

Implementing and Expanding a Virtualized Environment

"Virtualization has helped our organization lower costs, increase agility, and reduce energy consumption."

Bill Sunderland
Virtualization Program Manager,
Intel IT

Steve Anderson
Systems Engineer,
Intel IT

Executive Overview

In 2005, Intel IT began planning, engineering, and implementing a virtualized business computing production environment as part of our overall data center strategy. We are currently accelerating virtualization adoption for general purpose applications, with a goal of virtualizing 70 to 80 percent of our office and enterprise environments within two years.

We started by standardizing our virtualization infrastructure architecture using proven technology and measuring return on investment (ROI). From the outset, we have taken a cautious approach—conducting a thorough consolidation and ROI analysis; carefully evaluating resource utilization, available platforms, and risk factors; developing a comprehensive architecture design including hosts, storage, backup and restore (BAR), networking, and management; and addressing necessary business process changes. Our initial ramp focused wholly on older servers running non-mission-critical applications.

To date, we have virtualized 10 to 20 percent of the servers in our office and enterprise environments. This effort has yielded benefits that include automated deployment, faster recovery, and an average server consolidation ratio of 10:1 with calculated available capacity of 15:1. This has helped our organization lower costs, increase agility, and reduce energy consumption.

By integrating two-socket servers based on the Intel® Xeon® processor 5500 series, we see opportunities to more than double our consolidation ratios up to 20:1 based on internal benchmark and performance tests.

CONTENTS

Business Challenge	2
Solution.....	2
Examining Return on Investment	3
Projected Savings.....	3
Break-Even Point.....	4
Host Platform Selection	4
Resource Utilization in the Current Environment	4
Virtualization Host Platform.....	5
Virtualization Architecture	6
Virtualization Host.....	7
Network.....	7
Storage Architecture	7
Backup and Restore.....	7
Virtual Machine and Host Management.....	8
Resource Monitoring and Management	8
Physical to Virtual Conversion.....	9
Building the Environment	9
Engineering.....	9
Business Process	10
Key Lessons	10
Virtual Machine Management Servers.....	10
Physical To Virtual Conversions.....	11
Application Response Times.....	11
Virtual Machine Live Migration.....	11
Storage Area Network	11
Conclusion.....	11
Acronyms.....	11

IT@INTEL

The IT@Intel program enables senior staff to share IT-related experiences with industry colleagues. The program creates a variety of content—from white papers and briefs to online videos and social media—by working with our IT subject-matter experts. This content demonstrates industry leadership by sharing our own IT challenges and solutions.

Through social media activities, IT@Intel engages with the IT industry on topics ranging from data centers to server and client systems. The public can discuss hot topics with Intel IT professionals by visiting www.intel.com/IT.

BUSINESS CHALLENGE

Intel IT saw server virtualization and consolidation as a way to realize cost savings in hardware, software, and administration; achieve greater agility; and reduce energy consumption.

Even for a technology company like Intel, implementing server virtualization was initially daunting. To set the stage for success, Intel IT spent time upfront analyzing the opportunities for virtualization, determining how we would measure ROI, developing a virtualization architecture, and implementing a production environment.

Virtualizing our business-computing environment presented significant challenges. The Intel Enterprise and Office worldwide environment includes approximately 6,000 to 7,000 physical business-computing servers primarily running Microsoft Windows* and a wide variety of applications. While we knew that virtualizing and consolidating these physical servers could deliver significant cost savings and other benefits, the sheer size and sophistication of the environment made this a complex undertaking.

From the outset, we decided to standardize our virtualization infrastructure architecture using proven technology. We also wanted to conduct a detailed ROI analysis and consolidation assessment to guide the rollout. We are currently in the process of accelerating our deployment. Our goal is to virtualize 70 to 80 percent of our office and enterprise server applications over the next two years.

SOLUTION

In 2006, we conducted a detailed ROI analysis, examining each aspect of our environment that would be affected by virtualization.

Examining Return on Investment

Our analysis looked at current costs of deploying physical servers versus those in a virtualized environment and assumed we would consolidate multiple physical servers into virtual machines (VMs) on virtualization host servers. To calculate ROI at that time, our finance analysts used an extremely conservative consolidation ratio of 4:1 (although we expected to achieve much higher consolidation ratios in practice), with a target of consolidating approximately 4,800 older business-computing servers.

Our analysis included:

- **Server capital costs.** We assessed the cost of the average physical server in our non-virtualized environment versus those of two-socket and four-socket server platforms that we could use as virtualization hosts.
- **Data center utility costs.** We looked at the cost of electricity for IT equipment power and cooling, and did not include gains from using existing infrastructure to support additional computing capacity.
- **VM hypervisor license and support.** We based our analysis on the cost of a four-year license for the VM hypervisor, including support. We also included the cost of centralized server-management

software licenses provided by the VM hypervisor supplier.

