

Virtualization

Virtualization Best Practices



Latest Revision: May 2, 2008

Contents

Chapter 1: Deciding on Virtualization..... 1

What do we mean by	1
Why use Virtualization.....	2
Areas Where virtualization Provides Quick Value.....	3
When to Avoid Virtualization.....	4
Common Advantages.....	4
Facilities Management	4
Security / Business Continuity	5
Resource Management	5
Management.....	6
Common Disadvantages.....	8
Security / Risks	8
Performance	9
Management.....	12

Chapter 2: Implementation Strategies 15

Taking Advantage of the Advantages.....	16
Deciding Where to Start	16
Deciding on the Appropriate Virtualization Engine	17
Identifying Management Tools and Requirements for that Selection.....	21
Identifying Change Control Methods.....	22
Identifying Data Storage Resources and Limitations	23
Defining Pertinent Maintenance Tasks.....	25
A Word about Virtual Appliances and Virtual Desktops	25
Determining Costs as well as Return on Investment	27

Chapter 3: Maintenance Considerations..... 31

Optimization Strategies	32
Allocating Dedicated Resources	32
Over-allocating / Sharing Resources.....	33
Re-allocating existing resources between VMs	34
Move an existing VM to a new host	34
Optimizing and Tuning the individual VMs	35

Appendix A: References and Links.....37

Appendix B: Virtualization Checklists47

Identify Virtualization Candidates 47

- List and Prioritize Servers/Applications/Solutions 47
- Generic Questions for Guest Systems 48
- Configuration of Guest System 50
- Operating System with Tightened Security / Limited Functionality 53

Guidelines for Host Systems 54

- General Questions for Host Systems..... 54
- Basic Hardware Configuration 55
- Disk Configuration 57
- Management Infrastructure 59
- Cost Analysis 61

Glossary63

Chapter 1: Deciding on Virtualization

There is a lot of potential value that can be derived from virtualization but the key to maximizing that value is to understand what the expectations and business drivers are. Each client planning virtualization adopts one or more goals based upon product claims, press coverage, and industry peer input. One common virtualization goal is to improve the way in which IT manages its resources. This improvement may take the form of increased peak capacity, improved resilience, reduced configuration costs, or reduced systems management errors. One or more of these subordinate goals is often attainable. Perceived virtualization savings may also take the form of reducing the amount of hardware that IT purchases or manages. Although there may be some debate as to where to start and which goal saves the most, there is little doubt that virtualization has value and is here to stay.

This Best Practice document describes how to take advantage of virtualization by identifying a number of scenarios where virtualization can provide value, as well as best practices related to specific goals.

Note: This document is a work in progress. Updates will be provided as they become available. If you have any comments or feedback on this document please forward them to ImpCDFeedback@ca.com.

What do we mean by “virtualization”?

Virtualization is a broad term that can mean different things depending on who is using/interpreting it. Wikipedia (<http://en.wikipedia.org/wiki/Virtualization>) provides the following definition:

" a technique for hiding the physical characteristics of computing resources from the way in which other systems, applications, or end users interact with those resources. This includes making a single physical resource (such as a server, an operating system, an application, or storage device) appear to function as multiple logical resources; or it can include making multiple physical resources (such as storage devices or servers) appear as a single logical resource."

Or, to put it more succinctly, "hiding of hardware detail and reducing system management effort, through encapsulation."

There are two primary approaches to virtualization:

- **Platform Virtualization**

This initial form of virtualization refers to a single server which hosts one or more “virtual guest machines.” This is also often referred to as “Server Virtualization.”

- **Resource Virtualization**

Later virtualization efforts expanded this definition to include virtualization of specific system resources, such as storage and network resources. This can be done within a host server or across multiple servers (using a SAN, for example). Modern blade enclosures/servers often combine platform and resource virtualization, sharing storage, network, and other infrastructure across physical servers.

This document focuses primarily on Platform Virtualization with some coverage of Resource Virtualization; however, unless otherwise specified, the topic should be assumed to refer to Platform Virtualization.

Why use Virtualization

There are both advantages and disadvantages to using virtualization in your environment. It is critical that you understand what virtualization can offer in conjunction with your required level of skills/commitment. It is up to you to reconcile these factors with your expectations of how virtualization can be used in your environment. Finally, it is critical to understand that:

applications aren't suddenly going to require less resources just because they are virtualized

On the contrary, virtualization **adds** overhead and a virtualized application will use more resources than before. A virtualized application will not run faster unless it is hosted on faster hardware than it was run on originally. Thus, attempting to virtualize using your existing hardware is typically a “bad” idea.

The actual amount of additional overhead depends on a number of factors including the type of application being virtualized, the type of virtualization engine being used, what kind of hardware is available, and how it will be configured/used. According to lab tests in the April 2007 issue of “Network Computing” the overhead for an ESX server is typically less than 10% but ranges from 6% to 20 % (your results may vary and each user should monitor their own overhead). To view this article in its entirety, go to the following link <http://www.networkcomputing.com/showArticle.jhtml?articleID=198700359>.

It is important to make sure you have enough storage space, memory, CPU, network bandwidth and other resources to handle the applications plus the virtualization overhead. If the applications are business critical, you should plan for worst case scenarios; however avoid dedicating more resources than necessary since this will negatively impact other virtual machines on this host.

Areas Where virtualization Provides Quick Value

Following is an overview of scenarios where virtualization can provide quick value to your environment.

1. Organization needs a library of servers with different configurations, such as:

- Software Development (Test scenarios)
- Quality Assurance
- Software support where it's important to be able to quickly and easily reproduce a relatively large number of environments.
- Demo centers / demo scenarios

This is a common starting point for many companies since it is easy to realize significant value and the risks are typically minimal. Value is seen in reduced time to provision servers as well as in reduction of errors.

2. Consolidating selected business applications deployed to:

- Lightly used servers. This typically includes:
 - Service Providers (xSP) that have multiple small clients.
 - Multiple mid-tier managers originally implemented on separate servers for political, organizational or legal reasons.

In many cases isolation provided by virtualization is sufficient, especially if the data is separated onto private disk systems; however, it is critical to verify that virtualization satisfies the organization's isolation/separation requirements. Value is seen in reduced hardware and management costs.

- Servers with predictable resource consumption profiles. This will allow you to plan the distribution of work for virtualized servers. In these cases, keep in mind that:
 - Special care is required for applications that require lots of I/O.
 - Applications that require different sets of resources at the same time can coexist on the same physical server.
 - Applications that require the same resources at different times can also coexist on the same physical server.

In each of these cases value is realized through a reduction in the number of servers resulting in both hardware maintenance and management cost savings. Additional details on attaining rapid ROI through virtualization are provided later in this document.

Unless there is some other significant justification for virtualization in your environment, if the project does not fall into one of these categories you may find that it is hard to save money on the project. Although there may be other good reasons to consider virtualization it is critical that you understand what you are trying to accomplish with this project.

When to Avoid Virtualization

Regardless of whether it is *possible* to virtualize servers and applications in your environment, there are certain situations in which the potential risks far outweigh any advantages that might be gained.

Applications that make frequent and unpredictable demands on a large parts of the system's available resources are **not** ideal candidates for virtualization. A couple of examples are:

- Large database servers. Virtualization of database servers is rarely beneficial. Database server utilization is better improved by employing multiple database instances.
- Application Virtualization type servers, such as Citrix, and other types of servers that already include their own techniques for virtualization.

Additional examples of more common disadvantages are listed later in this guide. In these situations, projects should be analyzed, on a case by case basis, to carefully calculate the risks connected to the application / system.

Common Advantages

Following are some of the more common advantages to virtualization. The relative importance of each will depend on your exact environment and requirements.

Facilities Management

- **Saving Datacenter Space**

A common problem with datacenters is that the ever increasing number of managed applications requires more and more servers which, in turn, require more and more floor space. Virtualizing a significant number of these applications may save you from having to move into a new larger datacenter, and, in fact might enable you to use a smaller datacenter or to allocate some of the existing space for other functions.

- **Hardware Cost Savings**

Server virtualization typically requires larger and more expensive servers, however, *when done correctly*, combining multiple, under-utilized servers into a single larger system can result in a significant cost savings (e.g., lower hardware purchase cost and lower hardware maintenance cost). Fewer servers require less supporting infrastructure in the form of floor space, air conditioning, racks, networks, wires, cables, power supplies and backup systems.

- **Reduced Energy Bills**

Reducing the number of servers in the datacenter reduces the electricity bill. This cost savings can be significant when you account for the servers, the monitors and the air conditioning required to keep them cool.

InformationWeek (December 18/25 2006) cites one example where an organization reduced costs by half - from \$7000 to \$3500 per month.

Security / Business Continuity

- **Easy to backup the complete image**

Copying a complete virtual environment (image) to a backup location or to a staging area to allow controlled upgrades of the application/operating system is a trivial task when virtualization is used.

- **Disaster Recovery / Business Continuity**

As previously noted, virtualized copies of the environment can easily be moved to off-site servers. Keep in mind, however, that many applications depend on fixed IP addresses and/or the availability of other resources (for example SANs). Therefore, it is critical that you identify these requirements beforehand, by conducting tests in an environment in which the complete original infrastructure is down (or simulated to be down). During these tests it is also important to verify that any other applications that rely on any of these services in the virtualized environment can find the new clones of those services.

- **Virtual appliance might enhance security**

Since you have complete control over the required resources in a virtual appliance you can enhance security by removing any components that aren't required for this specific application. When this is done correctly, it can greatly enhance security; however it might also complicate the process of patching the system.

- **Virtual Desktops provides enhanced control over security**

Virtual Desktop Infrastructure (such as VMware ACE 2) can allow simplified and enhanced control over security by limiting a user's access to specific resources and certain types of data. For example, by providing a trusted partner with a secured virtual desktop instance to access sensitive information you minimize the risk that data will leave the central server.

Resource Management

- **Simplifies Chargeback systems**

Decoupling services from physical servers simplifies chargeback systems by enabling you to delineate utility pricing based on a pay per use model. The utilization metrics required for the chargeback system are often the same as those required to manage load balancing between systems.

- **Optimized usage of existing hardware resources**

In most datacenters there are a large number of servers that rarely take full advantage of the available resources. By managing your virtualized environment wisely, you can enable multiple logical servers to share resources in a way that allows access to more resources when needed but to share them with other applications when they are idle or close to idle.

When planning your deployment, however, you should allow for a worst case scenario and identify which applications might need resources at the same time.

- **Faster deployment of new logical servers**

With the necessary hardware resources available, virtualization technology can significantly simplify the task of deploying certain types of servers. For example, you can deploy an additional web server fairly quickly and add it to the load balancer rotation as additional resources are required. This often reduces the provisioning time for a new server from days (or maybe even weeks) to hours.

- **Moving logical servers between hardware**

Virtualization enables you to manage server load more efficiently by allowing you to move complete virtual servers to new hardware whenever additional resources are needed.

This is especially easy to do if you have the tools and infrastructure for “hot migration” - which allows you to move logical servers while they are still running (for example VMOTION or Live Migration together with a SAN or iSCSI infrastructure).

Note: When planning for Hot Migration it is critical to understand that this typically isn't supported between machines with different CPU architectures. Current technology requires the CPU to have the same vendor, processor family and core stepping. In addition, the source and target servers need to have access to the same external resources, such as SAN and network.

- **More flexible infrastructure**

Since logical servers can move easily between hardware sources, a virtualized environment is one in which the hardware is completely decoupled from the operating systems and the software. The result is a very flexible infrastructure where the hardware can be used to support the services/applications that are most important at this moment. This abstraction allows you to reduce costs since hardware and software upgrades no longer are directly coupled to each other.

Management

- **More, smaller applications logically separated from each other**

Logically separating applications from each other through the use of virtual appliances can simplify support by reducing the likelihood of applications “colliding” with one another.

Without virtualization you might be tempted to run multiple smaller applications within the same OS but this quickly leads to a large number of application combinations to certify and support.

■ **Fewer servers**

Since virtualized environments typically have fewer physical servers they can be easier to manage, especially from a security point of view. However, in order to reap the full benefits of virtualization, it is critical to carefully manage and monitor the performance and health of the individual virtual machines, the host systems and the connected SAN systems.

You will also need to account for the many additional management issues related to virtualization (see "Management" on page 12).

