that allows a user to browse a catalog of IT services, add them to a shopping cart and submit the order. After an administrator approves the order, the cloud does the rest. This process reduces the time required to make those resources available to the customer from months to minutes.

The cloud also provides a user interface that allows both the user and the IT administrator to easily manage the provisioned resources through the life cycle of the service request. After a user's resources have been delivered by a cloud, the user can track the order, which typically consists of some number of servers and software, and view the health of those resources; add servers; change the installed software; remove servers; increase or decrease the allocated processing power, memory or storage; and even start, stop and restart servers. These are self-service functions that can be performed 24 hours a day and take only minutes to perform.

3.2. Virtualization

In the 1990s, the concept of virtualization was expanded beyond virtual servers to higher levels of abstraction first the virtual platform, including storage and network resources, and subsequently the virtual application, which has no specific underlying infrastructure. Utility computing offered clusters as virtual platforms for computing with a metered business model. More recently software as a service (SaaS) has raised the level of virtualization to the application, with a business model of charging not by the resources consumed but by the value of the application to subscribers. The concept of cloud computing has evolved from the concepts of grid, utility and SaaS.

It is an emerging model through which users can gain access to their applications from anywhere, at any time, through their connected devices. These applications reside in massively scalable data centers where compute resources can be dynamically provisioned and shared to achieve significant economies of scale. Companies can choose to share these resources using public or private clouds, depending on their specific needs. Public clouds expose services to customers, businesses and consumers on the Internet. Private clouds are generally restricted to use within a company behind a firewall and have fewer security exposures as a result.

The strength of a cloud is its infrastructure management, enabled by the maturity and progress of virtualization technology to manage and better utilize the underlying resources through automatic provisioning, re-imaging, workload rebalancing, monitoring, systematic change request handling and a dynamic and automated security and resiliency platform. How can we adopt cloud computing technologies and management techniques to improve the efficiency and flexibility of their own data centers and other computing environments?

Data centers will be virtualized, efficiently managed centers, which will employ some of the tools and techniques adopted by Web-centric clouds, generalized for adoption by a broader range of customers and enhanced to support secure transactional workloads. With this highly efficient and shared infrastructure, it becomes possible for companies to respond rapidly to new business needs, to interpret large amounts of information in real time and to make sound business decisions based on moment-in-time insights. The data center that supports a dynamic infrastructure is an evolutionary new model that provides an innovative, efficient and flexible approach in helping to align IT with business goals.

Virtualization refers to the abstraction of logical resources away from their underlying physical resources in order to improve agility and flexibility, reduce costs and thus enhance business value. In a virtualized environment, computing environments can be dynamically created, expanded, shrunk or moved as demand varies. Virtualization is therefore extremely well suited to a dynamic cloud infrastructure, because it provides important advantages in sharing, manageability and isolation (that is, multiple users and applications can share physical resources without affecting one another). Virtualization allows a set of underutilized physical servers to be consolidated into a smaller number of more fully utilized physical servers, contributing to significant cost savings. There are many forms of virtualization commonly in use in today's IT infrastructures, and virtualization can mean different things to different people, depending on the context. A common interpretation of server virtualization is the mapping of a single physical resource to multiple logical representations or partitions. Logical partitions (LPARs) and virtual machines (VMs) are examples of this definition;

How does server virtualization work? In most cases, server virtualization is accomplished by the use of a hypervisor to logically assign and separate physical resources. The hypervisor allows a guest operating system, running on the virtual machine, to function as if it were solely in control of the hardware, unaware that other guests are sharing it. Each guest operating system is protected from the others and is thus unaffected by any instability or configuration issues of the others. Today, hypervisors are becoming a ubiquitous virtualization layer on client and server systems. There are two major types of hypervisors: bare-metal and hosted hypervisors. *Bare-metal hypervisors* A bare-metal hypervisor runs directly on server hardware to provide virtual machines with fine-grained timesharing of resources.

3.3. Cloud and Virtualization

One of the initial steps toward cloud computing is incorporating virtualization, which is separating the hardware from the software. In the past, transitions of this magnitude meant rewriting code, such as the transition from the mainframe to UNIX. Fortunately, the transition to VMware does not require the rewrite of code, and this has fueled the speed of the move toward virtualization software. There still will be challenges in this transition but, overall, the consolidation of servers into the virtual world has been fairly rapid with many applications making a seamless transition.

The journey to get to cloud computing begins with virtualization with the cloud OS providing infrastructure and application services. The infrastructure services are the ability to virtualize server, storage, and the network, as well as application services that provide availability and security for the applications that are being utilized in the cloud environment.

VMware's vSphere satisfies the initial step of virtualization, the separation of the hardware and the software. The next step is adding some of the many cloud applications that include how to do charge-backs and other application software.

These cloud-like capabilities include billing for usage, the ability to do self-service, and many others. Charging for consumption, even if it is internal, will lead to better management, with the ability to keep track of what services the consumer is utilizing. In addition, with cloud computing, there is the ability to program in more self-service by the end user in order to keep costs down.

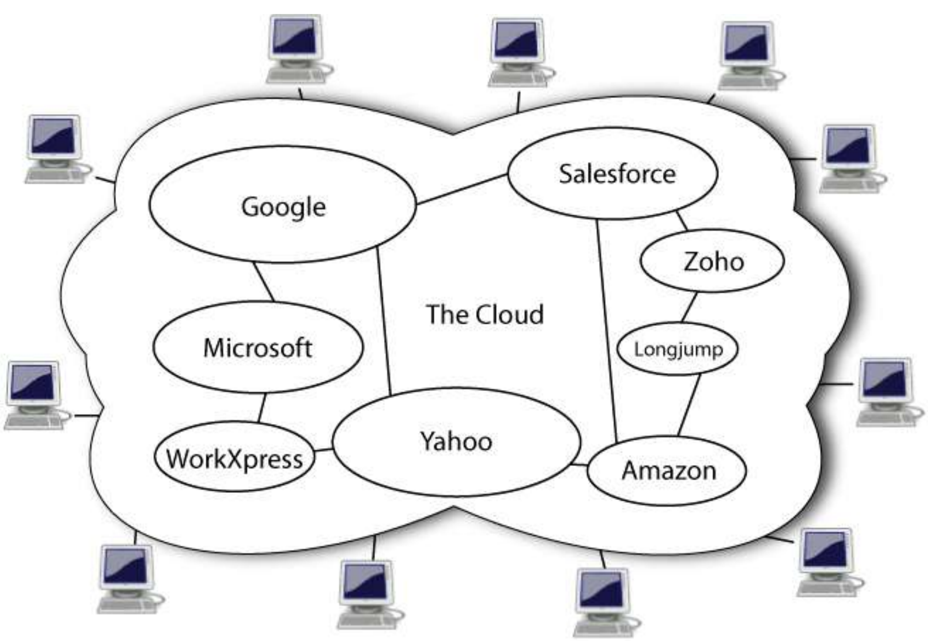

Cloud computing requires the use of virtualization, which is the separation of the hardware and the software using virtualization software such as VMware's vSphere.

Defining the different types of cloud computing provides us with knowledge as to what cloud computing has to offer. Whether you are a consumer or producer will define your definition of cloud computing. The public cloud is really geared more for the individual consumer or small company, while the private cloud is geared more for a medium-to large-company. In addition, the private cloud is branching out to incorporate the ability to have some data and applications serviced from the public cloud.

Cloud computing relies heavily on virtualization. The services are built on top of a virtualization layers which help the service providers to manage the service and offer standardized platform to the users. Virtualization is in fact another key element of cloud computing, it enables the service provider to offer the homogeneous service simultaneously to all customers, something that cannot be achieved, for example, in grid computing.

"Wilber is probably taking this Cloud computering too seriously.

3.4. Summary

How can virtualization and cloud computing be used to lower total cost without sacrificing performance, availability, security, and operational efficiency in the data center? What are the design considerations surrounding network convergence, which is often positioned as the network architecture for server virtualization?

