



Virtualization has become a hot topic. Cloud computing is the latest and most prominent application of this time-honored idea, which is almost as old as the computing field itself. The term "cloud" seems to have originated with someone's drawing of the Internet as a puffy cloud hiding many servers and connections. A user can receive a service from the cloud without ever knowing which machine (or machines) rendered the service, where it was located, or how many redundant copies of its data there are. The cloud realizes the old dream of a computer utility, first articulated at MIT in the early 1960s.

One of the big concerns about the cloud is that it may assign many computational processes to one machine, thereby making that machine a bottleneck and giving poor response time. Faouzi Kamoun addresses this concern head on, and assures us that in most cases the virtualization used in the cloud and elsewhere improves performance. He also addresses a misconception made prominent in a Dilbert cartoon, when the boss said he wanted to virtualize the servers so as to save electricity.

*Peter Denning
Editor*

Virtualizing the Datacenter Without Compromising Server Performance

By Faouzi Kamoun

College of Information Technology, University of Dubai
P.O. Box 14143, Dubai, UAE
fkamoun@ud.ac.ae

1. Introduction

Physical servers' processing and computing powers have been increasing for the past years, giving rise to a new breed of multi-core and multi-bit processors, with unprecedented amounts of memory and disk space. At the same time, organizations are faced with the conflicting goals of scaling up their IT infrastructure to accommodate more applications, workloads and users, while consolidating server assets for better efficiency. Further, in most enterprise datacenters, the underutilized resources of physical servers are becoming a liability because of inefficient power consumption and space utilization, and excessive Total Cost of Ownership (TCO). For instance, it is estimated that on average, 90 percent of Windows-based production servers run below 10 percent average utilization [1, page 4]. In a recent research report [2], IDC estimated that server overcapacity is costing IT organizations over \$140 billion.

Virtual Machine (VM) server technology provides a partitioning technique to run multiple and *isolated* virtual servers on a single physical device, thus optimizing hardware usage. A salient feature of virtualization stems from the fact that each virtual server runs independently under its own operating system that is completely separate from the server's primary (host) operating system. This way, each VM can be powered up or down, halted, and resumed independently from other VMs. The isolation feature enables each virtual server application to run independently of other applications, running on different virtual servers [3]. This is in contrast to having multiple applications sharing a single host operating system. The isolation among the VMs is further reinforced with the recent advent of multi-core processors. In fact, by assigning each core one or more VMs, physical separation of multiple VMs become possible. Each virtual server, running a guest OS, is presented with its virtual hardware, which includes virtual hard drive, Network Interface Card (NIC), video card, peripherals ports, controllers, disk drives, CPU,

and memory, among others [1, page 9]. The VMs and the shared resources are managed by a thin software layer, called the virtual machine monitor (VMM) or hypervisor, running on the host server. Figure 1 depicts a simplified diagram illustrating the basic concept of virtualization.




Figure 1. Illustrating the basic concept of server virtualization.

The concept of virtualization is not new and it has been used successfully by IBM within its mainframe environments since the mid 1960s. The technology however has gained renewed attention lately as recent technological advances in multi-core and virtualization-enabled processors, hardware-based I/O architectures, advanced memory access mechanisms, and VM management tools are making virtualization possible on industry-standard servers. Virtualization technology should not also be confused with cloud computing. In fact, cloud computing is a computing paradigm; an operational model which enables dynamically scalable and shared resources (like processors and storage) to be provided, on demand, as a service over the Internet. Although virtualization is not required for the delivery of cloud computing services, it plays a major role in enabling practical, agile, scalable, and low-cost cloud computing infrastructures either inside the data center (internal/private cloud) or outside the data center (external cloud). When coupled with virtualization, the cloud computing model enables physical resources to be virtualized and shared, thus enabling higher utilization rates, while reducing investment in dedicated hardware and associated space occupancy and power costs. This is another reason why virtualization is gaining momentum.

While there are compelling technical and economic reasons for IT organizations to virtualize their datacenters, there are additional management and performance issues that need to be addressed before organizations can take full advantage of latest server virtualization technologies. The main goal of this paper is three-fold: (1) shed light on the virtues and performance-related issues of server virtualization, (2) discuss some latest design solutions and

best industry practices to tackle these performance issues, and (3) provide insights into the future perspectives of server virtualization.