- **LAN.** An infrastructure utilization assessment found that our typical physical servers use three network ports. We calculated the cost of switch capacity, based on the number of switch ports used by each server and the total cost of the switch and the Ethernet cables.
- **Storage area network.** Many of our physical servers have two Fibre Channel (FC) connections to a storage area network (SAN). We found that each host in a virtualized solution would also use only two FC ports, even though it would support multiple virtualized server workloads. We included the cost of FC cables in our calculations.
- **Engineering and support headcount.** We did not include any headcount reductions, as we assumed that the efficiencies would be offset by additional engineering complexity and support needs in the virtualized environment.
- **Automation and productivity gains.** Our ROI calculations did not include productivity gains due to virtualization, although we did expect a range of productivity benefits, such as the ability to provision systems in minutes rather than days or months, perform hardware

maintenance and patch upgrades without down time, automatically load-balance workloads, and recover faster from hardware failures. We also expected that each support person would be able to manage a larger number of VMs due to virtualization and automation.

In addition, we expected that virtualization would result in more efficient resource utilization, such as higher memory and CPU utilization. By virtualizing, we also anticipated being able to consolidate dissimilar applications onto a single physical server but did not factor this into our ROI calculations.

PROJECTED SAVINGS

This 2006 analysis showed that virtualizing our business-computing environment could generate estimated savings of USD 17.6 million to 27.7 million over five years. We conducted another ROI analysis in 2008 after the initial roll-out to determine whether we should refresh existing aging physical servers with new ones or use physical to virtual migration to convert them into VMs. We looked solely at hard costs and found that for 4,800 servers, assuming a 15:1 consolidation ratio, we could achieve approximately USD 23 million in savings over four years by virtualizing. We also identified the specific sources of positive and negative ROI, as shown in Figure 1.

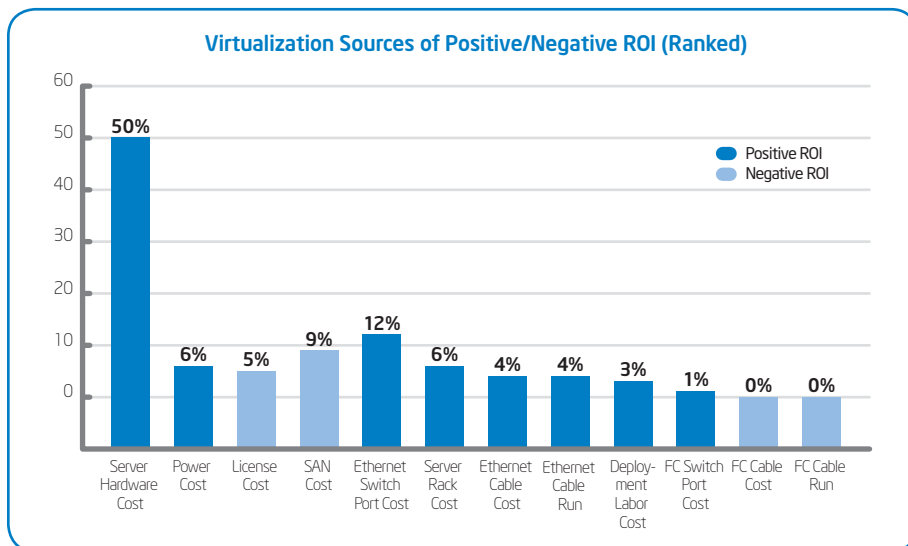


Figure 1. Virtualization sources of positive and negative ROI.

The top savings factor for virtualization is server hardware, followed by Ethernet switch ports and power savings. By identifying our primary negative ROI factors of virtualization license and SAN costs, we were able to evaluate whether we could reduce those costs for greater ROI. We realized that we were using expensive FC hard disk drives for our entire virtualized environment, but 70 percent was pre-production. We decided to use less expensive, lower performance serial ATA (SATA) hard disk drives for our pre-production environments, which enabled a 68 percent cost savings.

BREAK-EVEN POINT

To further analyze virtualization ROI, we also examined and updated a previous ROI study looking at consolidation ratios that had been performed by Intel’s manufacturing-computing group. We updated the study to include two-socket and four-socket servers based on the latest Intel® Xeon® processors and found that a migration from our current physical server environment to a virtualized

environment would result in positive ROI, even at very low consolidation levels. We found that the break-even point using two-socket servers as virtualization hosts would be at approximately 2.25 VMs per host, as shown in Figure 2. With four-socket hosts, the break-even point occurred at 3.25 VMs per host.

Host Platform Selection

The first critical step in defining our architecture design was selecting the right host platform. To make this decision, we analyzed resource utilization in our current non-virtualized environment, considered different consolidation scenarios, and examined our platform options.