Cloud computing and server virtualization strategies abound. Cloud computing can be private, public, or a hybrid of the two. Although this paper briefly discusses public cloud computing, the focus is on cloud computing in the data center. Virtualization can be applied to servers, storage, networks, and desktops. Mainframe and UNIX server virtualization technologies are mature and widely deployed. But x86 server virtualization is emerging rapidly and this paper discusses the impact of x86 server virtualization on the data center network.

Network convergence was originally applied to the convergence of voice, video, and data on the same telco network. But when network convergence is applied to the data center, it means transporting IP network and block-level storage traffic on the same physical network.

Although this is a separate topic from virtualization and cloud computing, some have positioned it as essential for virtualization to succeed. Two high-level architectures are presented, one with network convergence and one without. Physical network convergence can introduce complications for availability, security, and operations if it is applied indiscriminately and without careful attention to design considerations.

This paper discusses the key considerations for engineering a network for converged traffic, so that the network designer can make an informed decision about how best to meet the goal of lower total cost without compromising business requirements for performance, availability, security, and operational efficiency.

Server virtualization and cloud computing highlights immediate access to unlimited amounts of computing, storage, and network bandwidth.

Virtualization and cloud computing cannot exist without data centers and the physical hardware they house. To support virtualization, the data center architect has to harden the network against failures and make it adaptable and flexible without disrupting traffic and to do this while continuing to support existing data center assets.

A key conclusion is this: physical assets needed to support virtualization largely exist today. They will need to be extended to support server virtualization software where it makes sense and when it makes sense. *Doing so does not require upgrading the entire data center network from the access layer to the core.*

While the technology for virtualization is not new, the scale at which it is being deployed is unprecedented. The concept of a virtualized data center means that every aspect of every piece of hardware is abstracted from every other piece of hardware. Moving data from one array to another, moving applications from one server to another, moving network services from one switch to another essentially all add, move, or change operations—could be conducted without applications or users knowing anything about it. It is possible to build an abstraction layer that covers the entire data center itself. Solutions exist today to support nearly seamless application mobility across inter-state distances without disrupting applications or users being aware of this movement.

Virtualization means that the services provided by a hardware device are abstracted from the physical hardware. For example, if a disk provides the service of "storing data" for a server, the virtualization abstraction layer separates the server's access to the "storing data" service from the physical disk on which the data is stored. This allows the disk to be serviced or upgraded to increase capacity without disrupting the application while it is storing data. The "storing data" service the storage performs persists even when the hardware supporting the service changes.

A sound infrastructure for the next-generation data center should be robust enough to support virtualization of every component and the data center itself. Not all applications require or even benefit from every type of virtualization. Therefore, the infrastructure must be flexible enough to support both applications on dedicated hardware and applications on virtualized hardware. This implies that any network port can be configured for performance (bandwidth), security policies, and oversubscription ratios dynamically via management software.

In cloud computing, applications, computing and storage resources live somewhere in the network, or cloud. User's don't worry about the location and can rapidly access as much or as little of the computing, storage and networking capacity as they wish—paying for it by how much they use—just as they would with water or electricity services provided by utility companies.

Some features that apply to cloud computing are:

• Virtual infrastructure to provide resources. The data center itself becomes a dynamic pool of resources, enabled by virtualization technology. Applications are not constrained to a specific physical server and data is not captive to a single storage device. Operations focus on ensuring that adequate resources are available; the function of service provisioning handles what resources are allocated to an application or user.

• Service provisioning. Services must be provisioned with little or no effort on the part of the IT group responsible for maintaining the resource pools. Self-service portals that users can access let them request computing, storage, and network connectivity—all provided within minutes. This is a significant departure from the earlier IT model of project-based application deployment on dedicated hardware.

• Payment at time of use. Cloud computing supports quite a few innovative financial models, such as pay-as-you-go based on the resources used, and even a no-cost model in the case of some public cloud applications in which advertising pays for the infrastructure.

Cloud computing over the Internet is commonly called "public cloud computing." When used in the data center, it is commonly called "private cloud computing." The difference lies in who maintains control and responsibility for servers, storage, and networking infrastructure and ensures that application service levels are met. In public cloud computing, some or all aspects of operations and management are handled by a third party "as a service." Users can access an application or computing and storage using the Internet and the HTTP address of the service. Google Apps is a well-known example of public cloud computing, in which virtualization resides between the Internet connection and the data centers delivering the Google Apps service.

Clearly, public cloud computing is at an early stage in its evolution. However, all of the companies offering public cloud computing services have data centers; in fact, they are building some of the largest data centers in the world. They all have network architectures that demand flexibility, scalability, low operating cost, and high availability. They are built on top of products and technologies supplied by Brocade and others network vendors. These public cloud companies are building business on data center designs that virtualize computing, storage, and network equipment—which is the foundation of their IT investment.

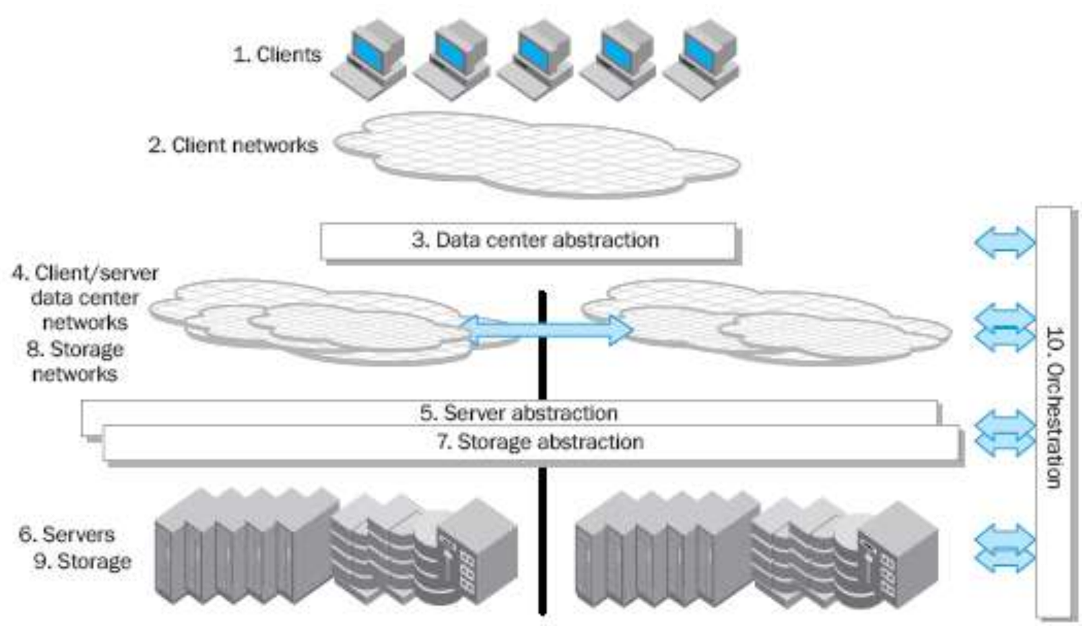

4. More about Virtualization

4.1. Introduction

Virtualization of business applications allows IT operations in companies of all sizes to reduce costs improve IT services and manage risk. The most dramatic cost savings are the result of reducing hardware, space and energy usage, as well as the productivity gains that also lead to cost savings. Service improvements include high availability of x86 application workloads and rapid provisioning of new services to meet the dynamic needs of the business. Virtualization can also mitigate risk to business operations. For example, when faults are detected or server load is too high, workloads can be moved out of harm's way.

The virtualization concept is expanding into many aspects of information technology. Servers, switches, storage, networking, and clients are all on a virtualization roadmap. But the virtualization movement is rooted in servers, and server virtualization will have the most profound impact on datacenter networks.

Server virtualization counters the trend to use purpose-specific appliances as servers. That trend succeeded because appliances are effective and easy to deploy. But in large numbers, appliances are inefficient. Because appliances are optimized for one function, many are dead-end devices subject to forklift upgrades. Perhaps worse, the proliferation of appliances has led to growing management burdens, sprawling datacenter networks, and high thermoelectric loads.

The solution to the underutilization problem is virtualization, and the essence of virtualization is an abstraction layer of software called the Hypervisor. The Hypervisor sits between hardware and the operating system. Virtualization allows multiple operating systems and applications to cohabitate on a physical computing platform. The graphic below shows virtualization supporting three logical servers on one platform. Server virtualization, especially when coupled with blade technology, increases computing and storage density while making IT assets more flexible.