2. The good Things about Server Virtualization

Organizations can reap many benefits by adopting virtual server technologies. We quickly review below some potential advantages that server virtualization has to offer.

2.1. Enhanced Hardware Utilization

A main benefit of virtual server technology is that it consolidates many underutilized servers into a fewer consolidated resources; thereby enabling higher hardware utilization. For instance, Bowdoin College spent about \$200,000 on its virtualization project by consolidating around 101 physical servers into 46 servers. This move also saved the College from having to buy about 60 servers and from having to double the size of the school's 500-square-foot datacenter. In another case, Web-hosting Company MaximumASP is expecting savings up to \$350,000 in hardware costs through the (8:1) consolidation of its 200 underutilized servers into 25 physical servers.

2.2. More Agile Provisioning and Deployment

Since in a virtual environment, virtual hard drives are represented by a set of encapsulated files which reside on the host machine, each of these files can be readily cloned and reused to deploy an additional virtual server. This feature expedites the provisioning of a new virtual server on an existing physical machine, with no additional hardware, software or reconfiguration requirements. Agile provisioning through virtualization also makes it easy to streamline software testing, training and development activities, across multiple environments, on a single physical server. This also enables developers to quickly pull the required infrastructure for testing and development with minimum usage of hardware resources. For example, Surgient, of Austin, Texas, a player in self-service virtualization automation and lab management, achieved 95% savings on software trials by using virtualization technology. Under this model, prospective customers can evaluate products through virtual training labs, running on remote virtual machines, without the need for software demo downloads or installations. Additional operational flexibility is also reflected by the ability to dynamically control memory, CPU and storage resources of each VM.

2.3. Lower Total Cost of Ownership

Through consolidation, virtualization can lead to significant CAPEX and OPEX cost savings and therefore lower TCO at the datacenter. This is reflected by (1) deferred purchase of new servers, (2) smaller datacenter footprint, (3) lower maintenance costs, (4) lower power, ventilation, cooling, rack, and cabling requirements, (5) lower disaster recovery costs, and (6) reduced server deployment costs. In particular, given the increasing trend in energy costs and high-density server deployments, a recent Gartner report [4] predicted that during the next five years, most U.S. datacenters will spend as much on energy (power and cooling) as they will on hardware infrastructure. However, a common misconception is that many people believe that by consolidating many physical servers into a single server, virtualization would certainly save energy. Such a belief is generally valid only if we assume that the physical servers are underutilized and thus are consuming energy while being idle most of the time. In this regards, many factors need to be taken into account. First, it is well known that servers consume substantial amount of power even when they are idle. Second, more energy is definitely needed for cooling the single virtualized server because of the additional overhead (processing load) due to higher utilization and higher heat dissipation (fortunately, most high-density modern

servers are built with high energy-efficiency in mind). Third, energy savings due to reduction in idle power consumption should not be overshadowed by the increase in energy consumption due to higher processing load. For these reasons, when servers are underutilized (which is typical in most datacenter environments), their consolidation does achieve energy savings and reduces utility bills. In addition, by exploiting the latest advances in multi-core processing and micro-architectures, the datacenter manager will be able to replace a number of legacy servers with a single host which is equipped with a more efficient multi-core processor. This upgrade will also increase the performance per watt of the system [5], thus leading to lower power and cooling requirements.

2.4. Enhanced Availability and Business Continuity

Since virtual machines are isolated from each others, the crash of a guest operating system has no effect on the host operating system or any other guest operating system. Further, since virtual servers are unaware of the underlying hardware, it becomes easier to transfer a virtual machine from one physical server to another. In addition, by taking a snapshot (virtual image) of a given virtual hard drive (which is merely a file residing on the host server), it becomes easier to perform backup and disaster recovery procedures. For instance, in case of a VM failure, the VM image can be replicated at a disaster recovery (DR) site. Since each VM is hardware independent, there is no need to duplicate the server hardware at the DR site. It is also possible to configure multiple VMs to perform workload re-balancing in order to meet application SLAs [5].