■ **Hardware maintenance**

The ability to move logical servers between hardware can simplify hardware upgrades by enabling you to build a new server, verify its functionality and compatibility with a copy/clone of the live image - all without affecting the existing application. When testing is done, you move the live applications over to the new server.

In a similar way you can easily and quickly have another physical server take over the role of hosting the applications when the original server has a hardware problem.

■ **Software maintenance**

With the right planning, change control for software maintenance can also be significantly enhanced through judicious use of virtualization. Since the complete logical machine can be copied and handled as a set of files you can easily set up separate areas for:

- Development
- Test / Quality Assurance (QA)
- Available Images / Gold Images
- Archive
- Configuration
- Production

Using a structure like this you can easily upgrade and test a new version in the development and QA area while still running the old version in production. When the new version is approved and you have a copy in the gold master library you can schedule a small maintenance window and transfer over to the new, already updated and verified, image. This topic is discussed in greater detail in the "Identifying Change Control Methods" section on page 22.

Common Disadvantages

Following is a summary of the common disadvantages associated with virtualization. These should be carefully weighed against the advantages to determine if virtualization is right for your environment.

Security / Risks

- **Internal Resistance**

A problem often encountered during a server consolidation project is that some parts of the organization might resist giving up control over their existing hardware or applications. It is very important to address this early to ensure cooperation from all application owners.

- **New, relatively unproven technology, tools and processes**

Though “virtualization” is not a new concept in software, the changes introduced in its most recent form, as well as the impact and implications of those changes, need to be clearly understood (and, thoroughly tested). This includes the introduction of new:

- **Abstraction layers**

Virtual Engines/hosts introduce a new abstraction layer that can potentially introduce new failures as well as security exposures. This is particularly true for engines which employ hypervisors since, by nature, hypervisors should be as lightweight and efficient as possible and therefore have limited error recovery and security implemented. This may, however, be mitigated by configuring the hypervisor or using a specific security-related virtual appliance or a plug-in to the hypervisor that manages the system.

- **Security Architecture**

Since applications have traditionally been tied to a specific piece of physical hardware and infrastructure it has been possible to design security around an environment in a fairly static way, making sure that the physical server and environment are secured. When dynamically allocated virtual servers are used it is necessary to track where the application is currently residing and this may require the use of dynamically configured security or, at minimum, the rule that the application only reside on secured host servers.

- **Immature and/or incomplete tools**

Dynamic environments also require new or enhanced tools for managing and securing virtualization. This includes the ability to:

- ◆ Ensure the complete environment is successfully patched. This applies to both guest and host systems.
- ◆ Analyze and manage the host OS and the virtual network to find and address bad configurations and other vulnerabilities.
- ◆ Analyze and secure traffic between VM’s on the same machine

- ◆ Be aware of and secure environments with Virtual Desktop Infrastructure (VDI) and Mobile VMs

- **Reorganization / consolidation of servers and applications**

Any major change to a server environment – reorganization, consolidation, modification – is disruptive and, by its very nature, a risk. However, if the change is managed in a controlled and well planned way, this risk can be managed as well and kept to a minimum.

- **Loss of logical servers**

Unless the images are well managed (see “Identifying Change Control Methods” section on page 22) it is easy to mistakenly delete a complete logical server (or set of logical servers) in a virtualized environment.

- **Consolidated datacenters**

Consolidating many servers into one big datacenter can provide a great advantage; however it also creates a huge single point of failure. Unless this datacenter is extremely secure, events such as a natural disaster (fire, flooding, etc.), power failure or sabotage, can cause a major disruption to all (or at least a significant part of) your IT infrastructure.

In addition when multiple smaller datacenters are combined into a single datacenter, it is critical to make sure that the supporting infrastructure can handle the additional load for normal situations as well as disaster recovery scenarios. This includes network infrastructure, power requirements, storage and backup requirements, cooling and floor space.

- **Real time or near real time requirements**

Care needs to be taken with applications that require real time or near real time response since the system clock on some virtualized systems can temporarily lag as much as 5-10 seconds if virtual machines are under heavy load. This is typically not a big problem, but it might be an issue in a system that requires real time or near real time response.

For more information see the VMware whitepaper “Timekeeping in VMware Virtual machines” (<http://www.vmware.com/resources/techresources/238>) and the VMware KB1420 concerning “Linux Guest Timing” (<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1420>). These papers focus on VMware hosts, however, similar problems exist for other solutions as well.

Performance

- **Overhead on resource consumption**

All virtualization technologies impose a performance penalty. If there are multiple virtualized machines (VMs) running at the same time, the amount of resource consumption - including the overhead – increases cumulatively and, often non-linearly, due to contention for resources that are simultaneously being used by multiple VMs.

The level of overhead depends on a large number of factors, however the typical increase for a hypervisor based system is around 10%. In the December 18/25, 2006 issue, *InformationWeek* cites this value as up to 15%, and a test in the April 2, 2007 issue of *Network Computing* estimates a 10% overhead (low of 6% and highs of 20%). Engines without hypervisor technology or other hardware assists have significantly higher overhead. Higher end blade solutions or high end hardware-based solutions (including LPAR or Containers) have relatively low overhead.

The penalty might be especially large in situations where:

- Multiple applications require heavy disc-access.

When multiple applications on the same host try to use any limited resource (typically, I/O), this will significantly affect performance. In disk I/O situations, this can often be handled by off-loading each application's data storage to a dedicated storage subsystem or using a high end dedicated storage solution (including SAN or iSCSI systems).

- Applications are performing many small transactions

Performing many small transactions (or combinations of large and small transactions) such as disk I/O, or network traffic, creates a significantly higher overhead than a few large transactions.

■ **Resource Allocation**

One advantage to virtualization is better utilization of existing resources – since, when correctly planned and managed, the virtualization engine can handle a degree of over-utilization thereby allowing applications to use resources that currently aren't in use by other applications. However, when mismanaged, this is also one of its biggest *dis*advantages since, under heavy load, the response time from virtual machines becomes unpredictable.

■ **Bottlenecks/Queuing Delays**

Poorly managed resource allocation can lead to bottlenecks. Although bottlenecks are not a new problem, they escalate more quickly in a virtualized environment since you have multiple logical servers potentially hitting the same resources.

There are few rules regarding what resources will be strained by which application. Bottlenecks depend, to a large degree, on what resources are available and how the application is used in a specific environment. However, typical examples include:

- **CPU constraints:** A common way to handle application CPU constraints in a non-virtualized environment is to allocate additional CPUs to the application. In a virtualized environment, however, this is not necessarily a good idea since the virtualization host typically waits until it has *all* allocated CPUs available before it can assign *any* resources to the virtual machine. In other words, in a virtualized environment it may be a better idea to assign fewer fast CPUs to each virtual machine than to over allocate CPUs. Of course, if you virtualize an application that is designed for multiple CPUs you should specify

that number of CPUs for the virtual machine and ensure you have that many physical CPUs available to the virtualization host. If you run an application and its dedicated database on one virtual system you often need to define two CPUs to permit both to be active. If you simultaneously run several virtual guests, each defined with multiple CPUs, then you should have, at least, the sum of defined CPUs free on the virtualization host to avoid a serious performance impact.

- **Disk I/O constraints:** DBAs know that databases often perform a large number of small read/write transactions to disk. This can severely affect performance unless you separate various functions to different physical disks/stripe sets. For this reason, basic storage tuning is very important in a virtualized environment. In fact, because disk intensive applications often compete for access to shared disks, it is common to dedicate disk space on separately allocated disks/stripe sets. SANs can provide disk resource virtualization and are commonly employed for virtualized applications. Properly sizing SAN cache and managing disk allocation effectively is critical.
- **Memory constraints.** Web servers and database servers often take advantage of available memory for caching. This can significantly improve performance by reducing disk I/Os.

Virtual machines, however, often have a problem sharing memory in an efficient and fast way and it is critical, therefore, that you provide enough memory for each application. The challenge is to do this without extensively over allocating resources, and making the solution more expensive than necessary.

This can be handled by using a management tool that monitors the application's memory requirement and using this knowledge to decide which applications fit together on a specific host.

Note: A limited over allocation of memory is often desirable since modern enterprise level virtualization hosts have advanced memory management functions that allow applications to share memory between different virtual machines and can dynamically borrow memory between the virtual machines. For additional details, consult the "Memory Resource Management in VMware ESX Server" (http://www.vmware.com/pdf/usenix_resource_mgmt.pdf) document.

- **Network constraints:** Client-Server type applications, along with network management tools, often place a large number of small packets on the network. Even if these packets don't add significant overall load, they may severely affect performance of other virtual machines sharing a NIC if the round trip latency is poor. Dedicating NICs to production virtual machines is typically required to avoid wide swings in response time of network dependant applications.

Most of these constraints can be managed by adding more resources and/or reorganizing which virtual servers are running on which hosts.

It is also highly recommended that you carefully read the performance tuning and best practices documents and white papers that are available from host system vendors. For example, the “VMware tuning Best Practices for ESX Server 3” document, which is available from VMware (http://www.wmware.com/pdf/vi_performance_tuning.pdf) focuses primarily on VMware’s ESX server, however many of its performance tips are applicable to other solutions as well.

Note: Poorly written applications often consume multiple resources. Therefore, whenever possible, review in-house applications to determine if they can be enhanced to better use available resources. For example, a database engine given poorly written SQL will likely put unnecessary stress on CPU, Network, Disk I/O and Memory. CA CEM and Introscope can rapidly identify resource bottlenecks in web applications.

Management

■ Support

Although the ability to isolate applications through virtualized appliances can simplify support, the disadvantage is that many software vendors do not fully support virtualized environments or else require that the issue be reproduced in a non-virtualized environment.

■ Multiple Servers on one physical server

Having multiple virtual machines on a single physical machine can introduce the potential for:

- Security vulnerabilities (see the previous “Security / Risks” section).
- Increased failure impact. A single HW failure now affects multiple logical/virtual servers. This risk can be reduced through use of resilient server hardware and good disaster recovery routines. In fact, some clients choose resilient server hardware for virtualization to reduce risk of failure when comparing multiple failure prone small servers with a single large resilient server.
- Any reboot/maintenance to the host system interrupts all guest OS, however this can be managed using “hot migration” technologies (such as VMOTION or Live Migration).

- Bottlenecks (see details in the "Performance" section above):
 - ◆ CPU cycles, Memory capacity and Network are typically relatively easy to handle by analyzing the typical load of the various applications and proactively allocating the resources. The resources can be dynamically allocated but this should be avoided, whenever possible, since it typically slows down the applications to repeatedly reallocate resources.
 - ◆ I/O bottlenecks are a tougher and a growing problem with multiple virtual machines (especially if they are often heavily used at the same time). However, understanding the usage pattern will allow you proactively allocate dedicated physical resources or additional shares of existing resources to the virtual machines.

- **Specialized skill set required**

Even though management of a virtualized environment can often be streamlined and simplified it still requires new, and currently unusual, skills.

Personnel that don't fully understand the specific requirements for virtualized environments can cause severe damage to the infrastructure. Therefore it's important to allocate time/resources for training both before and throughout any virtualization project.

Chapter 2: Implementation Strategies

Planning is an important part of any project - even more so with virtualization. The key to a good plan is to understand, from the outset, what the goals are, in order to ensure that those expectations can, in fact, be met by the implementation. Before undertaking a virtualization project you should clearly identify the following:

- **Reason** for the project (i.e., what are the business drivers?)
- **What** you are trying to virtualize (e.g., specific functions or applications?)
- **How much** the project is expected to cost (and to save)
- **What risks** – both functional and financial – are expected and, more importantly, acceptable.
- **Scope** of the implementation (i.e., is it a single, focused project, or will there be multiple phases and milestones)
- **Other changes** that are anticipated in the environment and how they might impact or be impacted by virtualization

This is, by no means, an exhaustive list, nor does it apply to all projects, however if you don't have a good understanding of the answers to those questions *before* you start the project you are more likely to encounter some unexpected bumps during the deployment.

Whatever the answers to these questions may be, it is important to carefully assess the tasks at hand and to plan out the action steps accordingly. Keep the project plan open particularly when dealing with a larger scale project. This enables you to incorporate any lessons learned during the first steps/milestones when it is time to start subsequent steps and sub projects.