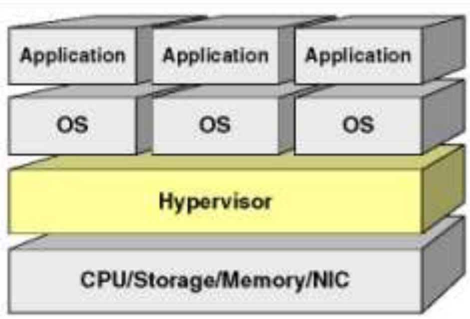

There are several elements of "Virtualization", such as:

- Server
- Storage
- Network

Virtualization is supported at the hardware (such as Intel) and software levels (VMware etc.)

Virtualization is the use of hardware and software technology to present a logical view of resources. This logical view is often strikingly different than the actual physical view. What does this really mean? System users may see the image of many different computers even though it is a single system. They may see many individual systems as a single computing resource. Individuals can be allowed to access computing solutions with devices that didn't exist when developers created the applications. Applications may appear to use devices that have long been considered obsolete even though none are actually present

There are many layers of technology that virtualize some portion of a computing environment. Each of these tools can be applied to making industry standard systems part of a larger, more efficient, more productive computing environment. Here's a graphical representation of those tools.

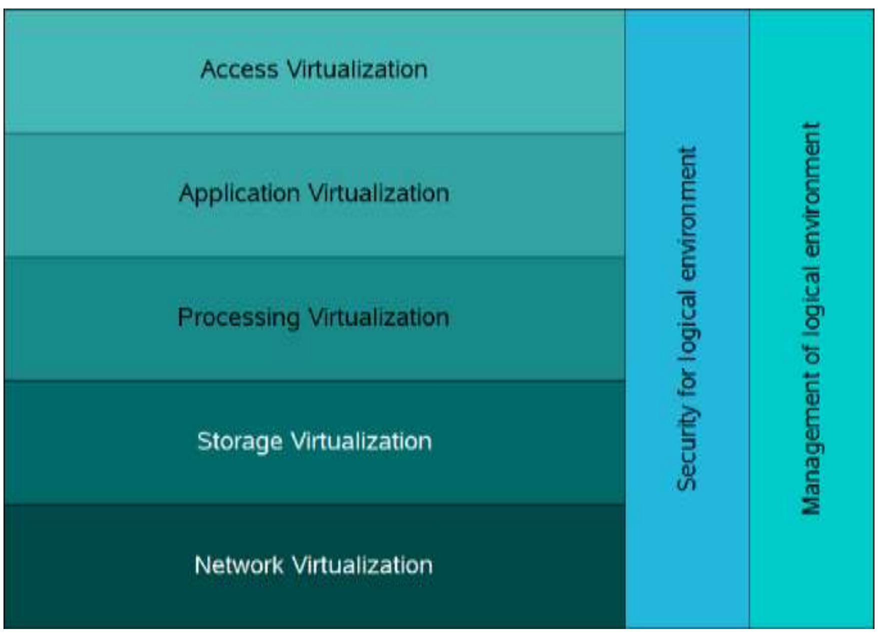

Processing virtualization hides the physical hardware configuration and makes it possible, on the one hand, to present a single system as if it were many or, on the other hand, to present many systems as if it were a single resource. Tools such as Intel® Virtualization Technology (Intel® VT) really help with the tasks found in this layer.

Suppliers, such as Intel, are focusing a great deal of investment on virtualization technology at all levels of the model. They are also working with suppliers of systems, operating system software, data management software, application development as well as the suppliers of application development framework software in order to offer organizations a highly optimized set of virtualization solutions at the lowest possible cost. Through the efforts of Intel to increase virtualization optimization and decrease power consumption, IT managers have the capability to increase their overall system utilization while decreasing costs by 50% or greater. Here's a few of the likely improvements virtualization technology will provide in the near future.

- Optimal use of an organization's systems will be assured because applications, application components and data will be moved to the most appropriate environment on a moment by moment basis
- Organizations will find it much easier to add processing power as needed to meet their own service level objectives
- New technology will co-exist and work efficiently with more established technology.
- Applications will be accessible from nearly any type of network-enabled device, over just about any network, form nearly anywhere without organizations being forced to re-implement or re-design their applications
- Application performance, scalability and reliability will increasingly be built into the environment rather than being a matter of tedious or complex design
- Applications and data will be increasingly secure and protected thus removing the fear IT management has of security breaches, malicious Email messages and the like.
- Individual software developers will no longer have to care which system is working for them, where it is located or what type of software is supporting them. They'll be able to focus on the task at hand rather than being asked to take on the role of system operators.

Virtualization technology, combined with advances in processor, I/O and other technology will make it possible for high volume, industry standard systems to be deployed for everything from decision support to collaborative environments to high volume transactional applications to high performance modeling

4.2. Server Virtualization

Few technologies have become a fundamental part of the datacenter as quickly as server virtualization. That's because the basic value proposition is so easy to grasp: when you run many logical servers on a single physical server, you get a lot more out of your hardware, so you can invest in fewer physical servers to handle the same set of workloads.

The details, of course, are more complicated. The hypervisor, a thin layer of software upon which you deploy virtual servers, is generally wrapped into a complete software solution that incurs some combination of licensing, support, and/or maintenance costs (depending on which virtualization software you chose). And very likely you will need to upgrade to server processors that support virtualization.

On the other hand, reducing the number of servers yields indirect cost savings — less space to rent, less cooling to pay for, and of course lower power consumption. Even more compelling is virtualization's inherent agility. As workloads shift, you can spin up and spin down virtual servers with ease, scaling to meet new application demands on the fly.

The path to rolling out a virtualized infrastructure has its share of pitfalls. You need to justify the initial cost and disruption in a way that does not create unrealistic expectations. And you need to know how to proceed with your rollout, to minimize risk and ensure performance stays at acceptable levels.

It's pretty easy to sell server virtualization. Who doesn't want to get the most possible use out of server hardware? in fact, the basic idea is so compelling, you need to be careful not to oversell. Make sure you account for the likely capital equipment, deployment, training, and maintenance costs. The real savings achieved by virtualization, as with so many other new technologies, tend to accrue over time.

Most virtualization deployments require new hardware, mainly because hypervisors require newer processors that support virtualization. so the best time to roll out virtualization is when you need to add servers to your existing infrastructure or when it's time to replace aging hardware.

The superior efficiency of newer servers will help make your case. Begin by calculating the power consumption and cooling levels the current infrastructure requires. (ideally, this should be conducted on a server-by-server basis, which can be time consuming, but will result in far more accurate numbers) then check the same specs for the hardware you plan to buy to get an idea of any power and cooling cost savings.

Add the fact that you will be using fewer physical servers for the same workloads, and your proposed virtualized infrastructure will look very, very good compared to the existing one. if the new hardware is sufficiently powerful, you may be able to run many logical servers on each physical unit.

Unfortunately, determining how many virtual servers will fit on a physical host is never an exact science. But there are tools that can help. Some server consolidation tools will even allow you to specify the make and model of your current and planned hardware, and will monitor your existing infrastructure for a period of time.

Armed with all that data, you can run reports that show exactly how many virtualization hosts you'll need, what type, and your expected ratio of virtual servers to physical hosts. Some will even calculate the expected power consumption and cooling capacity for the new infrastructure.

Investigate the options available from VMware, Microsoft, and others in order to get the most accurate data before you leap into any virtualization project.

One major benefit of embarking on a virtualization project is that it gives it the opportunity to jettison old hardware and old frameworks. There's never a better time to inspect the whole infrastructure and identify components that have fallen through the cracks, aren't necessary anymore, or can easily be folded into other tools or projects.

As you step through the planning stages of virtualization, you should pay close attention to anything that can be culled from the back-room herd without too much pain. it will ease the transition and cut down on the number of servers that need to be migrated or rebuilt on the virtualized foundation.