RESOURCE UTILIZATION IN THE CURRENT ENVIRONMENT

In late 2005 and early 2006, we collected extensive server resource utilization data from our current non-virtualized environment. We used performance-monitoring tools to collect real-time utilization data for all IT enterprise

business-computing servers at five major data centers. We monitored CPU, memory, network I/O, and disk I/O, and collected utilization snapshots for each server every 10 minutes over 60 days. We then compiled and analyzed the data to determine the maximum, minimum, and average utilization for each resource. This enabled us to estimate the VM host resources that we would need when we migrated each physical server into a virtualized environment. We found that our current environment was woefully underutilized, with an average CPU utilization of 12 percent, and 75 percent of surveyed systems consumed less than 1 GB of memory—even though most had 2 GB or more of memory installed.

We then conducted an updated CPU/memory resource utilization study in 2009. In order to identify average processor utilization and maximum memory consumption, we collected data every 15 minutes for 12 weeks from approximately 10,000 Intel processor-based physical

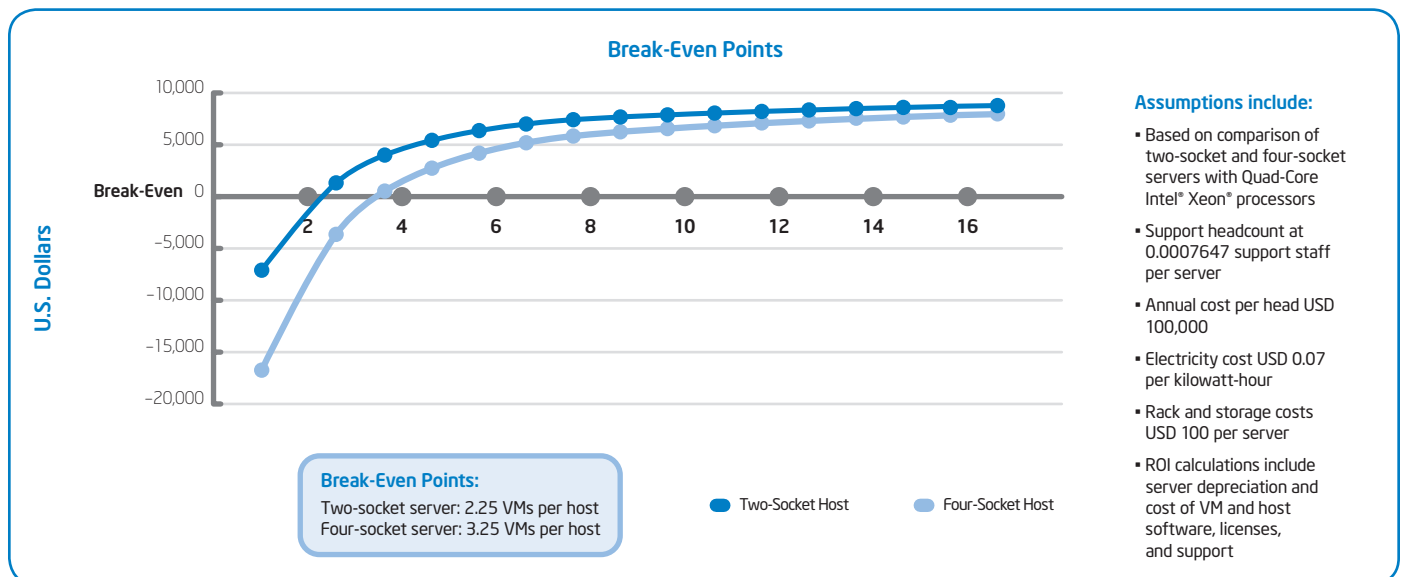


Figure 2. Virtualization ROI break-even point.

Figure 3a. Average memory consumed by quintile.