For example: If the first major milestone in one virtualization project is to implement a library of different virtualized test environments for the Quality Assurance (QA) team, it would then be beneficial to analyze the results from that experience before proceeding with the next milestone, which may be to incorporate various lightly used business critical applications into the environment.

Taking Advantage of the Advantages

As you can see there are many advantages to virtualization – but you need to understand the realities of those advantages, as well as how to counter the potential, and often related, disadvantages. The following sections help you take advantage of those advantages by:

- Deciding where to start
- Deciding on the appropriate virtualization engine
- Identifying Management Tools and Requirements for that selection
- Identifying change control methods
- Identifying data storage resources and limitations
- Defining pertinent maintenance tasks
- Determining costs as well as return on investment

Also included are a few helpful thoughts on virtual appliances and virtual desktops, detailed discussions of which are currently beyond the scope of this document.

Deciding Where to Start

There is no simple answer to this question since every situation is unique and every project requires an analysis of the organization's specific situation to realize the pertinent facts for that situation. Consult the previous "Areas Where virtualization Provides Quick Value" section on page 3 for tips on where virtualization quickly adds value.

Often the best place to start is to look at those situations where a library of different environments is required to be accessible on demand – for example, development, quality assurance, support and demo centers. Keep in mind the list of "Common Advantages" and "Common Disadvantages" identified earlier in this document (see page 4) and consider how these apply to your organization.

The next step is to carefully analyze the list of current (and potential) business applications to determine which would make a good fit for consolidation - for example, File and Print Servers, Web Servers and some carefully selected Business Applications. Whenever possible, it is best to start with servers that have low activity levels.

Note: Careful planning is especially critical for I/O intensive applications. Although these are generally not ideal candidates for virtualization, if you must virtualize them you should plan to use one or more dedicated physical disks (preferably implemented through SAN or iSCSI).

Deciding on the Appropriate Virtualization Engine

The following list represents some of the more interesting players in this area – either from a market share and/or a technology point of view. It is not intended to be a complete list.

- VMware (www.vmware.com)
- Microsoft (www.microsoft.com)
- Citrix Systems (<http://www.citrix.com/>)
- XenSource (www.xensource.com)
Acquired by Citrix in August 2007.
- Virtual Iron (www.virtualiron.com)
- Parallels (www.parallels.com)
Earlier sometimes referred to as SWsoft
- Sun Microsystems (<http://www.sun.com/>)
- InnoTek (www.virtualbox.org)
Acquired by Sun Microsystems in February 2008.
- Amazon EC2 (aws.amazon.com/ec2)

However, from a market share and market recognition point of view there are only a few players that currently have any direct weight.

In a survey conducted in February 2007, Forester Research asked 350 North American and European Technology managers a few questions related to virtualization. They found that:

- 40% of North American companies were virtualizing servers in 2006 (up from 29% 2005).
- When asked about their Vendor of choice:
 - 53% cited VMware
 - 9% cited Microsoft
 - 28% cited HP, IBM or DELLNone of these offer an exclusive virtualization software tool. In other words, it is likely that the actual tool isn't that important - as long as the hardware and solution provider supports the end result.
- 6% didn't know
- 4% mentioned others (one individual cited Xen).

Looking at this report it is obvious that VMware and Microsoft are the big players when it comes to the software virtualization tools, however it is also worthwhile to include XenSource (Note: August 15th 2007 Citrix Systems Inc announced a definitive agreement to acquire XenSource) in the discussion. XenSource, along with Virtual Iron, have been getting a great deal of attention in the press.

VMware is clearly the market leader and the front-runner at this point, however, in some situations, MS Virtual Server (and Hyper-V, when it is released) could potentially be more efficient than VMware and Xen solutions for Windows guests OS – due to its detailed knowledge of the Windows OS.

Microsoft currently has a major new virtualization engine called Hyper-V in beta (this was formerly referred to as “Viridian” or WSV). Microsoft made the feature complete release candidate available on March 19, 2008 and the final version is targeted for release in August 2008. Following are some important facts about Hyper-V:

- Built using true hypervisor technology
- Hyper-V host requires the x64 edition of Windows Server 2008,. If used on a Server Core installation this is a relatively thin layer.
- Host requires CPUs with hardware assisted virtualization (AMD-V, formerly called “Pacifica”, or Intel VT).
- Support for 32 or 64 bit guest OS across different server platforms (Windows and Xen-enabled Linux distributions). Microsoft will release a complete list of supported platforms before the product becomes generally available.
- Ability to support up to 4 virtual CPUs for each virtual machine (guest OS).
- Quick Migration enables rapid migration from one host to another with “minimal” downtime. The length of the downtime depends on a number of factors such as I/O performance to the SAN and the size of the virtual machines.
- Ability to take “snapshots” of running machines so that the user can easily revert to previously saved snapshots.
- Managed through Windows Management Infrastructure (WMI) and/or a published HyperCall API.

This simplifies integration with 3rd party management tools.

Following are some interesting planned or potential future enhancements for Windows Server 2008 and Hyper-V:

- Support for up to 64 CPUs on the host servers
- Live Migration, which enable migration of active machines to new hosts without any downtime.
- Hot add of resources such as storage, networking, memory and CPU.
- Ability to share memory between guest machines.

- A standalone Hyper-V Server that can be installed directly on the hardware, without a separate OS layer.

It is also noteworthy that, in July 2006, Microsoft announced that they are cooperating with XenSource to allow Xen-enabled guest operating systems, including Linux, to run on Windows Server "Longhorn" and that this will be supported by Microsoft.

Although a majority of the references and examples in this document refer to VMware that should not be interpreted as an endorsement of VMware over any other vendor. It is merely a reflection of the author's experiences. In determining the best tool for your project and environment you need to carefully consider what you are trying to accomplish, what type of functionality is required and when. For any larger scenario you need to consider the following items as well:

- Is there a particular solution for which you or other members of your team already have existing knowledge and experience?
- Does the solution support all functions and protocols that are required for your architecture?
- Do you anticipate a need for a specific format of Virtual Disks (sharing information/disks with other projects)
- Will the deployment be resource intensive? If so, a true hypervisor based architecture is highly recommended.
- What solutions have been tested or can be tested with the existing hardware platform and most critical applications.
- What is the expected total cost of ownership? For more information on this topic see "Determining Costs as well as Return on Investment" on page 27.

Required functions and protocols

It is important to clearly identify what functions you require from your virtualization engine because there is a big difference in the types of functions that various vendors ship with their different offerings.

Some of the more commonly required functions include:

- Support for Hot Migration, which allows you to easily, and without interruption, move a virtualized environment to another physical server.
- Support for SAN solution (Fibre Channel and/or iSCSI based). It is critical that the virtual images can reside on a SAN. If you are planning to use diskless systems you also need to make sure that the virtual engine can boot from the SAN.
- Support for multiple storage repositories which allows you to minimize the risk of storage contention.
- Support for all operating systems that are required by your organization (16/32/64 bits versions of Windows, UNIX and/or Linux).

- Ability for the system to access, utilize and distribute all required resources (including sufficient amount of RAM, CPU types, CPU count etc)
- Management tool or support for a management tool that can monitor performance and availability. These tools should, preferably, also integrate with your existing Enterprise Management System.

Virtual Disk Formats

Virtual Disk formats allow you to easily move a virtual disk between physical disks, however its important to understand that there are different standards available and that Microsoft and VMware are both using their own 'standard'.

Both Microsoft (VHD) and VMware (VMDK) have made their specifications for Virtual Disks open and free to use.

- VMware made the specification for its VMDK available in April 2006, however, they reserve the right to revise or rewrite it.
- Microsoft made its "Virtual Hard Disk Image Format (VHD)" freely available to the public in October 2006. (However, like VMware, they kept the right to revise or rewrite it.)

It is also noteworthy that XenSource uses Microsoft's VHD format.

Note: Since these formats now are open tools have been developed to convert between the formats. This typically works well for data disks however you need to be careful with system disks since different VM Engines emulate the system in different ways.

Hypervisor based Architecture

A hypervisor-based architecture employs technology that communicates more directly with the hardware without exchanging calls with an intermediary operating system. This means that the virtualization software requires fewer hardware resources (e.g., memory and storage) and also has a smaller overall footprint.

- Architectures with true Hypervisors makes the communication with multiple virtual machines significantly more efficient.
- Hypervisor-based solutions to a larger degree tend to take advantage of specific hardware functions supporting virtualization (available for both Intel and AMD chips).

All major players have or are planning to have a version based on hypervisor technology.

- VMware ESX Server is based on hypervisor architecture
- The open source project Xen (including XenSource Enterprise 3.x) has been a hypervisor based solution from its start.

- “Microsoft Virtual Server 2005 R2” runs as a service on top of the guest OS (WinXP, Win2003) and is not a true hypervisor based solution.
- Microsoft Hyper-V includes a hypervisor based Virtualization Engine and is planned to be released in August of 2008.

Windows Server 2008 x64 was shipped with a beta version (RC0) of Hyper-V and an update to a feature complete release candidate can be downloaded from <http://www.microsoft.com>. The product is expected to be generally available by August 2008.

Hardware Virtualization Engines and Applications

Virtualization often pushes the limits of resource consumption and the level of parallelism significantly further than most normal OS application combinations and these technologies are not yet highly standardized. Different guest OS and application combinations can stress different functions that a certain virtualization engine take advantage of when it is running on a certain hardware, while another application might stress other functions.

This combination of hardware, virtualization engine and application can often affect scalability and, just because a solution scales in a certain way on server A, there is no guarantee that it will scale the same way on Server B.

Identifying Management Tools and Requirements for that Selection

As previously noted, the ability to manage your environment becomes even more critical when virtualization is employed. Some of the more common management issues related to virtualization include the need to:

- Simplify creation of new virtual servers and migration of existing systems into a virtualized environment (for example, VMware Converter),
- Predict and track virtual environments that compete with each other for server and storage resources
- Predict and track performance utilization in real time as well as noting historical trends. This has to be done for individual environments, the host system as well as the SAN system, and preferably, in a way that allows correlation between these components.
- Manage “VM sprawl” By tracking where, why and how virtual applications are running and what resources they are using.
- Make sure the tools can cooperate and integrate with the existing enterprise management software.

If your management tool can handle most of these issues and includes a simple method to do hot migration (VMOTION, Live Migration or similar) you will have a good environment in which to do efficient load balancing between the servers. Although some management tasks can be automated, it is important to be able to predict, whenever possible, the amount of resources that are required *before* they are actually required. To accomplish this you will need a strong understanding of the business systems and occasional human intervention.

Identifying Change Control Methods

One of the major advantages of virtualization is its ability to quickly define and deploy new logical servers. This is useful in many situations – particularly when you need to access a large number of configurations for a limited time, but also in situations where you might need to deploy additional servers for scalability reasons.

Note: Virtualization makes it easy to quickly create a large number of different images, however unless there is a compelling reason to do so, it is advisable to minimize the number of variations of the base systems. This will greatly simplify the process of keeping these images up-to-date and secured.

Regardless of the underlying reasons, when you are building up a directory of images to use it is important to introduce the appropriate change control in order to best manage those images. There are many ways to do this however, regardless of what approach you take you should incorporate mechanisms for the following states:

- **Development**

Phase/library where the images are developed. This could include completely new images or updates to existing images.

When the image is ready for QA it is transferred to *Test / Quality Assurance*.

- **Test / Quality Assurance**

Phase/library where the newly updated/developed images are tested and verified. If further modification is needed they are transferred back to *Development*; otherwise they are subject to an approval process and, if approved, are transferred to the next stage - *Available Images*.

- **Available Images**

Library where the Gold Master images are stored when they are approved for usage.

When the image is to be deployed as a server a copy should be transferred to *Configuration*. If any changes need to be done to an existing image a copy of it should be transferred back to *Development*. Finally, if an image is replaced by an updated version or if, for any other reason, it should not be used, it should be transferred to *Archive*.

This is purely a storage area for the Master images of potential servers. These images should never be active or modified while in this stage.

- **Archive**

Library for Gold Master images that have been marked as “End of Life”

- **Configuration**

Phase where Gold Master images are configured before they are deployed for production.