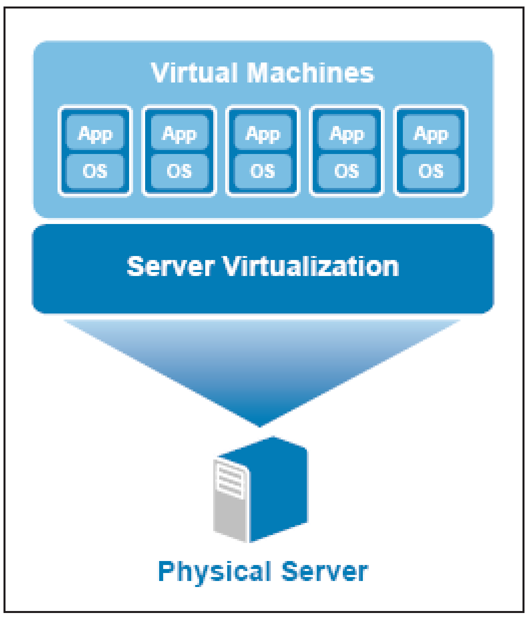

The obvious leader in virtualization platforms is VMware. They've been in the game the longest and easily have the most mature x86-based virtualization solution on the market. They're also expensive --you pay for all of those accolades.

The other major players are citrix's Xenserver and Microsoft's hyper-V. citrix's solution is based on the open-source Xen virtualization project and is a well-rounded offering, with features such as live migration.

Microsoft's hyper-V is arguably the cheapest option, but is well behind the other players in terms of features and robustness.

If you're a purely Microsoft shop and are looking to virtualize only a handful of windows servers, hyper-V may be attractive. If you're looking at deploying a larger virtualization environment that comprises a few different operating systems or more than a dozen or so virtual servers, you'll be better served looking elsewhere.

Eventually Microsoft may catch up to the competition in terms of features, management, and stability, but it's not quite there yet.

The other option is to roll your own. This method is certainly not for the faint of heart, and involves the use of several open-source tools, such as the open-source version of Xen, possibly the use of linux's KVM virtualization tools, or the use of VMware's free esXi embedded hypervisor. This is the cheapest solution by far, but it also requires highly skilled administrators and precludes many enterprise-level features such as live migrations and centralized management.

If you're going to be strictly virtualizing Linux servers and have the requisite skills available, this may be a good option. If not, it's definitely better to choose a commercial solution.

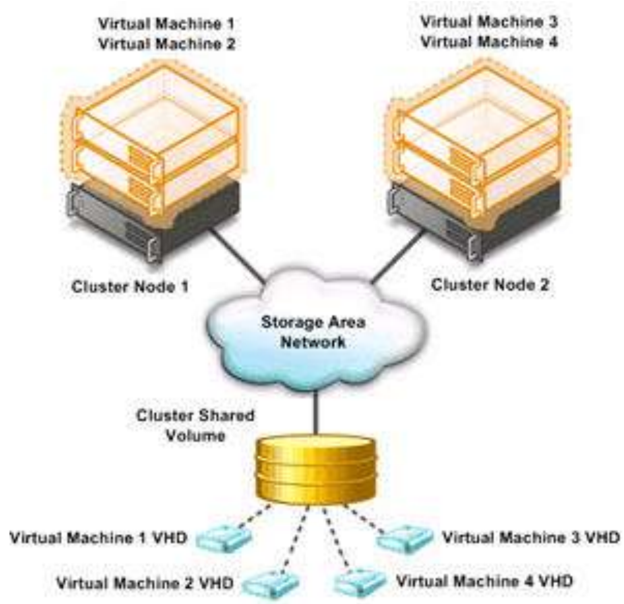

When doing due diligence for a proposed virtualized infrastructure, don't forget the applications. It may seem like a no-brainer that a given application will function on a virtual server, but you may run into problems with licensing and support. Although it's not as prevalent now as in the past, some ISVs refuse to support their applications if they're run on virtual servers.

In many cases this is somewhat of a cop-out, since there's no technical reason their products would have a problem with a virtual server, but that doesn't get the problem fixed when on the phone with product support. Make sure that your critical apps are okay to be virtualized before you make the cutover.

It's not just smaller applications that have problems like this. Microsoft doesn't support several of its key applications on some or all virtualization platforms. Also, there may be problems with keyed software that requires the presence of specific USB hardware license keys or other dongles to function.

Not all these issues can be sorted out, so make absolutely certain you're in the clear with your vendors before trying to pull them into the virtual pile.

After all the meetings, decisions, plans, budgets, and purchases are complete, the hard work starts: migrating.

All kinds of tools are available to assist in physical-to-virtual migrations and most of them are sufficiently mature to handle most tasks. However, performing a physical-to-virtual migration should not be the first resort, it should be the last.

If at all possible, rebuild your servers in the virtual environment and migrate the data, applications, and configurations. While physical-to-virtual tools can certainly work wonders, they can also cause odd problems down the line, especially with windows servers.

It should go without saying that you shouldn't migrate windows domain controllers. They're relatively simple to build and integrate into an Active Directory domain, and simple to demote.

Linux servers should also be rebuilt whenever possible, though physical-to-virtual migrations of Linux servers are generally more successful than windows migrations. Migrating applications and functions between Linux servers is generally far simpler than performing the same operations on windows.

In most cases, it's quicker to perform a fresh installation of Linux on a virtual server and then reinstall and reconfigure than it is to run a physical-to-virtual tool. If you're using templated virtual servers, it's even quicker.

Server virtualization can help increase IT flexibility and reduce equipment and labor costs by consolidating applications and workloads onto fewer physical servers. Rather than using five separate physical servers for five distinct applications, enterprises can employ VMware ESX Server™ virtualization software to run all five applications on the same physical server. On that physical server, administrators create distinct virtual machines that host operating systems and applications. Each virtual machine is completely isolated from the others and operates as if it has a dedicated set of resources. But in fact, the virtual machines share the same physical resources. Virtualization management tools such as VMware VirtualCenter™ enable IT administrators to dynamically adjust the size and number of virtual machines, as computing needs change.

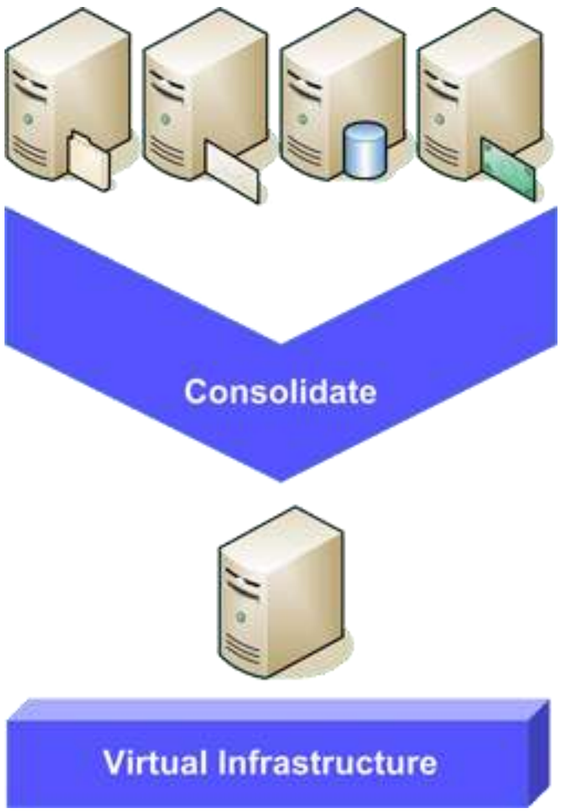

4.3. Storage Virtualization

Server virtualization offers several important advantages, but to realize the full benefits of server virtualization, enterprises must overcome significant storage challenges. In many IT infrastructures, those challenges arise because storage is directly attached to servers and cannot be shared beyond the physical server. Storage area networks (SANs) enable servers to share centralized resources but virtualizing storage on a SAN is only the first step.