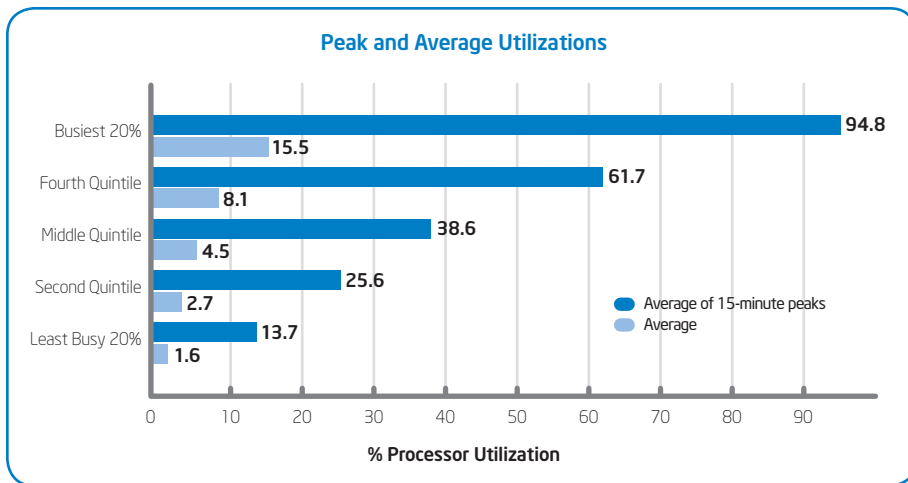
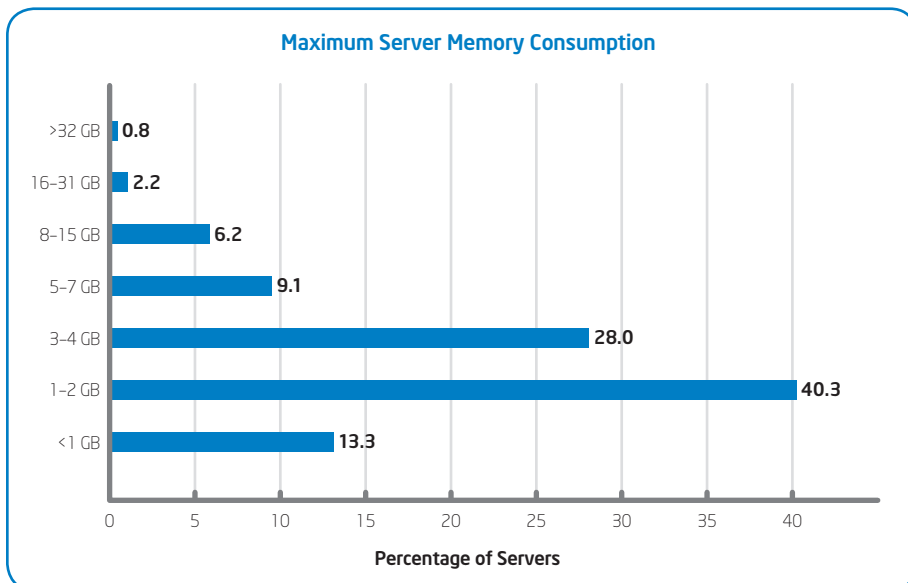


Figure 3b. Maximum server memory consumption (the difference between physical memory installed and the lowest memory available data point).



servers running Microsoft Windows*. We found that the median server experienced a peak CPU utilization of approximately 40 percent, and the long-term CPU utilization average for the majority of servers was below 10 percent. See Figure 3(a). Furthermore, approximately 50 percent of the servers surveyed consumed no more than 2 GB of RAM, and more than 25 percent ranged between 2 and 4 GB of RAM. See Figure 3(b). Understanding these points, we consolidated these underutilized physical servers onto a virtualized platform, resulting in lower overall total cost of ownership (TCO).

VIRTUALIZATION HOST PLATFORM

When selecting a host server platform, we first looked at the possibility of reusing existing servers and increasing their resource utilization through virtualization. We quickly concluded that reuse was not a good idea, as new machines offered much greater performance that would help us reach higher consolidation ratios and reduce overall costs. Selecting new systems enabled us to standardize on a single virtualization-hosting platform to simplify engineering, deployment, and support. It also allowed for live migration, which requires that hosts meet CPU compatibility

requirements, particularly for the support of 64-bit VMs.

We then needed to choose from two-socket, four-socket, eight-socket, and blade server platforms. To select the right platform, we had to balance cost, performance, and risk. The more VMs a host supports, the greater the risk if the host experiences failure. At the time, because virtualization technology was still relatively new, we decided that until the technology had matured, we wanted to limit consolidation ratios to a maximum of 20:1 to reduce risk.

While eight-socket and blade servers could accommodate larger numbers of VMs within a single chassis, we were concerned about the risk and the possible introduction of other technical issues such as network bandwidth limitations. This focused our decision on two-socket and four-socket servers, both of which accommodated quad-core processors at the time of the analysis. At this time, four-socket servers offered roughly twice the performance of two-socket servers at roughly twice the cost. Intel IT testing has found that four-socket servers offer other advantages such as more predictable scalability and are better suited for memory-intensive applications.¹ However, two-socket servers were less expensive than the larger servers and better fit our 20:1 maximum consolidation parameter.

As a result, we selected two-socket servers based on quad-core processors for our initial

virtualization efforts. With more than 1,600 VMs deployed, we have found our average consolidation ratio to be 10:1 with available capacity to achieve 15:1. In the process, we discovered that we needed to upgrade all of our hosts' memory to 32 GB, as memory bottleneck became the limiting factor.