This configuration step might be as simple as renaming the server and making sure it has the latest approved maintenance or it may include an automated step to apply additional software or prepare it to connect to existing external data. The use of automation for this task can help ensure a well defined baseline, and be sure that the automation also documents the exact state of the configured server. Once the server is configured it is moved into *Production*.

- **Production**

Phase where the images are deployed. This is normally the last stage for any specific image. If you decide to upgrade an image it should be based on the Gold Image located in the *Available Images* phase.

The only updates that should be done to images in the *Production* phase are those that are the result of normal production. This typically includes minor OS and application patches that don't require a reboot; however these should first be applied and verified on the gold master image.

If additional change control is required you might consider implementing staging areas in between one or more of the stages.

Identifying Data Storage Resources and Limitations

As previously noted one of the more important areas to analyze during a virtualization project is how to manage the data storage problem. When you introduce a project like this it is critical that you have a data storage infrastructure to support it. In many ways this can be considered a separate project, however, it is a requirement for any successful virtualization project and it is likely that it will save both money and, more importantly, make it easier to secure your business critical data.

A few challenges to look into in this area include:

- **Multiple I/O intensive applications using the same physical disk**

When multiple data-streams try to communicate to the same physical disk at the same time you are likely to run into severe delays. This is a well known problem within specific applications (for example databases), but you will most definitely run into this problem if multiple virtual machines share the same physical disk.

This risk can be managed by making sure each virtual machine (or better yet, each expected concurrent major data stream) has its own physical disk – preferably by implementing a SAN or iSCSI solution where the disks are correctly allocated to the various virtual machines. Implementing a SAN or iSCSI solution with many smaller disks is preferable in these situations since multiple disks normally provide faster overall performance due to the larger number of heads that can do simultaneous read/write operations.

- **Make disk space accessible from multiple host systems**

To simplify migration of virtual machines between different physical machines it is important that the data is stored in such a way that it is accessible in the same way from each one of the host systems.

This can be easily managed by moving the data storage from the host server to a SAN or an iSCSI solution.

- **Allocate enough storage for each virtual machine**

Like most other resources it is important to have enough data storage available for the application – and to account for both maintenance as well as expected growth.

On a traditional disk system this has been managed by allocating the server significantly more space than it really requires and, thereby, wasting huge amounts of disk space. With an external disk system it is easier to manage this by allocating additional space for a specific virtual machine when required. This can be managed manually, through the storage manager on a SAN or iSCSI solution, however, in some situations, it might also be worthwhile to look into a more formalized system for thin provisioning of storage assets.

True thin provisioning is a very interesting and efficient technology when you have a large number of users who might need access to a lot of storage that they often aren't going to use. Since you also have to give up control, to some degree, over who has access to a certain disk you also give up what may have earlier been predictable performance. Therefore, this needs to be *carefully analyzed* when I/O intensive applications are using the disk. With that in mind, this type of technology can be very efficient since read/write operations from one application might be spread out over multiple disks and, in this way, able to perform parallel I/O operations.

It is worthy to note that sites, like mySpace, are successfully using thin provisioning to allow a large number of users to have access to a large amount of disk - in a controlled manner.

Defining Pertinent Maintenance Tasks

One of the major advantages with virtualization is how it can, with the appropriate planning, reduce downtime for hardware and software maintenance to a minimum. Another advantage is that it can greatly simplify any maintenance tasks since they can be performed on an offline system while the production systems are running on another host.

To truly minimize the downtime associated with maintenance, it is important to have a "hot migration" solution, such as VMware VMOTION, or Microsoft Live Migration (not yet released) that allow you to move virtual machines between hosts while they are still active. Using this functionality, you can move all applications to a secondary system and then, without any interruption or stress, upgrade and test on the now offline original hardware. Once the upgrade is performed and tested, the application can be migrated back to the original system and you can move on to the next system that needs to be upgraded. All of this can be done without any interruption of the production system.

There are also significant advantages with virtualized environments for software maintenance tasks. Any updates to the production environment should go through a change management process to ensure that this update works as expected and to minimize the risk of any interruption of production.

In a traditional environment, this can be a complex task to tackle (in the right way), however, in a virtualized environment with a change management infrastructure implemented, the degree of complexity drops significantly. Although there may still be a brief interruption in production while the production system is switched over to the new and updated image, since this has been carefully tested, the risks are minimized. In addition, if something does go wrong you can switch back to the original system relatively easily.

Minor updates that do not interrupt production processing (e.g., security patches that do not require a reboot) can be applied directly to the production system once they have been tested and applied to the corresponding gold image. This will keep the gold image synchronized with the production and, at the same time, minimize the need for disruption in production. However, even in this situation, it is highly recommended that you first take a snapshot\backup so that you can quickly revert to the original state should a problem occur.

A Word about Virtual Appliances and Virtual Desktops

One of the great advantages with this technology is that, since the virtualized image should typically only be used for a very specialized application, it can be tuned to support that application in the best possible way. This is called a **Virtual Appliance** - and the OS and the environment often might be crippled so that it only supports this specific task.

As with most other solutions there are both advantages and disadvantages with this approach. A few advantages are:

- Tuning or disabling functions to optimize them for the specific task can enable the OS to run more efficiently and, therefore, lead to a faster application that uses fewer resources.
- Disabling certain features that the application does not require can result in a significantly more secure system since it minimizes the potential security holes.
- Support and maintenance may be simplified since there are fewer components that might interfere with each other (e.g., incompatible DLLs or Java versions)

On the other hand, disadvantages are:

- Unless the developer/administrator of the virtual appliance completely understands the requirements it is easy to completely or partially cripple the application's functionality by disabling a service that it requires for certain functions.
- Although Virtual Appliances may simplify support and maintenance, with a highly customized environment it can also be very hard to get support from the application and/or OS provider. Being asked to first reproduce the problem on a standard OS is not unusual.
- Breaking up applications into really small appliances might add a significant cost since common components, such as operating systems and security products, typically require a license for each appliance.

Although this document focuses on *server* virtualization it is worth mentioning the increasing push for **virtual desktop infrastructure** (VDI) as well.

It is a well known fact that most user desktops have extremely low utilization and, in fact, may not be used at all during many hours. However, since they do require quite a few resources when they are used, they still need to have access to those resources. At first glance this sounds like a perfect scenario for virtualization (and in some situations it might be) but there are a number of challenges to consider:

- Are the resources commonly used at approximately the same time?
As with server consolidation you need to analyze the resource usage patterns to determine if users commonly use their desktops at approximately the same time (for example at the beginning and/or end of day or around standard lunch hours)
- What is the network connectivity between the planned clients and the datacenter?
A solution like this obviously requires good low latency in the network connectivity between the clients and the datacenter. This is typically not a problem in a campus environment; however it might be a significant issue in a more distributed environment.

It is worth noting that this is not always a disadvantage. In fact, under certain circumstances (such as with client-server based applications) it can even be an advantage since the computing power is closer to the rest of the server infrastructure.

- **This is new technology**

Although similar solutions have been around for quite a while, they have not been widely accepted as a desktop standard. Since the concept is currently receiving new attention, we can probably expect to see significant enhancements in this area in the coming years. This is obviously an advantage, but it is also a risk since your investments might be leapfrogged by new technology.

Again, this document is not intended to dive into the depths of VDI, but it is important to be aware of it and to realize that many (but not all) of the advantages/disadvantages with server consolidation are also valid for desktop virtualization.

Determining Costs as well as Return on Investment

Gartner Group expects that the number of servers that are built on virtualization technology will grow from the current 540,000 to as many as 4 million by 2009. According to this study, 90% of the users are implementing virtualization to lower to cost of servers, office space and energy, however there are a lot of obstacles in the way and this gain is far from obvious.

A virtualization project can be an expensive exercise unless the business has proper control over the additional costs that may be incurred for additional licenses, newer, more advanced server environments, and infrastructure changes needed to support this new environment and make it coexist with the existing environment. In addition, it is very important to have a clear and well thought out strategy on how to manage this new technology.

Following are some of the more important factors to consider when analyzing the total cost for a virtualization project:

- **New high end servers to support virtualization**

In most situations the virtualized environment needs new and more costly hardware to handle multiple concurrent virtual servers. Such new servers need sufficient NICs with required bandwidth, additional memory, and CPUs.

- **Storage solution**

Unless the company has already invested in a storage solution that supports this environment (such as a high performance SAN or iSCSI solution) this might be a significant additional investment.

- **Adequate infrastructure**

You may need to review the datacenter infrastructure to determine if it contains adequate backbone bandwidth (especially where VMotion is part of the solution or one of the goals of virtualization is enhanced fault tolerance or disaster recovery), adequate power/cooling for the dense rack(s) used for virtualization servers.

- **Education and/or new hire**

To build, manage and support the virtualized environment in an efficient and secure way requires specialized knowledge that is currently unusual and in high demand.

- **Redeploying existing working infrastructure into Virtual Machines**

There is always the risk that you may be fixing something that does not need to be fixed and, if at all possible, this should be avoided, especially in the beginning of any virtualization project.

Based on earlier discussions (see "Areas Where virtualization Provides Quick Value" on page 3) it is generally a good idea to start with a situation where you need to house a library of servers (for example Development, QA, Support and/or Demo scenarios) as well as a secondary library for new projects.

- **Need for new management tools**

Good management tools are more important than ever when it comes to managing virtualized environments, since you need to be able to track (and, when possible, predict) resource consumption by individual virtual machines. This knowledge will allow you to redirect resources and move virtual machines to underutilized hosts before you experience performance problems.

- **License cost for virtualized applications**

Depending on the application and how virtualization is used, the licensing costs for OS and third party applications might be considerable.

It is important to make a careful analysis, based on the applications that should be virtualized. A few things to consider are:

- In a virtualized environment it is common to have many smaller servers (Virtual Appliances), each one of these likely needing an operating system, as well as a number of common applications such as antivirus, backup clients, back office connectivity etc...)
- Some software licenses are bound by how many physical CPUs the hosting server has and not how many VCPU are assigned to the virtual machine. This often makes the licensing more expensive since the host in a virtualized environment likely is to be a larger server.
- Some software is licensed to run on a specific physical server. This might reduce the advantages with the virtualization since you either need to invest in a license for each potential host or disable the VM load balancing functionality.

■ **License cost for virtualization engine**

Licenses for virtualization software typically have a significant price tag and this has to be taken into account. When you investigate prices make sure that the version you are looking at includes support for the tools and technologies you need. The price might differ depending on a large number of factors; a few of the more important are listed below.

- Support for SAN (iSCSI or Fibre Channel based)
- Hot migration support (VMOTION, Live Migration or similar)
- Size of Host server (commonly measured by number of CPUs or sockets)
- Included Upgrade and Support agreement

Chapter 3: Maintenance Considerations

It is important to track the performance profiles of the individual machines, as well as the virtualization hosts, both before and during the lifecycle of any virtualized environment in order to identify potential peak requirement periods. For example, consider an environment in which at the end of every workday, a majority of the users:

- Send an email, synchronize email folders and then logout from the mail server.
- Print out a large document to bring with them and study from home.
- Make backup copies of a number of important files to a folder on a file server.

If the users all share a common host system whose requirements are based on ordinary usage you might expect that these VMs would slow down substantially in response to the higher peak load resulting from the near simultaneous performance hit from these activities by hundreds or thousands of users.

By tracking the consumption of critical resources over time it is possible to identify patterns of resource usage by servers and use that knowledge to pair virtual machines that stress different types of resources or that stress the system at different points in time.

These performance profiles should be used to establish a baseline that can be monitored so that you can see how resource utilization changes over time. This knowledge can provide an excellent planning tool since it will highlight which virtual machines are requiring more or less resources over time.

VMware currently provides a number of tools to assist with management of application performance; however you might consider additional performance management tools as well. Regardless of the tools you select, you should also incorporate careful planning along with personal experiences with how certain applications are likely to require additional resources. This is especially important since many applications often have their peak load at the same time and having multiple machines request more resources at the same might put significant load on the host.

Optimization Strategies

Combining performance profiles with knowledge about the company's activities allows the system manager to optimize the resources for individual virtual machines as well as to make sure they are paired with other virtual machines that, if at all possible, are using different resources or experience peak load at different times.