Traditional SAN offerings can experience some of the same stumbling blocks as direct attached storage in a virtualized environment. In fact, some of the shortcomings of direct attach storage and traditional SANs are actually exacerbated by server virtualization. When virtualized server environments are coupled with traditional storage, it is more difficult for administrators to tailor storage requirements for each application. An application might benefit from a particular server interface, drive technology or speed, RAID configuration, or snapshot schedule. But in a virtualized server environment with direct attach storage; all of the applications running on the physical server are forced to use the same storage, with the same storage characteristics. Traditional SANs do not provide a solution. They offer only a limited number of volumes which are not enough to accommodate each of the large number of applications that can run on a virtualized server.

To take full advantage of server virtualization, IT departments need a storage solution that allows fully shared resources, easy and automated storage classification and migration and complete data protection.

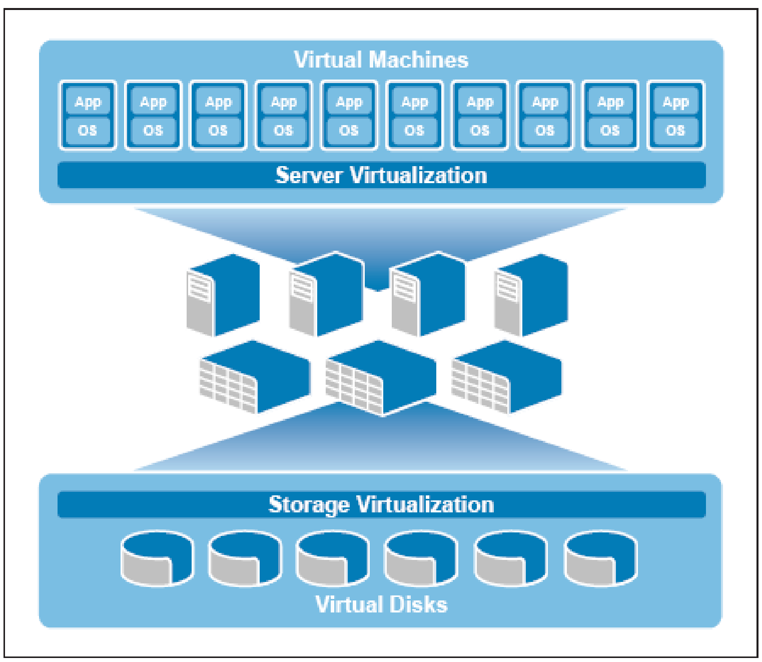

One of the benefits of consolidating direct attached storage into a storage network is the addition of advanced management and protection capabilities.

Without storage virtualization, the ability to provision, manage and back up or replicate storage resides in a storage controller, which supports a particular storage platform. Storage virtualization moves that logic into the storage network so services can be applied across multiple storage platforms, from different product families or even different vendors.

Storage virtualization falls into two categories: block-based solutions and file-based solutions. Both provide increased efficiency and flexibility in managing storage, and both can simplify data access for users.

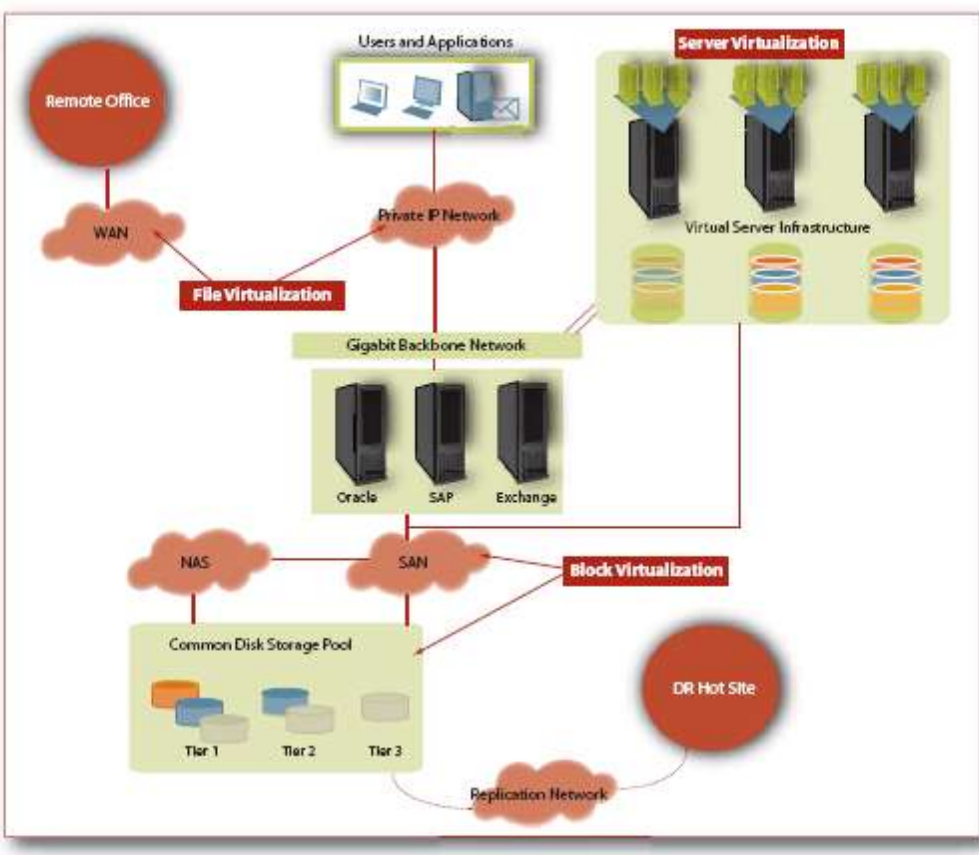

The primary benefits of block-based virtualization include:

- Heterogeneous volume management
- Heterogeneous data movement
- Heterogeneous copy and data protection

File-based virtualization can provide:

- A unified file system that remains unchanged for users when underlying physical storage is changed
- Global namespace allowing multiple NAS or Fibre Channel-attached storage devices to appear as a single file system
- Today, a majority of large enterprises have implemented server virtualization, and a growing number, perhaps as many as half, have implemented block-level (SAN) virtualization. Only a small minority are currently implementing file level virtualization.

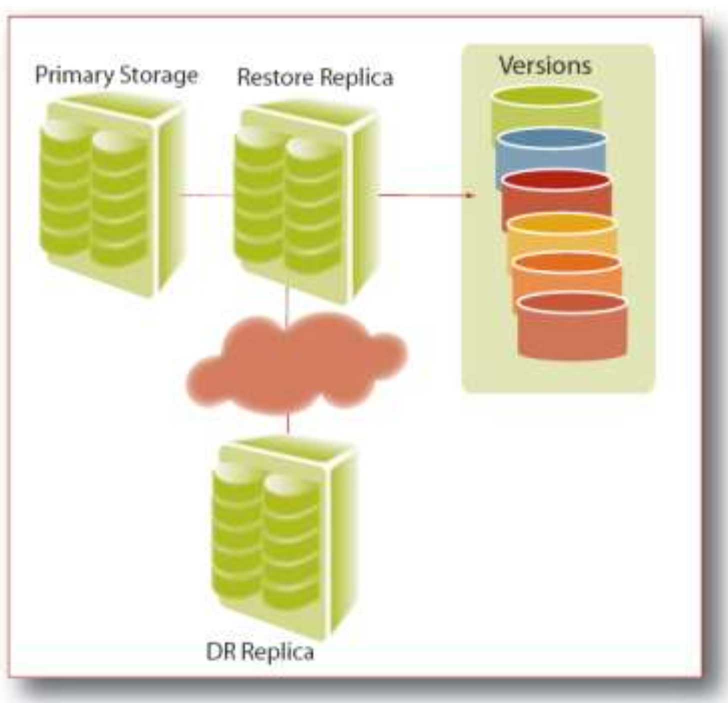

The adoption of virtualization technologies can help drive increased data protection.

Several block-based methods exist

Block-based virtualization moves logic from the individual storage device utilizing a higher layer of logic. This essentially makes the individual storage devices invisible to the user. Instead, blocks of information are addressed via LUNs or logical unit identifiers.

Block based storage virtualization can take several forms.

1) Storage controller-based

Virtualization within the storage controller is homogeneous and limited to systems provided by the vendors whose implementation is being used.

Pro: Works well in a homogeneous environment

Con: Not suited for heterogynous architectures

2) Server- or host-based virtualization

Storage is virtualized in the network by adding storage management software to an application server. Application servers and physical storage are tightly bound.