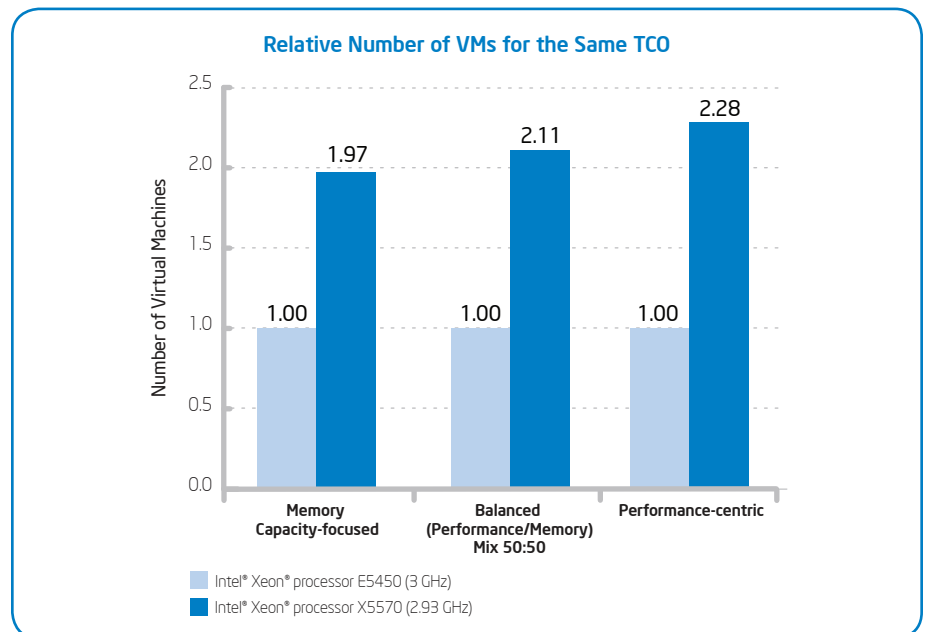
Now that virtualization technology has matured and Intel IT has had four years of hands-on experience, we believe that consolidation ratios of 20:1 to 30:1 offer acceptable risk. We see four-socket servers as a viable option for accommodating applications that have high resource needs; however, we are still primarily focused on virtualizing the numerous applications with low-end resource needs. Given the substantial performance advancements of the Intel Xeon processor 5500 series, we expect that the two-socket platform is capable of delivering our target consolidation ratios. Based on internal benchmark tests

using the vConsolidate* benchmark suite, we anticipate doubling our consolidation ratio while maintaining the same TCO using virtualization hosts built on these processors (see Figure 4).²

Virtualization Architecture

Designing our virtualization architecture was a complex undertaking. We started by assessing the infrastructure technologies currently in use within Intel IT—including storage and SAN, network, BAR, monitoring and alerting, capacity management, remote management, automation, and operating systems—in order to map them to a virtualized architecture. We then outlined our ideal virtualization infrastructure as if we were building a completely new environment. This allowed us to design an architecture that merged our current and ideal architectures into a realistic solution, with some compromises due to cost and required implementation effort.

Figure 4. Relative number of virtual machines for the same TCO for different virtualization scenarios.



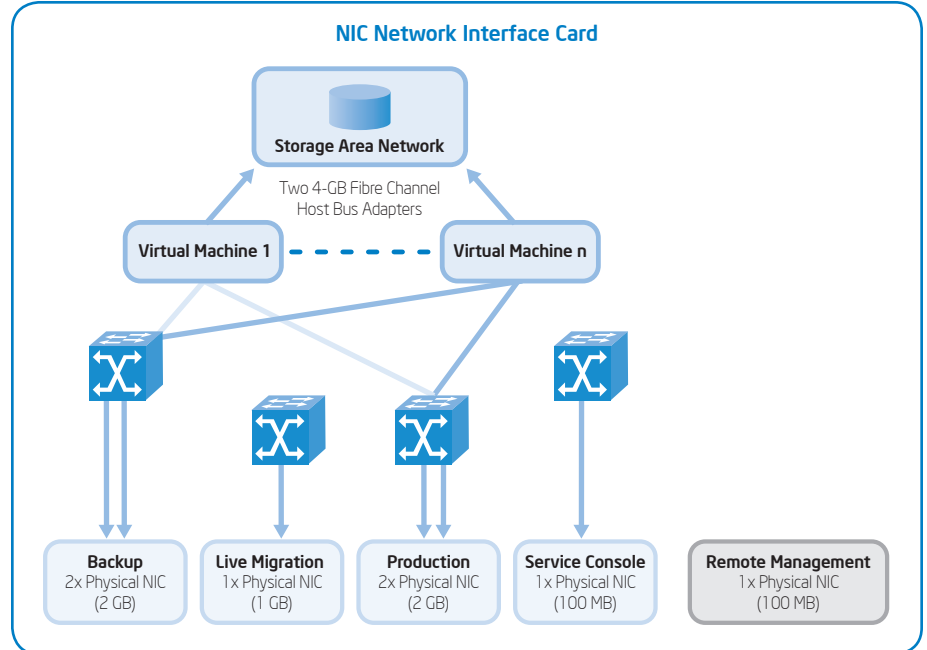


Figure 5. Virtualization host architecture showing network and storage.