Despite the fact that consolidating multiple VMs on a single physical server tends to lead to a single-point of failure for multiple applications, various high availability (such as clustering) and data replication mechanisms have been devised to circumvent this problem. For instance, to increase hardware availability, it is possible to cluster several virtual machines to a shared Internet SCSI (Small Computer System Interface) disk over a standard network connection. Take for example, the US automotive supplier Shiloh Industries. The company wanted to enhance the availability and business continuity of its datacenter applications through the adoption of multi-site disaster recovery architecture. To achieve this goal, Shiloh Industries opted for virtual servers that are connected to an iSCSI Storage Area Network (SAN). Alternatively, for better performance, the host operating systems can be clustered and if a VM on a particular virtual server fails, the server will migrate the VM to another node on the cluster [6, pages 269-295].

3. The Bad Things about Server Virtualization

Server virtualization brings a number of technical challenges that must be taken into account before embarking on any virtualization project. These are discussed below.

3.1. Performance Degradation

Virtualization introduces additional overhead to system's performance. The main system components which are affected by virtualization are CPU, memory, storage, networking and applications.

3.1.1. Virtualization, CPU Usage, and Network Performance

As virtual machines share the same network interface card, the network bandwidth will be allocated dynamically among the VMs. In case the aggregate demand for network bandwidth exceeds the capacity of the NIC, each VM will only get an equal fraction of the total NIC bandwidth. At the same time, to grant network access to each VM, the host CPU needs to run

an additional code to emulate the NIC. This additional processing time increases the requirement for host CPU resources in a virtual server environment, which might leave the virtual machines short of CPU resources. This constraint puts a limit on the number of VMs that can run on the same physical server. For instance, in [6, page 85], it is shown that when running eight VMs on a single processor server, the VMs become to a large extent short of CPU resources.

Additional strains on CPU resources are also triggered by the additional processing time needed to emulate not only the NIC but also any peripheral device inside the VM. Further, when the guest CPU load on a VM increases because of CPU-intensive tasks, the amount of resources available to emulate the NIC card decreases, which lowers throughput and thus degrades network performance.

3.1.2. Virtualization and Storage Performance

Storage performance in a virtual environment can be hindered if the virtual server does not possess enough physical storage capability or sufficient amount of processing power to emulate the storage controllers. Various types of virtual disk controllers can be configured and used by a virtual machine, including IDE, SCSI and RAID controllers. Further, multiple VMs, running on the same physical storage device, do not evenly share storage throughput [6, page 88]. Instead, the physical storage device provides the VMs with concurrent disk access. When many VMs simultaneously try to write to the storage device, end users might experience unacceptable latency. In this case, the capability of the storage device to withstand a large number of concurrent disk accesses significantly determines storage performance inside a virtual server. It is generally recommended to use SCSI and RAID controllers instead of IDE storage to better handle a large number of concurrent disk access requests.

3.1.3. Virtualization and Memory Performance

Virtualization does not add much overhead to memory performance. What is more important, instead, is to find out the most appropriate amount of memory to be allocated to each virtual machine and to the host operating system.

3.1.4. Virtualization and I/O Bottleneck

In a typical virtual server environment, the hypervisor provides each VM with a virtual NIC (vNIC) instance and implements a virtual switched network (VSN) to enable the vNICs to communicate with the shared conventional Ethernet NIC [7]. For these reasons, virtualization has generally been tagged as an inappropriate candidate for I/O hungry and delay-sensitive applications because of the extra hypervisor overhead layer.

3.1.5. Virtualization, Chatty Applications and Latency

Many server virtualization initiatives involve consolidating multiple datacenters or moving application servers across the Wide Area Network (WAN). In this scenario, even if bandwidth is abundant, network latency due to propagation delay can bring performance to an unacceptable level. This is particularly true for chatty protocols such as HTTP, CIFS, MAPI, TDS, and NFS, where, for a given session, the number of roundtrips that packets need to perform across the WAN can increase dramatically [8]. Several techniques can be used to minimize the number of roundtrips for chatty applications across the WAN. These include TCP transport layer optimization, and layer 7 application layer optimization. Most of these techniques are still however vendor-specific and proprietary by nature.

3.2. Scalability Constraints

Memory, CPU, storage and workload constraints limit the number of virtual machines that can run on the same virtual server, while delivering acceptable application performance. As a result before organizations move a physical server into a virtual environment, they must assess their computing, storage and performance requirements, perform thorough capacity planning, and choose the right hardware and networking configuration in order to ensure that the performance of the application running inside the virtual machine is not compromised.