It is important that the system manager have a deep understanding of the overhead of various virtual machines and how tuning options affect the overall health of the environment. Typical options for optimizing virtual resources include:

- Allocating dedicated resources
- Over-allocating\sharing resources
- Re-allocating existing resources between VMs
- Moving an existing VM to a new host
- Optimizing and tuning the individual VMs

There are several documents available on this topic and it is highly recommended that the system manager become familiar with them. Following are two examples:

- "Memory Resource Management in VMware ESX Server"
(www.vmware.com/pdf/usenix_resource_mgmt.pdf)
- "Performance Tuning Best Practices for ESX Server 3"
(http://www.vmware.com/pdf/vi_performance_tuning.pdf). This document references an additional document which includes further details.

Allocating Dedicated Resources

Although most modern enterprise class virtualization hosts include intelligent methods for sharing over-allocated resources among guests, there are a number of physical limitations to monitor. In situations where resources are known bottlenecks, but where predictable response times are required, it is a good idea to dedicate hardware for individual virtual machines – or, at least, to make sure that there are always resources available.

It is important to identify which virtual machines require which type of dedicated hardware in order to manage resource consumption. Potential candidates include:

- **Physical Disks**

Disk I/O is a common bottleneck and it is highly recommended that you have dedicated physical disks (preferably through SAN or iSCSI).

Certain technologies for 'thin provisioning of data storage assets' spread the load over multiple physical discs in an efficient way, however, this needs to be carefully analyzed before use with business critical applications.

- **RAM memory**

Even though many modern virtualization hosts have a very intelligent memory management system it might be unwise to extensively over allocate memory if your applications are memory constrained.

With that in mind, if you have somewhat similar environments and/or if the applications are not fully utilizing memory you can typically employ some over allocation and rely on technologies, such as memory sharing, or advanced techniques for dynamic allocation of memory (for example ballooning within VMware).

- **Network Cards (NICs)**

NICs are often shared between virtual machines, however, if a guaranteed response time is required you should consider allocating dedicated NICs to your critical and network intensive applications.

In addition, unless management traffic and hot migration events are infrequent, it is considered best practice to have dedicated NICs for the Service Console and hot migration tools, such as VMOTION or Live migration.

Over-allocating / Sharing Resources

To best take advantage of the hardware, it is often possible to over allocate resources so that various guest OSs share a limited set of resources. In these situations the host system is responsible for scheduling currently available resources between the logical systems.

As noted previously, it is best to proceed cautiously if any of the virtual machines are resource intensive and you require predictable response times. However, when there are resources available, the hosts/hypervisors are typically fairly good at efficiently sharing the over allocated resources.

In addition to the NICs and RAM memory, which are described in the previous section, **CPUs** are likely candidates for shared\over allocated resources. The available CPUs are almost always shared between the virtual machines and automatically scheduled between the active virtual machines.

Note: Even though you can assign multiple CPUs to a virtual machine, this is not a good idea unless the application truly utilizes multiple CPUs. The main reason for this is that the host will wait until it has all allocated resources available before being dispatched. In addition, a guest system with many CPUs will tie up those CPUs even if it is only using one or two. This is one of the reasons why it often is better to have multiple smaller VMs than a few huge logical virtual servers. For example, if a virtual machine is assigned 4 CPUs the host can not assign any additional CPUs until all 4 CPU are available. Furthermore, it is critical to ensure that no single virtual machine can have dedicated access to all of a host system's resources of a certain type as this would block anyone else, including the host kernel, from accessing those resources. For example, never allow a virtual machine to have 4 VCPU on a 4 CPU host or have access to 4 Logical Network cards on a host system with 4 physical NICs.

Re-allocating existing resources between VMs

Although similar to the previous two options, reallocating existing resources between VMs requires a more thorough analysis of the ongoing and expected resource consumption of the logical servers. VMware includes utilities to handle this automatically but, typically, this requires manual intervention – and this is where performance profiles and baselines can be truly useful. The key here is that we are talking about *reallocating* resources. If you add resources to one virtual machine you typically need to remove the corresponding resource from another logical machine on the same host.

Move an existing VM to a new host

After analyzing the performance profiles for both virtual machines and the actual host system you might consider reorganizing which virtual machines are running on which host systems - or you may decide to move all virtual machines to a new host to allow for maintenance.

If you have systems with “hot migration” (such as VMOTION, Live Migration or similar) and a well planned out disk storage system on SAN or iSCSI this is not a very complex task. However, it should normally not be done to handle a short term peak. Rather, the ideal is to be proactive and to plan ahead to see when a guest OS needs to be transferred to a new system.

Regardless of the optimization approach you take, it is important to realize that virtualization always adds overhead. You cannot simply add the resource requirements for the individual guests to derive an estimate for the host system. You also need to consider peak period requirements and remember that applications hit their peaks at the same time or during similar situations.

Other important items to consider are:

- According to Network Computing (April 2nd 2007) the overhead of VMware ESX is typically less than 10% (low of 6%, highs of 20%).
- Make sure you provide sufficient memory for memory intensive applications (such as IIS, and SQL depending on the query type). However, avoid over-allocating it on applications that do not use it.
- Having multiple applications which employ intensive disk-access (where the bottleneck is I/O) will significantly affect performance unless you can off-load each application's data storage to a dedicated physical disk
- To be able to efficiently manage the VMs you need to have the virtual disks on separate disk systems using high performance iSCSI or SAN solutions. This will significantly simplify your management of virtual machines and their disk requirement, both by allowing you to grow/shrink required disk space and, more importantly, to simplify any migration of virtual machines between host systems.

Optimizing and Tuning the individual VMs

Although all of these measures are important, you should not overlook the importance of analyzing the actual VM machine to determine if it would benefit from basic internal tuning steps.

In a non-virtualized environment this type of tuning is recommended, but often not critical; however, since any unnecessary usage of resources in a virtualized environment penalizes other applications it becomes critical if the virtualized environment is to run efficiently.

A few examples are:

- **Disable/uninstall screensavers and other unnecessary applications**

During the creation of any server system you should avoid installing, or even uninstall, any applications that are not going to be used. This applies to all types of application and utilities since this minimizes the risk for security exposures. However, it is *especially* important for tools that might perform heavy calculations, such as 3D screensavers or multimedia tools.

Note: No server needs a 3D screensaver and very few servers require a media player. This true for both virtualized and non-virtualized servers.

- **Disable services and daemons that aren't in use and/or required**

One of the great advantages with virtualized servers is that it is easier to have complete control over the system and, therefore, it is easier to remove unused services and daemons from "Virtual Appliance" If this is done correctly it can enhance both security and performance significantly.

Note that this does require quite a good understanding of both the application as well as the guest operating systems and, any changes that are made require a significant quality assurance effort.

■ **Disable/disconnect unnecessary devices in host and guest systems**

Devices connected to the guest and / or the host system consume various types of resources and should, therefore, be disabled or disconnected on systems that aren't using them.

A few examples of resources that might be considered for this are:

- Serial (COM) and Parallel (LPT) ports
- Floppy/CD/DVD drives
- USB / FireWire Adapters
- Network Interface Cards

Disabling these devices frees up IRQ resources and eliminates conflicts. Many of these resources also consume CPU through polling especially in a virtualized environment,. Finally, having multiple virtual machines polling the same physical device can sometimes cause delays due to contention.

Note: When these changes are made on a host that is part of a load balancing scheme it is important to make sure that none of the potential guests requires the specific resource. If not, the migration of this guest will likely fail.

■ **Schedule supporting back office jobs during off-peak hours**

This is something that is likely to have been done long before any virtualization project was implemented so the chances are good that a lot of the basic data on when to run various procedures is already available. However, this task is significantly more complex in a virtualized environment since you now need to take into consideration the load from all other virtual machines running on the same host. And, to complicate this further, since the organization is also likely to be using some type of dynamic load balancing (such as "hot migration") this is a moving target.

As a result, scheduling procedures, such as antivirus updates and backups, in a virtualized environment is not a trivial task. However, a good performance reporting tool can provide a huge advantage when it comes to estimating the load on both the individual virtual machines as well as the host system while the various back office procedures are running.

Appendix A: References and Links

The following resources were used in this document:

- **VMware Whitepapers, technical notes etc...**

- Proceedings of the 5th Symposium on Operating Systems Design and Implementation (Memory Resource Management in VMware ESX Server) by Carl A. Waldspurger, VMware, Inc., available at:

http://www.vmware.com/pdf/usenix_resource_mgmt.pdf

- Performance Tuning Best Practices for ESX Server 3, available at

http://www.vmware.com/pdf/vi_performance_tuning.pdf

This includes:

- Overview VMware specific technical resources

www.vmware.com/resources/

www.vmware.com/support/pubs/vi_pubs.html

- Mismatched HALs – KB Article 1077

(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1077>)

- IRQ Sharing – KB article 1290

(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1290>)

- Idle Loop – KB article 1730

(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1730>)

- CPU Utilization – KB Article 2032

(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=2032>)

- Network Throughput between Virtual Machines – KB article 1428

(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1428>)

- Queue depth on QLogic – KB article 1267

(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1267>)

- Guest Storage Drivers - KB article 9645697
(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=9645697>)
 - Disk Outstanding commands parameter - KB article 1268
(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1268>)
 - Linux Guest Timing - KB article 1420
(<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1420>)
 - Timekeeping in VMware Virtual Machines
(www.vmware.com/resources/techresources/238)
 - Virtual SMP Best Practices – “Best Practices using Virtual SMP”
(www.vmware.com/resources/techresources/240)
 - VMFS partitions – “Recommendations for aligning VMFS Partitions”
(www.vmware.com/resources/techresources/608)
 - Esxtop – Using esxtop to Troubleshoot Performance Problems”
(www.vmware.com/resources/techresources/436)
 - SAN best practices – “SAN Configuration Guide”
(www.vmware.com/pdf/vi3_301_201_san_cfg.pdf)
 - NFS best practices – “Server Configuration Guide” and iSCSI best practices – “Server Configuration Guide”
(www.vmware.com/pdf/vi3_301_201_server_config.pdf)
 - Memory allocation and Swap Space – “Resource Management Guide” (www.vmware.com/pdf/vi3_301_201_resource_mgmt.pdf)
- **3PAR Thin Provisioning**
<http://www.3par.com/documents/3PAR-tp-wp-01.2.pdf>
 - **Clabby Analytics Research**
“CA's Unicenter Advanced Systems Management: Virtualization Cluster Management for Heterogeneous Environments”
<http://resources.technewsworld.com/technewsworld/search/viewabstract/87259/index.jsp>
http://ca.com/files/IndustryAnalystReports/cas_advancedsystemsmanagement_final.pdf
 - **SWSOFT Virtuozzo - White Paper**
“Top Ten Considerations for choosing a server virtualization technology”
<http://whitepapers.zdnet.com/whitepaper.aspx?docid=148305>
 - **Baseline** (Where leadership meets technology)
 - March 2007 (issue 070) “Virtualization Beyond the Buzz” by Michael Vizard <http://www.baselinemag.com/article2/0,1540,2100143,00.asp>

- April 2007 (Issue 071)
<http://www.baselinemag.com/article2/0,1540,2113063,00.asp>
 - "Virtualization: For Servers, a Disappearing Act" (Less hardware is good, but immature technology isn't. With new products on the way, companies may find better methods to speed up virtual server deployments and cut costs.) by Brian P. Watson
 - "Microsoft: Ramping Up" by Brian P. Watson
 - "VMware: Speeding Ahead" by Brian P. Watson
 - "XenSource: Aiming High" by Brian P. Watson
- **CIO** – Business Technology Leadership
 - June 15th 2007
"Thinking Inside the Boxes" or "How Server Virtualization Tools Can Balance Data Center Loads" by Katherine Walsh
www.cio.com/article/117256/
 - "ABC: An Introduction to virtualization"
www.cio.com/article/40701/
 - "Taking Virtual Servers to the next level"
www.cio.com/article/122950/
 - "Server Virtualization Snapshot" www.cio.com/article/106950/
 - "The Virtues of Virtualization" www.cio.com/article/11855/
 - "The Benefits of Consolidation and Virtualization"
www.cio.com/article/21970/
 - July 15th 2007
"Virtual Possibilities" or "Taking Virtual Servers to the Next Level" by Thomas Wailgum
www.cio.com/article/122950/ - same article as "Taking Virtual Servers to the next level" from June 15th
- **CommunicationsNews** May 2007:
 - "Virtualization Takes Hold" (Enterprise data centers embrace new/old technology to better utilize resources, but security, compliance and mobility issues must be addressed.) by Jeff Jilg
http://comnews.com/stories/articles/0507/0507coverstory_virtulization.htm
 - "The Namespace Option" by Panos Tsirigotic:
http://comnews.com/stories/articles/0507/0507namespace_option.htm
 - "Virtual machines go mobile" (Optimization software shrinks files to improve portability and performance.)
http://comnews.com/stories/articles/0507/0507v_machines_go_mobile.htm