VIRTUALIZATION HOST

Our virtualization host architecture design is based on a two-socket virtualization host platform and includes storage and networking, as shown in Figure 5.

NETWORK

Our network architecture consists of seven network ports for each host:

- **One 1-GB port dedicated to the VM live migration software.** We considered sharing ports for BAR and live migration, but we were concerned that when backups were intensively using bandwidth, live migrations could fail.
- **One 100-MB port for the service console.** Performance testing and subsequent production use have shown that 100 MB is adequate for console network traffic in normal use. However, we found that a 1-GB port is advisable when conducting physical to virtual (P2V) conversions.
- **One 100-MB port for remote management.** This was adequate in our physical server environment, and we have found that it suffices in the virtualized environment, as well.

- **Two 1-GB ports bonded together for production network traffic.** We needed high bandwidth because virtualization consolidates multiple server workloads onto a single host. Performance testing showed that 2 GB is adequate to support 16:1 consolidation ratios.
- **Two 1-GB ports bonded together for backup traffic.** Intel's BAR implementation requires all combined SAN data from the VMs on the virtualization host to flow through the host's BAR network. This demands 2 GB of bandwidth to meet our BAR service level agreements (SLAs).

STORAGE ARCHITECTURE

VM live migration calls for a SAN approach. The data used by each VM's applications must reside on the SAN, accessible by all hosts, in order to allow live migration to occur. We store most data and software on the SAN, with the exception of virtualization guest VM operating systems, which reside on the server's local storage.

Analysis of performance data has shown that disk I/O has never been an issue in any of our environments, even when running numerous VMs with disk-intensive applications on a single host.

BACKUP AND RESTORE

Backups are performed from within the individual VMs, with the frequency and strategy determined by the owner of the VM and by our operations BAR team. We found that it was not cost-effective to replace our existing BAR physical infrastructure. As a result, we had to determine the maximum amount of data we could back up and restore, given that the combined data from the VMs on each host has to pass from the SAN through the host's network before being offloaded to tape. We found that the maximum was 4 TB per host before we encountered issues with BAR SLAs. As we increase the number of VMs per host, we anticipate that we will begin to run into this limitation. To address this, we have launched a proof of concept (PoC) effort to identify a more efficient method for BAR, such as SAN to SAN to tape.

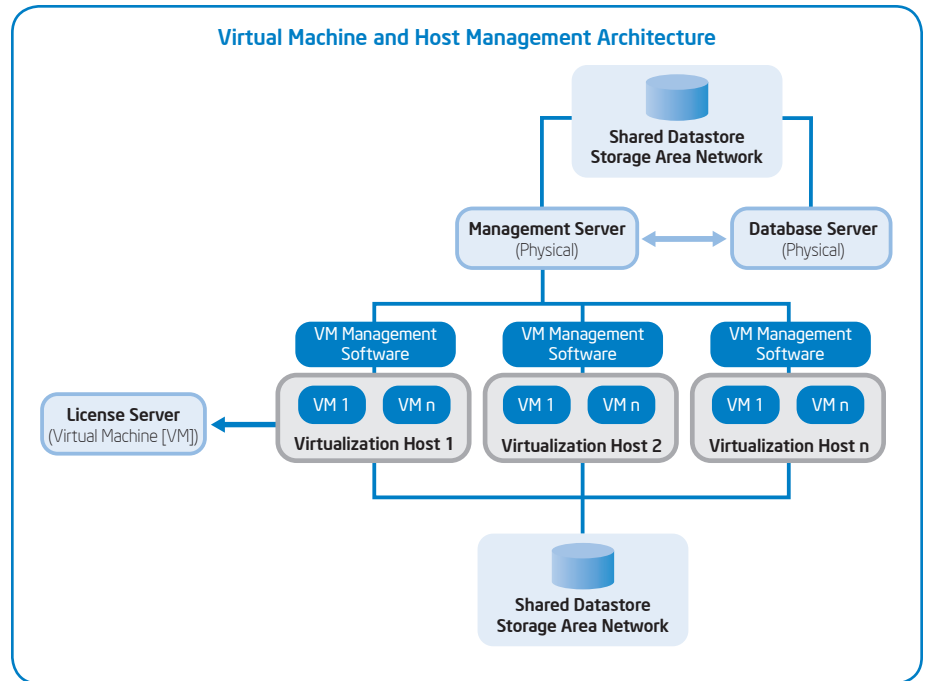


Figure 6. Virtual machine and host management architecture.

VIRTUAL MACHINE AND HOST MANAGEMENT

Our centralized VM and host server-management software provided by our VM hypervisor vendor performs a wide range of functions, including VM and host configuration, provisioning, monitoring, migration, and resource management. Therefore, the design and implementation of our VM and host-management architecture affected the efficiency of our entire virtualized environment.