3.3. Tracking Dynamic Virtual Machines

Since multiple virtual servers can co-exist within the same physical server, asset management in a virtual environment can be a real challenge. Further, it is a common maintenance practice to move VMs from one physical server to another. Consequently, as the server infrastructure becomes both virtual and dynamic, the task of managing and controlling changes, as well as re-provisioning applications and network services associated with moving VMs, becomes a real challenge [9].

As network management philosophy is shifting from infrastructure management to service management, it becomes even harder in a virtual environment to make a correlation between a hardware failure and the affected applications. For instance, the failure of a link connecting a physical server to a backbone switch will affect all the VMs residing on that physical device, as well as the applications that are running on these virtual machines. Real-time tracking of these dynamic relationships is not a trivial task.

3.4. Potential Security Vulnerabilities

The consolidation of multiple servers inside the same physical device introduces new security vulnerabilities. This is due to the fact that if a hacker can compromise the security of the hypervisor, he might get access to all virtual machines running on the physical server. Further a malicious code infecting the host OS of a physical sever can potentially infiltrate all the applications running on the VMs.

A recent Gartner research report [10] highlighted that security tools for virtual server environments are still immature and that “many organizations mistakenly assume that their approach for securing virtual machines (VMs) will be the same as securing any OS and thus plan to apply their existing configuration guidelines, standards and tools. While this is a start, simply applying the technologies and best practices for securing physical servers won’t provide sufficient protections for VMs.” For instance, most of the security tools and policies have been put in place within the context of a physical word, characterized by fixed servers that are identified by unique physical attributes. Unfortunately, a virtual architecture is hidden and dynamic by nature. VMs can be migrated easily from one host to another and they can be cloned easily. For the above reasons, it is important to design the right security perimeters and policies, intrusion prevention systems, access controls, and best VM security practices, while ensuring that the host OS in the hypervisor is properly configured hardened and patched for known vulnerabilities.

4. Virtualizing While Safeguarding Performance

Major players in the virtualization industry are pursuing different approaches to remedy the inherent negative impact of virtualization on server performance. Accordingly, hardware, software and firmware solutions have been proposed. These solutions are expected to evolve

with a rapid pace in the years to come. The main approaches to circumvent key performance issues and enable near-native performance of virtualized servers are outlined below.

4.1. Selective Virtualization

The current state of virtualization technology reflects ongoing endeavors among the major virtualization players to resolve the prevailing I/O bottlenecks, which are introduced by the hypervisor layer. Until these issues are fully resolved, it remains safer to keep mission-critical, transactional applications that are I/O and latency-sensitive (such as ERP systems), away from any virtualization initiative. It is further recommended that applications which exhibit peak utilization around the same period should not be virtualized inside the same physical server.

4.2. Direct Assignment of Physical NICs to VMs

One approach to address the I/O bottleneck issue is to dedicate separate physical NICs to VMs. As shown in figure 2, each VM (equipped with a NIC driver) is allowed to exchange data with a dedicated physical NIC (typically a 1GbE interface). Note that the VMM is now excluded from the I/O data path. This interaction is made possible via a hardware-based Direct Memory Access (DMA) Remapping function. The DMA module maps system memory access to the target VM. While the VMM (hypervisor) is bypassed as far as the data path is concerned, it still needs to control the data flow to ensure complete isolation of the VMs' DMA requests [7].




Figure 2. Dedicating separate NICs to VMs.

A key limitation of the direct NIC assignment approach is the additional cost in terms of multiple NICs and more cabling and the lack of flexibility in supporting advanced virtualization features such as seamless migration of VMs from one physical server to another. The DMA layer also introduces additional latency in I/O data path. When an additional NIC card is installed as a backup, this approach can provide good reliability in case of a NIC failure.

4.3. Firmware-based I/O Virtualization (IOV)

A firmware-based IOV approach provides management tools that map the links between VMs and the shared NIC port. This enables IT administrators to create virtual I/O channels that can be used by individual guest VMs. A combination of microprocessors and firmware provides the basic building block to isolate I/O channels to multiple VMs. This I/O virtualization approach is

criticized for its inability to truly separate the I/O channels. As a result, reset or re-initialization of an individual firmware I/O channel will impact all remaining channels [11].

4.4. Hardware-based I/O Virtualization (IOV)

Recent advances in multi-channel, hardware-based I/O architectures are enabling true hardware-based virtualization of the I/O subsystem. These advances promise to address most of the I/O virtualization bottleneck issues. As shown in figure 3, instead of having a single NIC card that is being shared among contending guest OS's through the hypervisor, IOV assigns each VM a truly independent I/O channel. This channel is physically implemented as a separate hardware path, which is built in the silicon, inside the NIC's core structure [11].