- "Security Rules Have Changed" (New solutions are necessary to protect virtualized networks.) by John Peterson
http://comnews.com/stories/articles/0507/0507security_rules.htm
- "The phantom menace: Security (Protecting the virtualized network is different from previous network security practices)" by Allwyn Sequeira
http://comnews.com/stories/articles/0507/0507the_phantom.htm
- Related Online material:
 - ◆ "Virtualized versus traditional blade platforms" (Server consolidation has become a key strategy for reducing data center costs.)
http://comnews.com/stories/articles/0507web/0507tehuti_networks.htm
 - ◆ "Virtualization and IT service assurance" (Proper monitoring and management of virtualized environments is essential to address performance problems.)
http://comnews.com/stories/articles/0507web/0507network_general.htm
 - ◆ "Virtually insecure" (Address the security implications of a disruptive technology.)
http://comnews.com/stories/articles/0507web/0507bt_ins.htm
 - ◆ "Common mistakes when consolidating servers" (Avoid potential pitfalls by planning ahead.)
<http://comnews.com/stories/articles/0507web/0507cirba.htm>
- **ComputerWorld**
November 2007 "CA Launches virtualization management tool" by Matt Hambleon
<http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9005389>
- **Custom Systems Magazine**
March 2007 "Fitting More Apps into fewer boxes" by David Gilbert.
- **eWeek** (http://www.eweek.com/previous_issues/)
 - November 27th 2006
"Unicenter Goes Virtual" By Paula Musich
<http://www.eweek.com/article2/0,1759,2064622,00.asp>
 - March 12th 2007
 - "Competition heats up" (VMware, Virtual Iron, Parallels roll out enhanced virtualization offerings) by Scott Ferguson
 - "Isilon boosts storage management" (Data migration between multiple tiers of clustered storage) by Chris Preimesberger

- April 2nd/9th 2007
 - “Xen Expansion” (Open Source Virtualization project moves from bleeding edge to deployment ready) by Jason Brooks
<http://www.eweek.com/article2/0,1895,2107810,00.asp>
 - “Iron-clad server virtualization” (Virtual Iron 3.5 offers a flexible low cost solution)
<http://www.eweek.com/article2/0,1895,2107811,00.asp>
 - Linux Kernel to add VMI
<http://www.eweek.com/article2/0,1895,2107818,00.asp>
- May 7th 2007
 - “Stretching your resources” – (Amazon.com elastic compute cloud allows ad-hoc VM deployments) by Jason Brooks
<http://www.eweek.com/article2/0,1895,2124973,00.asp>
 - Enomalism helps manage VMs
<http://www.eweek.com/article2/0,1895,2124901,00.asp>
 - esxRanger ably backs up VMs
<http://www.eweek.com/article2/0,1895,2124944,00.asp>
- May 21st 2007
http://www.eweek.com/current_issue/0,1913,i=1936,00.asp
 - “Virtualization comes into focus” (desktop users look for better pc management) by Scott Ferguson
<http://www.eweek.com/article2/0,1895,2131524,00.asp>
 - VMware adds Linux’s Paravirt-ops Virtualization
<http://www.eweek.com/article2/0,1895,2128344,00.asp>
 - Virtual Iron, provision networks take on desktop virtualization
<http://www.eweek.com/article2/0,1895,2114414,00.asp>
 - uXcomm Takes on Virtualization with acquisition
http://www.eweek.com/print_article2/0,1217,a=204426,00.asp
 - XenSource prepares latest Virtualization Release
<http://www.eweek.com/article2/0,1895,2109856,00.asp>
 - Novell adds Virtuozzo Virtualization to SLES
<http://www.eweek.com/article2/0,1895,2104847,00.asp>
- June 11th 2007
http://www.eweek.com/current_issue/0,1913,i=1943,00.asp
 - “Expanding Virtualization’s reach” or “Vendors are attempting to expand the reach of virtualization” – (VMware rolls out utility computing for hosting companies) by Jeffrey Burt & Scott Ferguson
www.eweek-digital.com/eweeek/20070611_stnd/data/eweeek20070611_stnd-dl.pdf

“VMware expands Virtualization Options” by Jason Brooks

<http://www.eweek.com/article2/0,1895,2142436,00.asp>

“VMware releases full version of ACE 2 Software” by Scott Ferguson

<http://www.eweek.com/article2/0,1895,2130437,00.asp>

“VMware tries it hands at utility computing” by Scott Ferguson

<http://www.eweek.com/article2/0,1895,2142149,00.asp>

– “ClearCube, VMware partner on Virtualized Desktops” by Scott Ferguson

<http://www.eweek.com/article2/0,1895,2136424,00.asp>

– “VMware latest Virtualization Software supports Vista” by Scott Ferguson

<http://www.eweek.com/article2/0,1895,2128067,00.asp>

– “VMware’s VI3 Suite Delivers on Virtualizations promise” by Jason Brooks

<http://www.eweek.com/article2/0,1895,2095585,00.asp>

– “IBM Launches new Virtualization Tools” by Scott Ferguson

<http://www.eweek.com/article2/0,1895,2094523,00.asp>

– “Adoption of Virtualization Continues to grow: Report” by Scott Ferguson

<http://www.eweek.com/article2/0,1895,2094148,00.asp>

■ **Information Week**

■ Dec 18th/25th 2006 “Virtual Payoff In the Real World” by Charles Babcock

– “Making the best of Both Worlds (Virtualization software brings the same management challenges as physical serves to virtual machines)” by Charles Babcock

<http://informationweek.com/1119/virtual.htm>

– “Microsoft Opens Virtualization Standard in Gambit Against VMware (Microsoft has previously opened specs in only Web services arena, where it faces tougher open standards competition.)” by Charles Babcock

<http://informationweek.com/1119/vmware.htm>

– “Microsoft And Xen Team For Virtualization (With this move Microsoft is expanding its support of virtualization in its most advanced software, and it’s doing so in a way that uses fewer system resources than past approaches)” by Charles Babcock

<http://informationweek.com/1119/xen.htm>

- March 12th 2007 "Desktop Virtualization: VMware Eyes New Pastures" (Virtualizes hundreds or thousands of PCs and Notebooks, sends a set of virtualized files over the network-executed locally, Remote Control, Pocket ACE allows virtual desktops to be loaded on iPods or flash drives...) by Charles Babcock
 - "VMware Accuses Microsoft of Restricting its Customers (Microsoft's virtualization technology shuts out third parties, the vendor claims)" by Charles Babcock (March 10 2007)
<http://informationweek.com/1128/vmware.htm>
- April 9th 2007 "Citrix Seizes The Moment For Desktop Virtualization" (VMware ACE 2, Citrix Desktop Server to serve desktops for the end users) by Charles Babcock
 - "VMware Looks To Virtualize The Desktop (Virtualization may make desktops and laptops more secure and easier to manage)" by Charles Babcock (March 10 2007)
<http://informationweek.com/1129/vmware.htm>
- June 11th 2007
"How 9 Hot Technologies can blow up in your face" / "Virtualization Threats Ahead" and VMwares new equation (Virtualization + SAAS) by Charles Babcock (cbabcock@cmp.com)
<http://www.informationweek.com/software/showArticle.jhtml?articleID=199902576>
- June 25th 2007 "Virtual Desktop May Take Awhile To Become Real" by Charles Babcock
<http://www.informationweek.com/story/showArticle.jhtml?articleID=200000171>
 - VMWARE'S APPROACH
<http://informationweek.com/1129/vmware.htm/>
 - HP, VMware Teaming Up To Use Server Virtualization To Replace PCs
<http://www.informationweek.com/showArticle.jhtml;jsessionid=FT44IMWWJZSXSQSNLRSKHSCJUNN2JVN?articleID=193600083>
- July 16th 2007 "A Virtualization Bargain" by Joe Hernick (jhernick@nwc.com)
www.informationweek.com/1146/xenenterprise.htm
 - Full Assessment: Find our testers notebook.
www.informationweek.com/1146/notebook.htm
 - More Numbers: Get the complete result of our tests:
www.informationweek.com/1146/benchmark.htm
 - Take a Look: Screenshots at XenEnterprise:
www.informationweek.com/1146/gallery_xen.htm
 - Microsoft Play: Viridian coming soon, short some key features:
<http://nwc.com/go/ms-virtual/>

- What's next? All eyes are on I/O virtualization technology:
<http://nwc.com/go/iov/>
- July 23rd 2007
"Virtual Machines in motion: Live Migration Adds to Appeal" by Charles Babcock
<http://www.informationweek.com/story/showArticle.jhtml?articleID=201200284>
- It all adds up: The practical realities of virtual data centers:
www.informationweek.com/1146/datacenters.htm
- Virtual XEN: open Source virtualization software is a bargain:
(same article as "A Virtualization bargain" from July 16th)
- Friends in high places: Intel Invests \$218 million as VMware preps IPO:
www.informationweek.com/1145/intel.htm
- August 13th 2007
"Virtualization: Key To Linux's Future or a Linux Killer" by Antones Gonsalves (antoneg@pacbell.net)
www.informationweek.com/showArticle.jhtml?articleID=201400215
- Distribution lists
 - http://www.informationweek.com/software/showArticle.jhtml?articleID=199202005&cid=RSSfeed_IWK_News "SUN SEES VIRTUALIZATION WITHOUT VMWARE OR XEN" (The company's investment in its Live-Star project could translate into improved software and security distribution models.)
- **InfoStor**
 - March 2007 - "HP Enhances Virtualization, security"
http://www.infostor.com/display_article/284541/23/ARTCL/none/none/HP-enhances-virtualization,-security
By Kevin Komiega
 - June 2007 - "EMC Unites SRM, VMware" by Dave Simpson
http://www.infostor.com/display_article/292582/23/ARTCL/none/none/EMC-unites-SRM,-VMware/?dcmp=ISWNUS_ARCH
- **InfoWorld**
 - March 26th 2007 "Desktop Virtualizers Vie for Position (Four Competing solutions demonstrate potential, need to grow)" Product reviewed include VMware Workstation 6.0 Beta 3, Parallels Workstation for Windows 2.2, Microsoft Virtual PC 2007, InnoTek VirtualBox 1.3

- Distribution Lists (www.infoworld.com (From Newsletter))
 - PolyServe WhitePaper: EXTEND VIRTUALIZATION TO YOUR MOST IMPORTANT APPS
<http://newsletter.infoworld.com/t?ctl=16F1519:62116BDDDBABB2BF17523F0755AA7BD34EFF29049075316B4> and
http://www.polyserve.com/pdf/Virtualization_Utility_Whitepaper.pdf?utm_source=InfoWorld&utm_medium=Newsletter&utm_content=Virtualization
 - INFOWORLD VIRTUALIZATION REPORT PODCAST (set of podcasts)
<http://www.infoworld.com/weblog/podcasts/dellvirtualization.html>
 - INTERVIEW: VIRTUALIZATION MANAGEMENT Q&A WITH OPSWARE CTO TIM HOWES
http://weblog.infoworld.com/virtualization/archives/2007/04/interview_virtualization.html?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05
 - MICROSOFT ANNOUNCEMENTS AROUND VIRTUALIZATION PRODUCTS (podcast)
http://weblog.infoworld.com/virtualization/archives/2007/04/microsoft_announcements_around_virtualization_products.html?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05
 - VMWARE WORKSTATION 6.0 REACHES RELEASE CANDIDATE STAGE
http://weblog.infoworld.com/virtualization/archives/2007/03/vmware_workstation_6_0_reaches_release_candidate_stage.html?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05
 - INTEL'S PENRYN PROCESSOR FAMILY BOOSTS VIRTUALIZATION PERFORMANCE
http://weblog.infoworld.com/virtualization/archives/2007/03/intels_penryn_processor_family_boosts_virtualization_performance.html?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05
 - MICROSOFT ANNOUNCED SCVMM BETA 2 FEATURES
http://weblog.infoworld.com/virtualization/archives/2007/03/microsoft_announced_scvmm_beta_2_features.html?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05
 - NEARLY HALF OF VIRTUALIZATION PROJECTS ARE UNSUCCESSFUL
http://weblog.infoworld.com/virtualization/archives/2007/03/nearly_half_of_virtualization_projects_are_unsuccessful.html?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05?source=NLC-VIRTUALIZATION_REPORT&cgd=2007-04-05
- **Network Computing** April 2nd 2007 "The Virtualization Drag (VMware can help you save money and gain flexibility, but there will be trade-offs. We took to the lab to determine where performance hits will come from, and how to minimize them)"
<http://www.networkcomputing.com/showArticle.jhtml?articleID=198700359>
by Michael Caton