Our design is shown in Figure 6. The design balances cost, availability, scalability, and location. Because we determined that we could afford up to 24 hours of down time with minimal impact, we were able to reduce costs by eliminating the need for clustering.

We run our centralized VM and host management software on a physical non-virtualized server, and we run the virtualization host software license server separately as a VM within our virtualized environment. This means that our license server keeps working even if

the management software is unavailable, allowing us to continue deploying VMs and hosts.

We also offloaded the database used by the management software to another physical server. Performance tests showed that this helped both the database and the management software servers run more efficiently, so that the system performed better during periods of heavy concurrent administrative use, such as when conducting multiple simultaneous deployments.

We originally expected to deploy only one management server within each major region, as each server was capable of managing 200 hosts and a total of 2,000 VMs. Because of the anticipated growth of our data centers and our commitment to accelerating deployment of virtualization overall, we decided to deploy one server per major data center. This also eliminated the concern about the impact of wide area network bandwidth between data centers.

RESOURCE MONITORING AND MANAGEMENT

We use a range of other tools to monitor and manage resources within our environment. We monitor host resource utilization to determine whether it remains within acceptable limits:

- Daily average of less than 70 percent host CPU utilization
- Less than 5 percent swap utilization
- Less than 65 percent production network virtual switch utilization

While we consider a brief peak to be acceptable, sustained excessive resource use is not. In practice, dynamic resource balancing limits the extent to which this occurs, though it considers only CPU and memory utilization—not disk and network I/O—when deciding how to balance VMs among hosts.

We use internally developed capacity, performance, and management agents to track CPU, memory, disk I/O, and network I/O performance for each VM. We created asset management scripts that run on the




Figure 7. Physical to virtual conversion.

management server every five minutes and update our internal asset management database with information about newly provisioned VMs, those migrated to new hosts, or those removed.

PHYSICAL TO VIRTUAL CONVERSION

We defined a clear process and architecture for physical to virtual (P2V) conversion that migrates each physical server into the virtualized environment, as shown in Figure 7.

We decided to use a manual conversion process to accommodate the different physical server configurations in our data centers. We refined this process as we completed more conversions, and we tested and included batch conversion capabilities. These batch conversions allowed us to convert a larger number of systems but required additional planning and customer communication.

We found that once we had successfully completed a conversion, we needed to allocate resources to validate the results of the conversion activity as well as application functionality and performance.

To reduce the touch time on each converted system, we automated the removal of hardware-specific software support using Windows scripting technology. This has improved both VM performance and stability, as the system is no longer attempting to manage hardware not present in the system.

We also found that we could successfully conduct multiple simultaneous conversions into virtual environments when multiple hosts were present. However, only one conversion should be attempted at any one time per host. Attempting multiple simultaneous conversions onto the same host significantly lengthened conversion time and could result in a stalled or failed conversion due to saturation of the service console port with network traffic.

Building the Environment

Once we had designed the architecture of our virtualized environment, we scoped and executed the engineering and business process activities needed to implement it.

ENGINEERING

Our engineering activities involved building and testing each part of our virtualized environment. This included running PoC tests to validate our approach. These PoC efforts tested virtual host architecture against various workloads to determine individual VM responsiveness. These results helped establish maximum consolidation ratios and tested specific use cases for Intel business groups.

Virtualization Host and Management Server.

We engineered, tested, and delivered build documentation for our selected two-socket server platform based on Intel® Xeon® processors, as well as the servers running our VM and host-management software and associated database. Through performance testing, we determined that these servers required 4 GB of memory (12 GB for database servers), dual 1-GB Ethernet ports, and two 72-GB hard disk drives.

Dynamic Resource Balancing. We engineered and documented the dynamic resource-balancing capability, and tested basic functionality to help ensure that this feature worked as desired. This included establishing default settings for normal CPU priority with unlimited ability to expand resources, if necessary; restricting the use of settings that allow VMs to reserve production environment resources; and closely monitoring clusters during early implementation to determine optimal settings. We initially used a partially automated mode in which the dynamic resource-balancing software places the VM on a host that has the required resources, but does not subsequently automatically move that VM between machines. Once we better understood the VMs' characteristics, we used a fully automated approach. To allow maximum flexibility, we generally avoided specifying rules that defined which VMs should or should not share a host.

Time Synchronization. Because VMs share time on a physical host, a VM cannot exactly duplicate the timing of a physical machine. The VM's clock could be out of synch with the host clock, causing problems for the application running on the VM. We solved this problem by running Network Time Protocol on the virtualization software host console to help ensure that the host is synchronized with the network. To synchronize the VM with the host, for each VM, we disabled the operating system (OS) time server synchronization service and used the VM supplier's software.

Provisioning. We determined which guest OSs the VMs would support and created automated provisioning solutions. We automated installations of 32-bit and 64-bit versions of Microsoft Windows for VMs and the hypervisor supplier's Linux*-based host OS. We accomplished this by creating images and scripts, and by using a third-party deployment tool.