Figure 3. Hardware-based I/O Virtualization (IOV) concept (adapted from [11]).

Combined with large 10 GbE pipes, the IOV approach has the potential to minimize CPU and hypervisor overheads; thus enabling virtual servers to perform fast I/O functions and support I/O intensive applications. Hardware-based IOV also promises to provide better guaranteed bandwidth and QoS for the virtualized applications through the hardware isolation of data paths inside the virtual server. It is also possible to borrow bandwidth from a given I/O channel and route it to specific applications, when needed.

Further, since each I/O channel is hardware independent (with its own Tx/Rx data paths, DMA engines and interrupts) it can be individually reset or reinitialized without affecting the remaining I/O channels. A common criticism to hardware-based I/O virtualization solutions is that they tend to lock customers to vendor-specific network interface cards.

Recently, the I/O virtualization working group of the PCI-SIG standard organization introduced I/O Virtualization (IOV) specifications, which can be used in conjunction with system virtualization technologies, to enable multiple OS's running on the same physical server to natively share PCIe devices [12]. In particular, the Single Root IOV (SR-IOV) Virtualization and Sharing 1.0 specification enables multiple VMs in a single Root Complex (i.e. host CPU chip set) to share PCIe IOV endpoints without compromising performance.

4.5. Consolidating Computing Resources in a Shared Resource Pool

Another approach to optimize the usage of computing resources and enhance performance is to allocate virtual machines to a resource pool, rather than to a dedicated physical server. As illustrated in figure 4, a distributed architecture allows multiple physical servers to be consolidated into a single resource pool

This pool offers processor, memory, disk, and networking resources to multiple VMs. A Distributed Resource Scheduler (DSR) is used to dynamically balance VM workloads across the resource pool, requesting additional resources from the pool during heavy load conditions or upon request. The DSR can be combined with virtual management software to dynamically migrate VMs from one host to another, thus enabling applications to meet their target service levels [13]. This way, it becomes also possible to free-up resources and consolidate light workload into a fewer number of physical servers.

Though the above approach has the potential to dynamically adapt the usage of computing resources to changing workload conditions, it is criticized for the additional complexity and poor visibility in managing performance issues. For instance, the approach makes it harder for administrators to track where a particular application is running and which resources this application is using. This also adds another layer of complexity in asset management and fault correlation. Advanced VM management tools are being developed to address this concern.




Figure 4. Multiple VMs sharing a common resource pool.

4.6. Management Tools for Virtual Environments

While virtualization technology has been evolving at a rapid pace for the past few years, the development of effective VM management tools is still lagging behind [14]. These tools are however most needed to help administrators monitor the performance of the applications running on the VMs, manage SLAs, and set business priorities.

Intelligent management tools can, for instance, assist in optimizing the allocation of physical and virtual resources to individual VMs, in response to increased demand or to meet SLAs. These tools can for example monitor VMs' performance and utilization and optimize resource configurations accordingly. Core system components, such as CPU, memory and hard disks can have their utilization monitored and managed to adapt to changing workload conditions. For instance, automated dynamic load balancing and reconfiguration tools can help individual VMs

make best use of unused system resources during high workload situations.

Gathered statistics about memory usage can provide guidance to selecting the right amount of memory to be allocated to VMs and to the host operating system. In addition, management tools that can monitor the health of the virtual servers, predict their failures and generate alerts are essential for establishing a proactive virtual server management strategy.

Capacity planning, modeling and simulation tools will be very helpful in assessing beforehand whether the existing physical server hardware will be capable of sustaining a given set of applications, running on multiple VMs and whether application performance targets will be met. Capacity tools can also assist in minimizing the under-provisioning and over-provisioning of server resources to VMs.

4.7. Advances in Processor and Memory Technologies

A new breed of low-power, multi-core processors are contributing towards alleviating some of the performance bottlenecks in virtualized environments, especially during high workload situations. For instance, advances in core micro-architectures enable each core to execute more instructions per clock cycle, thus increasing throughput.