Pro: Leverages storage from heterogeneous storage devices

Con: Tied to individual application servers and can be difficult to manage

3) Symmetrical or in-band virtualization

The I/O actually passes through a Windows or UNIX/LINUX-based “appliance” that sits between the servers and Fibre Channel or iSCSI storage devices.

Pro: This is a popular approach that provides a true heterogeneous solution. It requires no changes in either servers or storage and is easy to implement. Because it actually handles the data blocks, it supports sophisticated data protection services.

Con: Because the appliance handles all data, scalability and reliability can be issues, and failover must be coordinated with hosts and storage.

4) Asymmetrical or out-of-band virtualization

An appliance communicates with the SAN but is not in the data path. The abstraction map resides on the appliance and is distributed to application servers, which, in turn, do the remapping.

Pro: Because the appliance does not directly process data, this architecture is highly scalable.

Con: This solution requires placement of "agents" on application servers, which can add to the management burden. Also, because data does not pass through the appliance, data protection must be done by the host agent, which can send copies of data requiring protection through the appliance. This makes data protection an in-band process.

5) Intelligent SAN devices (Split Path Acceleration of Intelligent Devices or SPAID)

This is a hybrid approach that embeds virtualization logic in the SAN. Data is handled in the network while using the virtualization manager for some processes.

Pro: True heterogeneous support of application servers and storage with high levels of scalability

Con: Lack of a standard for porting to SAN devices, although a standard has been proposed.

6) Consolidated storage controller

This approach leverages the controller logic of an enterprise storage platform but supports externally attached heterogeneous storage devices as well. This architecture is similar to an in-band solution in that all data must flow through the storage controller.

Pro: Single point of management and high availability. Some also consolidate multiple protocols (FC, iSCSI, NAS).

Con: Still an in-band solution, which can limit throughput.

Emergence of file-based storage virtualization

The benefits of storage virtualization are not limited to block-based devices connected to a Fibre Channel or iSCSI storage network. File virtualization is an emerging and very promising technology for managing the growing volume of unstructured file data within organizations. Traditionally, large network attached storage (NAS) systems have been limited by their handling of such large amounts of file data. Because servers store data in a volume defined by the NAS storage, upgrades or changes to storage could cause changes in directory names, requiring corresponding changes enabling application servers to find the new location.

File virtualization eliminates this problem by creating a unified global virtual namespace, spanning file system containers residing on NAS systems and enterprise servers. Abstracting the physical location of files via the global virtual namespace enables storage to be expanded or migrated with no effect on the application servers. This is particularly important for unstructured files, allowing them to be moved, where and when necessary, without having to inform users or disrupting access. Network attached storage (NAS) systems can support file-based virtualization and, on the back end, can leverage SAN-attached storage from multiple vendors, provisioning storage to servers via NAS services.

Key services of file based virtualization include:

- **Seamless migration of data.** Data can be moved at file-level granularity with complete transparency to users and applications. File-level migration can be more specific than bulk-oriented block-level migration. The process can be automated if appropriate; automated migration eliminates disruption of business and the costs of migration-related client reconfiguration.

- **Global file manager/global namespace.** By abstracting the logical address from a file's physical location, this capability simplifies file access and data management and eliminates geographic limitations.

- **Storage tiering.** Files can be automatically moved among storage tiers, across multiple platforms and vendors with no disruption of operations. This policy-driven flexibility significantly reduces both capital and operational expenses and infrastructure costs.

Most organizations are planning and implementing some type of virtualization, if not multiple instances of server and storage virtualization. Each of these variants, even taken in isolation, has significant benefits. There is, however, great potential synergy between server virtualization and storage virtualization.

This is due to the consolidation, optimization and protection that storage virtualization can provide to the virtualized server environment.

At the same time, the virtualized server environment provides opportunities—and can even be a catalyst—for making storage utilization and management more efficient, and thereby driving increased consolidation, data protection and optimization in the storage environment

To take advantage of advanced functionality in a virtualized server environment, organizations must leverage a shared storage solution. In many cases, organizations are consolidating servers that have direct attached storage (DAS) with server virtualization. This provides an opportunity to consolidate the DAS islands onto a robust enterprise storage environment.

The benefits include a single management point, increased sharing among applications, and the ability to leverage the advanced functions commonly included with enterprise storage

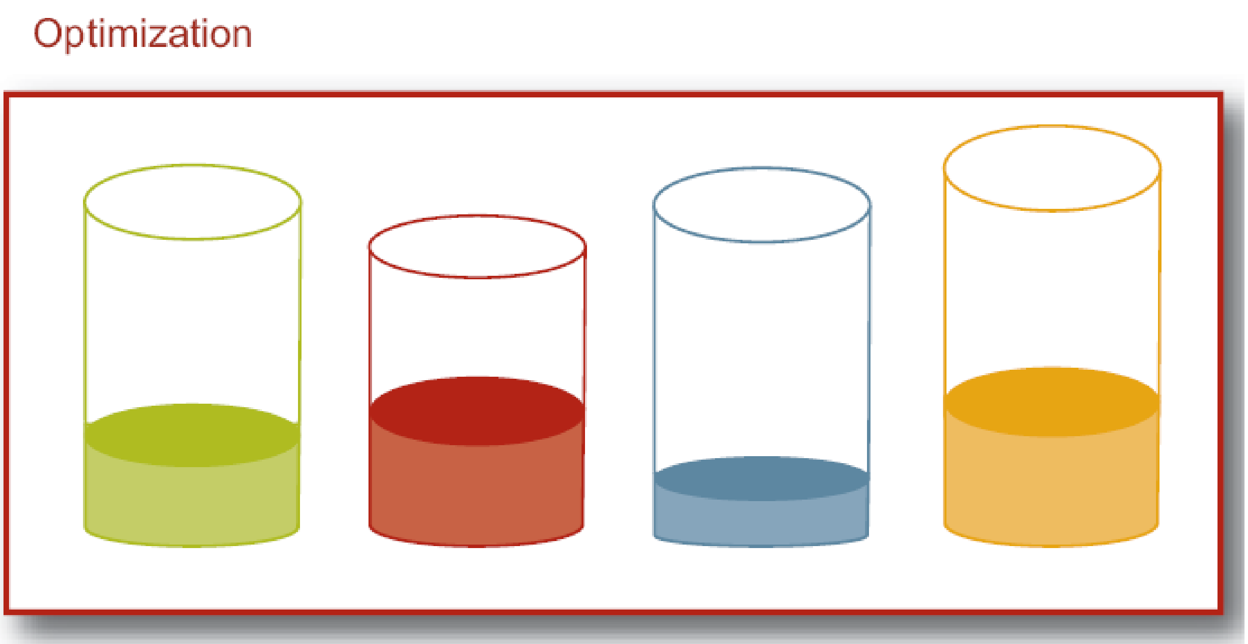

Virtualization supports optimization from both a server and storage standpoint.

4.4. Network Virtualization

Virtualization is changing the economics of the datacenter by making computing resources more flexible and efficient. To capitalize on these benefits, most datacenters must make fundamental changes to their Top-of-Rack or End-of-Row network architectures.

A virtualization strategy is more likely to succeed if the performance of the datacenter network can be assured before services are enabled and if the cable plant foundation for the Top of Rack or End of Row networks is certified to meet the needs of today and tomorrow.