Inventory. Intel has security and business-related requirements for effective asset management, which can be challenging in virtualized environments. Dynamic resource balancing automatically moves VMs between hosts, making it difficult to track the location of each VM without manually extracting it from the VM management software. To meet our corporate requirements, we wrote an SQL script to automatically track creation, deletion, and physical movement of VMs.

Patching. We engineered, tested, and documented patching processes. We use fully automated patching for VMs running on Microsoft Windows, and for our management server and associated database. Currently, we perform manual patching of hosts but are in the process of engineering an automated, down-the-wire solution.

BUSINESS PROCESS

Implementing a virtualized environment involves significant changes to business processes as well as technology. For Intel IT, this included providing training for all personnel supporting the virtualized environments and ensuring that we had all of the security procedures in place to mitigate risk, including controls to enforce network security.

Many organizations fear that virtualization will result in an unexpected proliferation of VMs, often known as VM sprawl. To avoid this, we defined a process for deploying VMs. A hosting services team within our operations group is responsible for handling all service requests, including landing applications and determining whether a new VM or a new physical server is needed. This centralized model also helps ensure that we make optimum use of resources. Another key to controlling VM sprawl is that we charge each business group for each of their VM deployments on a monthly basis.

Key Lessons

We learned many valuable lessons while architecting and deploying our virtualized environment.

VIRTUAL MACHINE MANAGEMENT SERVERS

Because we rely heavily on our management servers for routine duties, dynamic resource balancing, asset management, and performance trending, we initially considered a high-availability clustered server design. However, we found that the complexity of this solution caused problems. Since down time of up to 24 hours had minimal impact on our environment, we decided not to use clustering, which also reduced costs.

We originally planned to have our VM software license server run on the physical management server. However, this meant that when the management software was down, we could not add VMs or hosts to our environment, which delayed deployments. Also, VM licenses expired if the management server was down for too long. Because of these potential issues, we decided to run our license server as a VM within our virtualized environment.

Our original design called for the database used by the management software to run on the same server as the management software. However, we found that this caused the system to slow to a crawl when we were deploying multiple VMs concurrently because the database and management software were contending for the same resources. We solved this by offloading the database to its own physical server.

PHYSICAL TO VIRTUAL CONVERSIONS

We discovered that customized OSs can make it difficult to migrate legacy systems into a virtualized environment because they affect the behavior of P2V tools. We also found that P2V conversion takes up to four times longer when attempted on a legacy system without optimal network settings. We alleviate this problem by making sure that all NICs in the systems to be converted are set to full duplex, regardless of their supported network speed.

APPLICATION RESPONSE TIMES

Applications that require real-time response or are sensitive to delays may not perform well in a virtualized environment because they have to wait for the resources they need to execute. We found that it is also a good idea to turn off graphical screen savers because they consume unnecessary resources.

VIRTUAL MACHINE LIVE MIGRATION

Our live migration software requires that all VMs in a cluster use the same subnet. We can use multiple network switches only if they enable a rapid spanning tree. We also avoid dedicating physical hardware to VMs, because live migrations do not work

properly with these physical constraints. Live migration and dynamic resource balancing require network bandwidth of 1 GB, as well as hosts with identical or compatible CPU architectures, particularly for support of 64-bit VMs.

STORAGE AREA NETWORK

To enable live migration, all hosts in a cluster need to have visibility into all shared disk units (logical unit numbers). Host assignment to SAN frame ports is on a round-robin basis to avoid starving hosts of data access.

CONCLUSION

Based on the extensive analysis, planning, and testing that we conducted in our initial virtualization roll-out, Intel IT is now accelerating our deployment, with the ultimate goal of virtualizing 70 to 80 percent of our business-computing environment. Virtualization has delivered significant benefits to our organization, including faster recovery, automated deployment, and significant cost savings. We expect to achieve increasing gains as our environment grows and matures.

ACRONYMS

BAR	backup and restore
FC	Fibre Channel
LUN	logical unit number
NIC	network interface card
NTP	Network Time Protocol
PoC	proof of concept
P2V	physical to virtual
ROI	return on investment
SAN	storage area network
SLA	service level agreement
TCO	total cost of ownership
VM	virtual machine

For more straight talk on current topics from Intel’s IT leaders, visit www.intel.com/it

¹ "Comparing Two- and Four-Socket Platforms for Server Virtualization," Intel IT, March 2008.

² "Virtualization with the Intel® Xeon® Processor 5500 Series: A Proof of Concept," Intel IT, June 2009.

This paper is for informational purposes only.

THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE.


Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, Intel logo, and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2010 Intel Corporation. All rights reserved.

Printed in USA
0110/KAR/LAI/PDF

 Please Recycle
323153-001US