- **Network World**

- November 27th 2006
"CA Unicenter to manage Virtual Servers" –by Denise Dubie
<http://www.networkworld.com/news/2006/112706-ca-unicenter-virtual-servers.html>
- June 11th 2007
"10 Free virtualization tools worth a look" (including links) by Dennis Connor
<http://www.networkworld.com/news/2007/060507-free-virtualization-tools.html>
 - Why Virtualization is cool - Virtualization is heavily over hyped
<http://www.networkworld.com/podcasts/panorama/2007/05/052107pan-winblad.html>
Podcast/MP3 – Ann Winblad interviewed by Beth Schultz
 - Microsoft "Shoot to high" on virtualization, says exec by John Fontana, May 17th 2007
<http://www.networkworld.com/news/2007/051706-microsoft-virtualization.html>
 - "Microsoft lays out Windows server road map" by John Fontana, (May 16th 2007)
<http://www.networkworld.com/news/2007/051607-microsoft-windows-road-map.html>

- **Microsoft**

- www.microsoft.com/technet/technetmag/issues/2006/07/InsideMSFT/
Inside Microsoft.COM (Managing SQL Server 2005 Peer-to-Peer replication) by David Lindquist

- **Smart Enterprise**

- Spring 2007
"Virtually Efficient" (Virtual servers can bring greater efficiency and flexibility. But CIOs must first have a strategy to effectively manage these new environments.) by Amy Larsen DeCarlo
<http://smartenterprisemag.com/articles/2007spring/enterpriseitmanagement.jhtml>

Appendix B: Virtualization Checklists

Any virtualization project requires thorough planning and careful consideration of which applications to virtualize, in what order and pace to virtualize them, what hardware to use, how to configure the environment and who should be responsible for various parts of the project.

This appendix presents several checklists that can be used as a starting point for planning your virtualization project. These checklists are far from complete, however, our intent is to highlight some of the more important questions that the project team needs to consider. It is highly recommended that the project team conduct additional brainstorming meetings to further evaluate other factors that are important for your organization's particular needs.

Identify Virtualization Candidates

A key part of any virtualization project deciding which servers / applications should be virtualized and prioritizing when and in what order they should be virtualized.

It is highly recommended that all of these considerations be included in the plan, but the actual virtualization implementation should be realized in phases using an open project plan that can take advantage of any lessons learned.

List and Prioritize Servers/Applications/Solutions

List all Servers / Applications / Solutions that you are planning to virtualize:

Server / Solution	Description	Timeframe	Priority

For each one of these, answer and/or consider the following questions:

Generic Questions for Guest Systems

Item	Comments
<p>List all applications that will be included in this Virtual Machine.</p> <p>Note: Unless the applications require or benefit significantly from being co-located it is recommended that you separate applications into private guest systems.</p>	
<p>Acceptance from primary application owner/users?</p>	
<p>Status/prior experience with these applications? For Example:</p> <ul style="list-style-type: none"> ▪ Is it a new application? ▪ Will a new release of the application to be installed in virtual environment? ▪ Will an identical release of the application be moved into a virtualized environment? <p>Note: If possible/applicable gather performance data from an existing implementation of the application. Also, take note of other caveats and experiences from earlier implementations.</p>	
<p>What are the business drivers for virtualization?</p> <ul style="list-style-type: none"> • Saving Energy • Saving on Hardware Budget • Saving Floor Space • High Availability • Maintenance / Management • Load Balancing / Flexibility • Security • Other... <p>Note: Document all applicable business drivers, but prioritize them to indicate order of importance.</p>	

Item	Comments
<p>Classify/describe these applications and identify where they fit into "Areas where virtualization provides quick value" as described in "Virtualization Best Practices".</p> <p>Examples could be:</p> <ul style="list-style-type: none"> • Test Environments / QA / Support • Demo Environment • Low utilized application server • Application Server with predictable load • Application Server with unpredictable load • Application Server with consistent high load • Other... 	
<p>Specify expected load on the system:</p> <ul style="list-style-type: none"> • If possible reference back to historic data • Specify peak period and expected fluctuation during off-peak periods 	
<p>Will the virtual machine always be active/enabled?</p> <p>If not, document the expected schema or rules of engagement when it will be active.</p>	
<p>Does the application have real time or near real time requirements?</p> <p>Note: Virtual machines are typically not suited for real time applications.</p>	

Configuration of Guest System

Item	Comments
Operating System for the virtual machine	
Ensure that the OS listed above is supported by the host system.	
How many VCPUs should be assigned to the guest? Caution: Multiple CPUs can be used; however their configuration should be fact based. Do not over allocate the number of VCPUs without a good reason. Also make sure the host has at least twice as many CPUs as the VM with the most assigned VCPUs.	
Configure virtual machines that only use one CPU to use a Uniprocessor version of the HAL/kernel.	
How much RAM will be associated with this guest? Caution: If a predictable result is required, ensure that all configured RAM is available when needed (for VMware look at Shares vs. Reservation of resources). Hypervisors often have an intelligent memory management system; however over-allocation of resources introduces a significant risk to affect performance.	
If feasible, avoid allocating more than 896 Mb to Linux based guest systems. This to make the memory management in the Linux kernel more efficient.	
How many NICs are needed by this guest and of what speed ? Will they be using dedicated or shared physical NICs? Note: Sharing physical NICs between guest systems is often acceptable, but dedicated NICs are recommended for network intensive applications that require predictable performance.	

Item	Comments
<p>Use a minimalist OS</p> <ul style="list-style-type: none">• If Windows Server 2008 is used consider the CORE version.• If using Linux/Using consider disabling the windows manager and other components that aren't required.	
<p>Disable or uninstall unnecessary applications.</p> <p>This can include:</p> <ul style="list-style-type: none">• All screensavers• Multimedia tools• Auto started but unused gadgets <p>Caution: Make sure these aren't required for any part of the application.</p>	
<p>Disable Services / Daemons that aren't required.</p> <p>A few examples to consider are:</p> <ul style="list-style-type: none">• Web Servers (IIS, Apache, etc)• Indexing Servers• SMTP/POP3• SSH/FTP/TELNET• File and Printer Sharing• Messenger <p>Caution: Make sure these aren't required for any part of the application.</p>	

Item	Comments
<p>Disable or disconnect unnecessary devices on guest and host systems.</p> <p>This will free up IRQs and minimize the risk for queuing and collisions. In addition it will also free up CPU resources since these would otherwise need to be polled by the virtual machines.</p> <p>Devices to consider are:</p> <ul style="list-style-type: none"> • Serial & Parallel ports • Floppy/CD/DVD/BlueRay Drives • USB/Firewire Adapters • Unused Network Interfaces <p>Caution: Make sure these aren't required for any part of the application.</p>	
<p>Schedule supporting back office jobs during off-peak hours.</p> <p>Note: Off peak should be calculated globally on the complete host and not just in each individual guest.</p>	

Operating System with Tightened Security / Limited Functionality

It is not unusual for Virtual Machines to have a tight lock down schema applied to them or to have certain functionality disabled to save resources and enhance security.

This can be done by blocking certain functions or by disabling functions, such as windows managers or various services /daemons. When this is done it is critical that you verify that this doesn't adversely affect the applications that are running on this guest system.

To reduce the chances for error, you should gather the following information:

- List all non-standard security lock down settings on the guest server

Examples of this could be:

- Encryption modules
- Modified Enhanced File system security
- Hardened requirement on passwords
- Limited access to Administrator / root account
- Local/Private Firewall functionality

- List all Services / Daemons that are not activated on the guest system

Examples of this could be:

- IIS
- SMTP/POP3
- SSH/FTP/TELNET
- File and Printer Sharing
- Messenger

- List all other functionality that is disabled on this system

Examples of this could be:

- Windows manager
- Multimedia applications
- USB / Firewire adapters
- Floppy/CD/DVD/BlueRay Devices

Note: The examples listed above are not, in any way, intended to be a complete list. Rather, they are merely representative of the type of changes that are of interest.

For each one of these list items verify that the changes do not affect any part of the hosted applications and that the application still is fully supported after these changes are made.

Guidelines for Host Systems

General Questions for Host Systems

Item	Comments
What type of virtualization engine is being used? Examples: ESX 3.0, Hyper-V, XenServer, etc...	
What guest systems will typically run on this host? Note: This should be mapped to the guest system specified in the section above.	
Are the guest-systems statically located on this host or are there plans to dynamically re-allocate them between hosts? If so, what methods are planned to be used for this? <ul style="list-style-type: none">• Manual Steps / Scripts• Quick Migration...• VMotion, XenMotion, Live Migration• Other Tools	

Basic Hardware Configuration

Item	Comments
<p>How many CPUs / Cores will be used?</p>	
<p>Total number of simultaneously used VCPUs allocated to guest systems from this host.</p> <p>Caution: Do NOT excessively over-commit process resources.</p> <p>The acceptable ratio depends on the load on the individual guest system. It is important to remember that the virtualization itself adds a 6-20% overhead.</p>	
<p>Is hyperthreading enabled in BIOS?</p> <p>Hyperthreading does not provide the same effect as a separate CPU/Core but it will provide a performance improvement.</p>	
<p>Are all hosts in a single Hot Migration pool (VMotion, Live Migration, XenMotion) using the same CPU architecture?</p> <p>Consult the documentation provided with the specific virtualization engine to determine the exact limits. Typically, the CPU requires the same architecture, but minor details, such as clock frequency, can differ.</p>	
<p>How much RAM will be available on this host?</p> <p>For predictable performance make sure the host has more RAM than the total amount of memory that will be used by the host system plus the sum of the guests systems.</p> <p>Note: Hypervisors have intelligent memory management techniques; however, over-allocation of memory can force swapping which significantly reduces the performance.</p>	

Item	Comments
<p>How many NICs are available and with what performance specification?</p> <p>Having multiple high performance server class NICs available is highly recommended.</p> <p>Dedicated NICs are highly recommended for guest systems with significant network requirements in order to guarantee a predictable performance.</p>	
<p>Ensure the NIC drivers in the host and guest are correctly configured when it comes to:</p> <ul style="list-style-type: none"> ▪ Auto-negotiated speed or the set speed ▪ Half or full duplex mode 	
<p>Multiple network adapters configured for NIC teaming provides increased performance and failover.</p>	
<p>Configure the system to have dedicated NIC for the host systems management part (for example the Service Console in VMware)</p>	
<p>To ensure best performance the recommendation is to use network adapters supporting the following hardware features:</p> <ul style="list-style-type: none"> ▪ TCP Checksum offload ▪ TCP segmentation offload (TSO) ▪ Handling of high memory DMA (i.e. 64-bit DMA addresses) ▪ Handle multiple Scatter/Gather elements per Tx frame 	

Disk Configuration

Item	Comments
<p>What type of disk system will be used ?</p> <p>For example: Fibre Channel or iSCSI based SAN, NAS or Direct Attached Storage.</p> <p>Note: A SAN solution is highly recommended in order to take full advantage of the virtualized environment.</p>	
<p>If an iSCSI based SAN is used:</p> <ul style="list-style-type: none"> • Use Gigabit network- if possible, 10 Gigabit Ethernet. • Multiple network adapters ensure better speed and redundancy. • Use iSCSI TOE NIC or an iSCSI HBA. The TCP/IP Offload Engine (TOE) or Host Bus Adapter (HBA) offloads the iSCSI and TCP/IP encapsulation from the CPU. An HBA also allows the system to boot from the SAN. <p>Note: A software-only initiator should be avoided since it consumes a significant amount of CPU.</p> <ul style="list-style-type: none"> • If iSCSI HBA is used, configure it for maximum queue depth. 	
<p>Tune the cache for the disk subsystem</p> <ul style="list-style-type: none"> • Ensure that Write through cache is enabled • Identify amount of available cache? <p>Caution: Write back or other write caching mechanisms are a potential point of failure which can, in a worst case, corrupt a database.</p>	