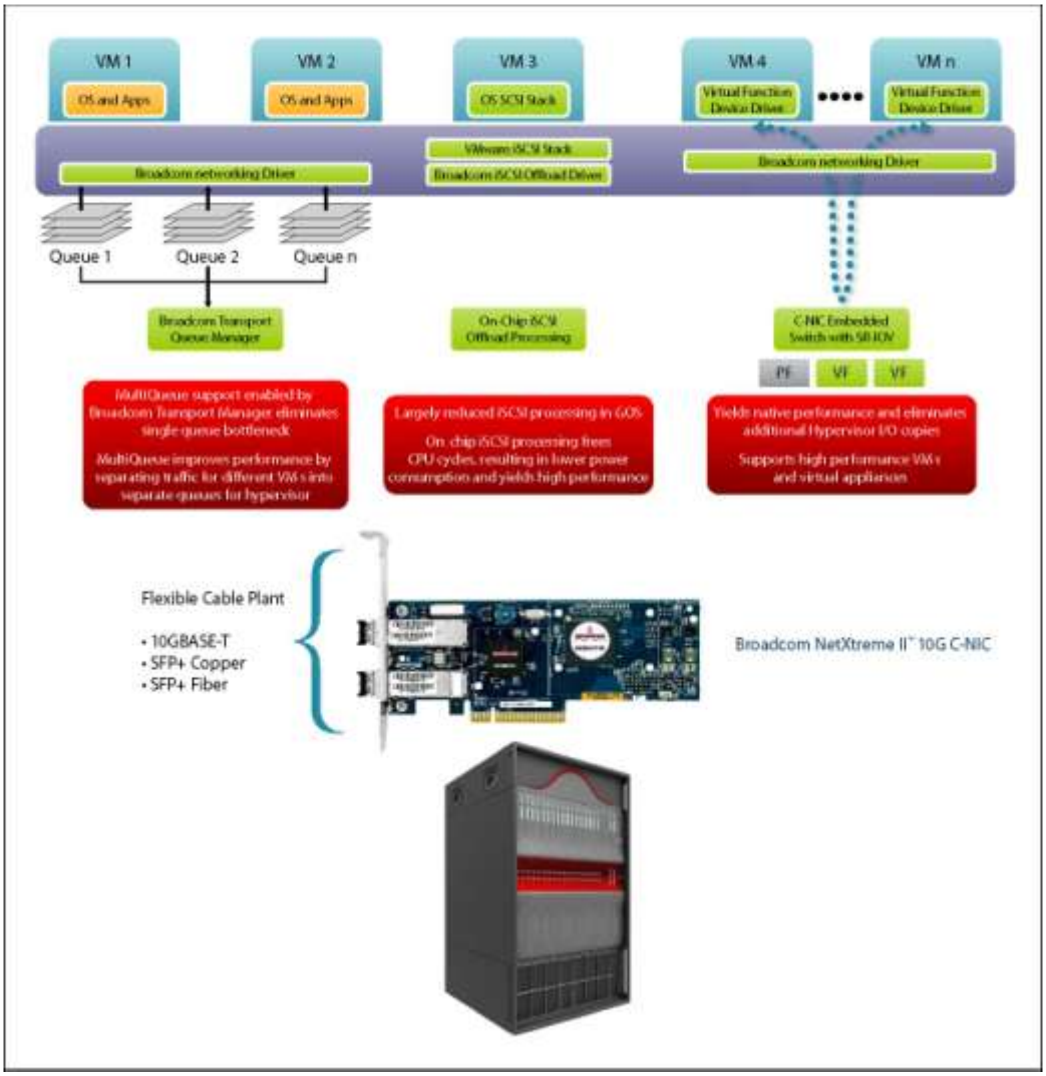

The physical network must adapt to the requirements and advantages created by virtualization, specifically higher utilization and higher bandwidth. To do this, forward-thinking network professionals have implemented the End-of-Row ("EoR") or Top-of-Rack ("ToR") topologies in their datacenter networks.

Option 1: The End of Row Topology

As a point of reference, Figure 1 depicts a pre-virtualized datacenter, where each asset (server, storage device, etc.) is individually linked to an Ethernet switch.

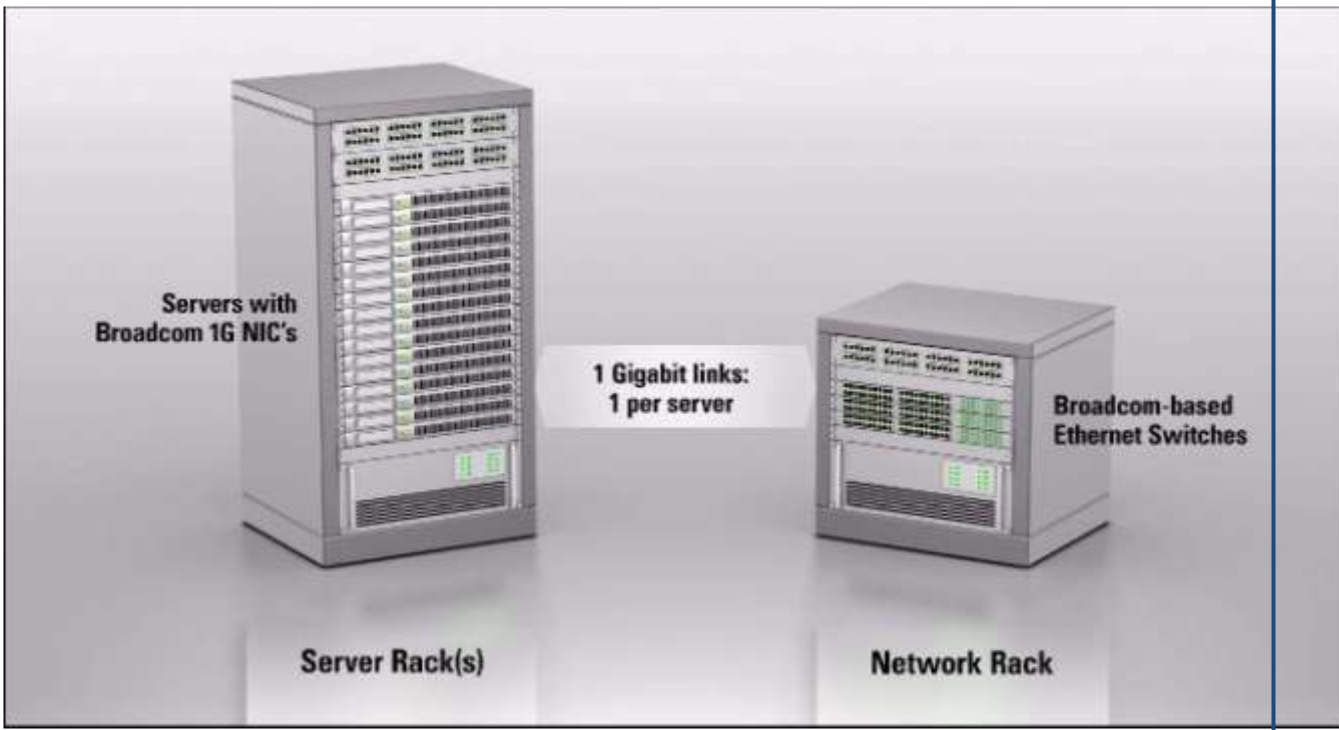

Figure 1: Conventional Topology

This topology uses structured cable connections that are difficult to modify. Since virtualization facilitates change, a network architecture that inhibits it is inherently problematic. The conventional topology is also dated, as a large number of 1 Gigabit links is incongruent with consolidated servers that need fewer, faster connections. This shift from many "thin roots" to fewer "thick roots" must be supported by the network.

The End-of-Row network topology, as shown in Figure 2, addresses the shortcomings of a conventional datacenter network by dedicating an Ethernet switch to each row of equipment racks. The virtualized assets in each rack, in each row, are linked to a switch in the EoR rack. That switch also provides a trunk connection to a datacenter concentrator.