Item	Comments
<p>What level of RAID are you using?</p> <ul style="list-style-type: none"> • None • RAID 0: Data Striping • RAID 1: Data Mirroring • RAID 5: Block Striping with distributed parity • RAID 1+0: Mirrored sets in a striped set <p>Note: A RAID 5 or RAID 1+0 system is recommended for any mission critical system. RAID 0 is sufficient for a cache or other systems that requires high speed but for which reliability isn't critical.</p>	
<p>Make sure any disk intensive applications or subsystems have access to a private disk (or, more precisely, a private disk arm). In general, it is often better to have many small disks than a few huge ones.</p> <p>Examples of application/sub-applications that should have private disks are:</p> <ul style="list-style-type: none"> • Operating System • Any major application (Depending on the application this might be separated onto multiple disks) • Temp areas / cache areas with high I/O • Database data files • Database Transaction logs • Database TempDb • Database table with high load 	
<p>Make sure the virtual machine's swap files are located on a high speed storage system.</p> <p>Note: Memory swapping should be avoided, whenever possible; however, when required this step minimizes the performance hit.</p>	

Item	Comments
<p>For each of these application/sub-applications identify the following:</p> <ul style="list-style-type: none"> • How many I/O per second (IOPS) is required? • What is maximum acceptable response times? <p>Ensure that the provided disk system can guarantee these requirements - even in a failover situation.</p>	
<p>Ensure that sufficient storage is allocated – or even better, can be easily added when required.</p> <p>Note: Plan so that you never utilize more than 80% of the available disk space.</p>	
<p>Ensure all devices included in a Hot Migration Pool (VMotion, Live/Quick Migration, XenMotion) with good performance can reach all required disk systems.</p>	

Management Infrastructure

Management of business critical applications has always been important. In a virtualized environment, however, it is even more important since each system can affect the others’ scalability and since an unplanned interruption of a host system now can affect multiple business critical applications.

With good management practices you can minimize these downtimes by finding out when you need more resources and quickly provisioning a new server or making sure the critical application gets all the resources it needs.

The first step is to list the management products you are planning to use to manage your virtual environment. The following checklist only lists a few common examples; specify any other tools where applicable.

Products / Management area	Comments
CA Advanced Systems Management (AMS)	
CA Data Center Automation (DCA) Manager	
VMware Virtual Center	

Products / Management area	Comments
Microsoft Virtual Machine Manager	
<p>Performance Management Tools</p> <p>These are tools that manage performance of guest systems as well as the complete host and, preferably, have the functionality to correlate this information.</p> <p>CA has several performance management tools, including;</p> <ul style="list-style-type: none"> ▪ CA Data Center Automation ▪ Unicenter Network and Systems Management (NSM) ▪ Wily Introscope & CEM 	
<p>Change Control / Change Management Tools</p> <ul style="list-style-type: none"> • CA Unicenter Service Desk • CA Cohesion • Harvest CCM 	
<p>Tools to manage Virtual Machine Sprawl</p> <ul style="list-style-type: none"> • CA CMDB • CA Advanced Systems Management (AMS) • CA Data Center Automation (DCA) Manager • Client Management Solution 	
Other Tools...	

Just as important as which tool you use are the procedures you use to control and manage your environment and how the virtual images are created and used. Make sure you have a well defined process for this, including several distinct stages for Development, Test/Quality Assurance, Gold Images, Archive of older images, Configuration and finally production.

Cost Analysis

Last, but not least, you need to consider the true costs connected to the virtualized environment. A complete analysis requires a close look at the organization's situation and specific requirements; however, the following checklist represents a good starting point and highlights several key points that you need to consider:

Cost Area	Comments
Server Hardware	
Storage Solution / Infrastructure	
Education and potential new hire	
Redeployment of existing and working solutions	
New Management Tools (license, hardware, services)	
Virtualization Engine and Tools	
License Costs <ul style="list-style-type: none"> • Need for additional licenses for Operating System and back office products (such as security and other management tools/agents). Virtualization often leads to the deployment of more, but smaller logical servers <ul style="list-style-type: none"> • Application licenses that are bound to the number of physical CPUs on the hosting server. • Licenses that are limited to a single named physical server. If you have any of these, determine if you need additional licenses to handle load migration between host systems 	

Glossary

Advanced Systems Management (CA ASM)

CA ASM is a solution for managing multiple heterogeneous virtual and physical server environments. It provides a common method for managing different types of virtualized environments as well as cluster environments.

ASM also allows the user to either configure the environment manually, either with or without optimization advice from the tool, or to enable it to automatically adapt to business needs in real time, continuously assessing and tuning system resources and services based upon your business policies. This ensures your IT investments are optimized.

Application Virtualization

There are several definitions for "Application Virtualization" some of which are significantly different from each other. Following are two common interpretations:

1. A technique that aims to improve stability and security of the systems by using abstraction layers between the application and the actual hardware. This technique is used in most modern operating system to enable protection of the operating system, as well as other applications, from an application that has been poorly written.
2. Application service virtualization is also referred to as Application Virtualization and refers to running software on a central server rather than distributed to each of the users' computers. Except for a very thin generic client and communication protocols (that often can be a standard webbrowser) no changes are made to the local computer's file system or registry.

Hot Migration

Generic name for technologies that allow virtualized environments to move between host systems while they are still running. This functionality greatly enhances the advantages of virtualization by minimizing (or even eliminating) application downtime when environments are moved to manage load balancing or during hardware maintenance.

VMware ESX servers use VMOTION to accomplish this, XenSource is planning to release XenMotion together with XenEnterprise 4.1 and Microsoft is planning to include Live Migration in a future release of Hyper-V.

Hypervisor

A hypervisor is a basic virtualization component that provides an abstraction layer between the hardware and the "guest" operating systems. A hypervisor has its own kernel and is installed directly on the hardware. It can be considered as a minimalist operating system that controls the communication between the guest OS and the hardware.

A virtualized environment without a true hypervisor (sometimes referred to as Type 2 Hypervisors) needs to have a primary OS in between the hardware and the Virtualization Engine, which can add significant overhead.

Hyper-V

Hyper-V is Microsoft's next server virtualization product. Hyper-V went beta in December 2007 and a feature complete release candidate is currently available for Windows Server 2008 users. It supports up to 4 CPUs per virtual machines, has a true Hypervisor architecture, support for Quick Migration and support for 32 and 64 bit guest systems running Windows or Linux.

A future version is expected to support host systems with up to 64 CPUs and will include features such as Live Migration and hot-add of resources such as storage, networking, memory and CPUs. Hyper-V is expected to be released in August of 2008.

iSCSI

iSCSI is a network protocol standard (transport layer) that allows use of SCSI protocol over TCP/IP networks. Building iSCSI based Storage Area Networks (SANs) using this protocol, together with Gigabit Ethernet, is a significantly cheaper alternative to the traditional Fibre Channel based SANs

Live Migration

Live Migration is Microsoft's solution for Hot Migration (See Hot Migration above).

Operating System (OS) Virtualization

OS virtualization allows isolated partitions, or virtual environments (VE), to be located on the same physical server and operating system. Multiple VEs share a common operating system that they communicate with through an "OS virtualization Layer", This layer is responsible for ensuring security and complete isolation of dedicated resources and data owned by a specific VE.

VEs are sometimes called "virtual private servers (VPS)", "jails," "guests," "zones," "vservers" or "containers", to name a few.

Paravirtualization

Paravirtualization is a technique where the "guest" operating systems utilize software interfaces through the hypervisor that aren't identical to the underlying hardware (though, typically, similar). This can make certain calls more efficient and it can certainly simplify the role of the virtual engine; however, it requires the operating system to be aware of and to use the specific Virtualization Engine interfaces.

Quick Migration

A Microsoft solution that enables migration of guest systems to other Hyper-V servers. This is not the same as Live migration since this solution requires some downtime. This downtime is typically short but it depends on factors such as I/O performance to the SAN and the size and usage of the virtual machines.

SAN (Storage Area Network)

A Storage Area Network (SAN) is an architecture that allows remote storage devices to be attached to servers in such a way that the operating system perceives it to be a locally attached device. This significantly enhances the virtual environments since you can now easily separate the storage of the virtual images from the rest of the hardware and, thereby, simplifies any type of Live Migration. Traditionally, SAN was implemented with Fibre Channel, however lately it has become more and more common with to see cheaper iSCSI solutions being used in this capacity.

Server Core

A Server Core installation of Microsoft Windows Server 2008 is a minimal installation that avoids extra overhead. This greatly limits the roles that can be performed by the server but it also improves security and reduces management costs. The Server Core is trimmed down in a number of ways but the most obvious is that the Graphical User interface is removed. It is, instead, administered from the command prompt or a remote administrative interface.

SMP – Symmetric Multiprocessing

Symmetric Multiprocessing is a computer architecture where two or more CPUs are connected to a common shared memory. Since the processors share memory the operating system can easily balance the workload by moving tasks between them. SMP is currently the most common architecture used by multiprocessor machines. Examples of competing architectures include NUMA (Non-Uniform Memory Access), ASMP (asymmetric multiprocessing) or computer clustered multiprocessing (e.g. Beowulf).

Thin Provisioning (of storage resources)

Thin provisioning is a technique that allows disk space to be allocated to servers or users on a “just enough” as well as a “just in time” basis. This addresses a problem where servers and users often exaggerate their current need to account for the need for additional space in the near future.

VDI – Virtual Desktop Infrastructure

Virtual Desktop Infrastructure (VDI) utilizes virtualization techniques to provide end users with their desktop environments. This technique is similar to server virtualization, however it also presents its own unique set of advantages and challenges.

VI3 – VMware Infrastructure 3

VI3 is VMware’s virtual server infrastructure and it consists of a VMware ESX Server and related distributed services, such as High Availability (HA), Distributed Resource Scheduler (DRS) and Consolidated Backup.

Virtual Appliance

Similar to ordinary software appliances virtual appliances are minimalist packages with everything you need to provide a specific service. A virtual appliance is delivered in the form of a pre-configured virtual machine that includes an OS and application that are optimized. This simplifies the install and can greatly enhance security, since unnecessary services, daemons, protocols, etc. can be disabled.

VHD (Virtual Hard Disk)

A virtual machine encapsulates an entire server or desktop environment in a file. The Microsoft VHD file format specifies a virtual machine hard disk that can reside on a native host file system encapsulated within a single file. Microsoft has made the VHD Image Format Specification available to third parties under a royalty-free license, however they still own the standard and they reserve the right to revise or rewrite it.

Viridian

Viridian is the codename Microsoft has been using for its expected server virtualization product. For additional details see “Hyper-V” above.

Virtual Machine

A virtual machine is a self contained software environment that works on top of a host operating system through a set of well defined interfaces. These interfaces can either be specific to this virtualization (see Paravirtualization) or can use an abstraction layer that catches the calls the virtualized software environment normally would communicate with its operating system and/or underlying hardware infrastructure.

VMFS (Virtual Machine File System)

VMFS is VMware’s cluster file system and it is used only with ESX server. VMFS allows multiple virtual machine disk images to be stored in a way that can be accessed (read/write) by multiple servers simultaneously (currently up to 32). VMFS is required for VMOTION to work.

VMDK (Virtual Machine Disk Format)

A virtual machine encapsulates an entire server or desktop environment in a file. VMware's VMDK specification describes and documents the virtual machine environment and how it is stored. VMware has made the specification for VMDK available to the public; however they retain the right to revise or rewrite it.

VMOTION

VMOTION is VMware's solution for Hot Migration (See Hot Migration above).

Windows Server Virtualization (WSV)

Windows Server Virtualization (WSV) is the name Microsoft previously used for its server virtualization product Hyper-V. For additional details see "Hyper-V" above.

XenMotion

XenMotion is XenSource solution for Hot Migration (See Hot Migration above).