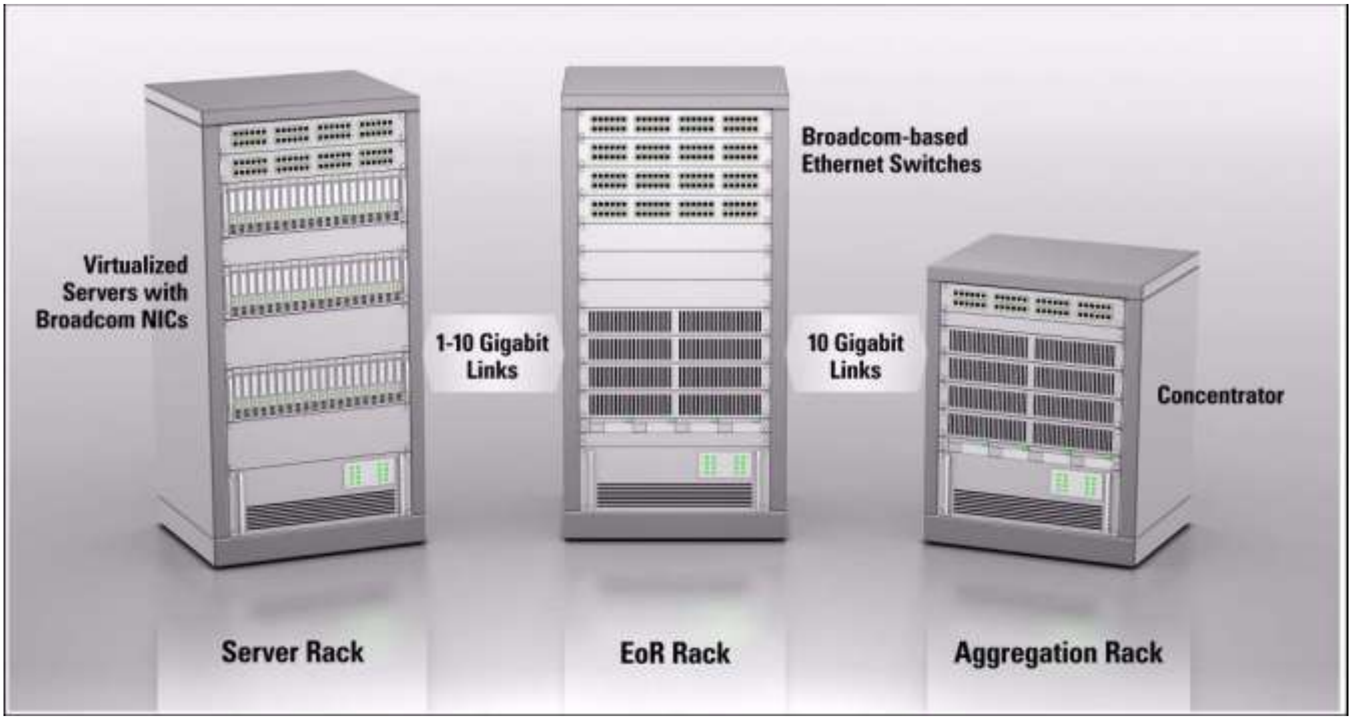

Figure 2: End-of-Row Topology

The EoR topology divides the switch fabric and physical connections from one tier into two, making the network more adaptable. EoR limits the length of the cables in the lower tier to the length of a row of racks. Shorter cables are generally easier to install and easier to change.

EoR topology confines the impact of asset reconfiguration to a row of racks, instead of across an entire datacenter. EoR may reuse some elements of the existing physical network, although major changes and upgrades are likely.

Option 2: The Top-of-Rack Topology

The Top-of-Rack topology is a bigger departure from the conventional architecture. It dedicates an Ethernet switch to every rack of servers. The ToR switch interconnects assets in each rack and provides a trunk connection to an aggregation point in the datacenter.

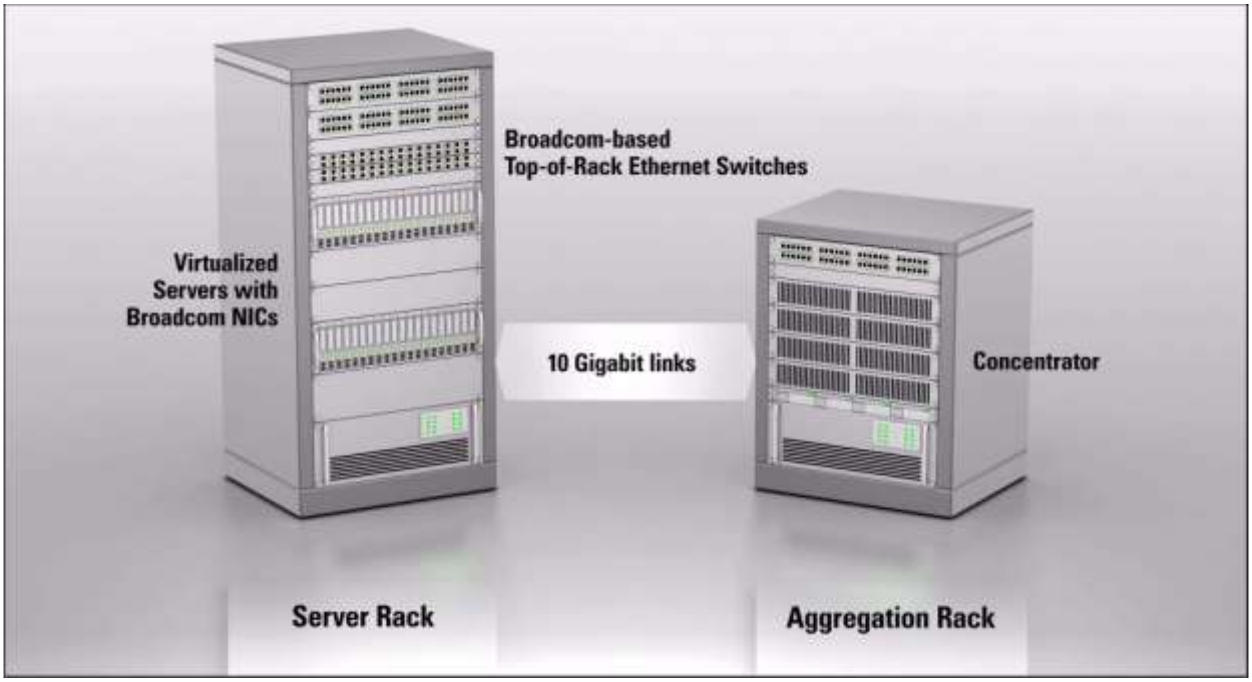

Figure 3: ToR Topology

Like EoR, the ToR topology divides the switch fabric and the physical connections into two tiers. The difference is the granularity of the lower tier. Where EoR creates modularity in a row of racks, ToR creates modularity in each individual rack.

Note that the ToR design does limit a server rack to only one switch. The graphic above shows two switches in a rack: one primary and one for redundancy. If the Ethernet switches are implemented as blades there could be even more switches in a rack. The Top of Rack versus End of Row discussion is complex. Many Layer 2 and Layer 3 networking issues must be considered. Some of the relative advantages of each topology are:

Top of Rack	End of Row
Less structured cable	Less disruptive to infrastructure
Easier to change/expand	Fewer switches, trunk connections
Most modular	Easier to manage/support

Topology and economics are the prime determinants of connectivity, but any option requires a reliable cable infrastructure. While cable is a passive technology and vendors provide warranties, you should bear in mind:

- Warranties end.
- A manufacturer's warranty probably excludes installation labor.
- The failure of any cable could mean the failure of services.
- Repairing a problem is always more expensive than preventing it.

To ensure successful delivery of virtualized services, the cable infrastructure should be subject to certification testing. Certification is a rigorous assessment, performed before the network is put into service, of connectors, installation workmanship, and cables. The result of certification tests are compared to industry standards, with the result being a “Pass” or “Fail” grade for every link. A link that passes certification meets the defined performance specifications. Links that do not pass are repaired, usually at the expense of the vendor or network installer.

Certification traditionally focuses on structured cabling, but if the Top-of-Rack topology is employed, it may include patch-cords used for Intra-Rack cabling.

As virtualization marches forward, it is triggering fundamental changes to datacenter networks that are unavoidable and in many cases, desirable. To deliver bandwidth reliability to virtualized assets and end users, 10 Gigabit Ethernet will be employed in the virtualized datacenter. 10 Gigabit Ethernet is important because it is a way to future-proof the datacenter network for years to come while meeting the immediate need to support virtualized servers and services.

10 Gigabit Ethernet can be implemented through a variety of copper and fiber options. The 10GBASET standard and new 10GBASE-T silicon open the doors for cost-efficient deployment of 10 Gigabit Ethernet across a virtualized datacenter. Whatever the choice in Layer 1 technology, transitioning to 10 Gigabit Ethernet requires forethought, careful planning and a methodology for test and troubleshooting.