Newly introduced processors are being and are expected to be further enhanced to support virtualization in many ways, including (1) Support for multiple logical CPUs and new privileged instructions to accelerate communication between the hypervisor and the VMs, (2) support for integrated I/O memory management units and DMA remapping to better support I/O virtualization, and (3) better support to handle VM interrupts [7]. In addition, recent advances in intelligent memory access and caching mechanisms are contributing towards reducing memory access latency and increasing memory access efficiency [5].

5. Future Perspectives

With the very recent economic downturn, enterprises will likely consider virtualization technology as a cost-cutting measure to reduce the Total Cost of Ownership (TCO) of their datacenters. To further reduce the TCO, enterprises will be tempted to favor open source virtualization solutions, as opposed to being locked in with proprietary virtualization software. The recent Solution of Open Virtualization (SOV) initiative, led by IBM, Intel, and Novell illustrates this growing trend towards open source virtualization solutions. SOV is based on the blending of IBM's system x and Blade Center Servers and management software, Intel Xeon multicore processors and Virtualization Technology (VT) and Novell SUSE Linux Enterprise Server.

In future, we also expect to see more partnerships and collaborations between server virtualization players and processor manufacturers such as Intel and AMD. This collaboration is needed to enable virtualization solutions make the best use of multi-core processor technologies.

As the virtualization technology matures and heads towards commoditization, priority will shift from functionality to optimized performance, secured deployments, seamless interoperability, software licensing considerations, clients' education and training, and automated system management tools. It is also clear that for many small and medium size business (SMB) organizations, the migration towards virtual servers would entail building or acquiring the appropriate expertise to deal with the additional complexity of the technology.

Current trends suggest that, in future, the extra layer of hypervisor software between the VMs and the physical hardware would eventually vanish. In fact hardware-based access

mechanisms can potentially address the prevailing overhead issue that is causing most of the I/O bottleneck and which is excluding many resource intensive enterprise applications from virtualization projects. This migration from software-based towards hardware-based (built-in) hypervisors will further push a virtual server to become a commodity hardware box, supporting multiple guest operating systems and associated applications.

Finally, in the coming years, we expect to witness the convergence of server, storage, desktop, and application virtualization. This convergence will lay down the basis for next-generation enterprise virtualization, cloud computing and Infrastructure as a Service (IaaS). It is not a surprise that, last year, Gartner listed virtualization and cloud computing as the top two technologies that will dominate the IT landscape for the next few years. We also expect that enterprises will explore the options to improve the nexus between data-center virtualization and cloud computing, with a main focus towards increasing the chances that these two technologies meet. Current trends suggest a potential hybrid virtual/cloud model for datacenter resources, whereby most mission-critical applications will be running inside the virtual enterprise datacenter, while other applications (with less security, and latency requirements) will be run in the cloud.

References

1. R. Dittner et al., *Virtualization with Microsoft virtual server 2005*, Rockland: Syngress Publishing, 1996.
2. IDC, "Virtualization and multicore innovations disrupt the worldwide server market", IDC Doc# 206035, March 2007.
3. B. Posey, "Server virtualization basics for network pros," June 2008. http://searchnetworking.techtarget.com/tip/0,289483,sid7_gci1297805,00.html.
4. R. Kumar, "U.S. data centers: The calm before the storm," Gartner RAS Core Research Note G00151687, Sept 2007
5. Novell, "Solution for open virtualization helps provide server consolidation," Novell, IBM, and Intel white paper, 2007, http://software.intel.com/sites/oss/pdf/open_virtualization.pdf.
6. B. Armstrong, *Professional Microsoft Virtual Server 2005*, Indianapolis: Wiley Publishing, 2007.
7. NetXen, "The future of Ethernet I/O virtualization is here today," NetXen whitepaper, 2007. <http://www.netxen.com/technology/pdfs/FutureofEthernet.pdf>.
8. G. Lawton, "Server virtualization and the network: Site consolidation's impact on latency, March 2008. http://searchnetworking.techtarget.com/tip/0,289483,sid7_gci1305910,00.html.
9. G. Lawton, "Virtual machines present dynamic environment issues for network pros" June 2008. http://searchnetworking.techtarget.com/tip/0,289483,sid7_gci1317936,00.html.
10. N. MacDonald, "Securing virtualization, virtualizing security," Gartner Symposium/ITxpo 2007: Emerging Trends, San Francisco, April 22-26 2007.
11. P. Levy, R. Chalaka and G. Scherer, "The missing piece of virtualization: Eliminating the I/O bottleneck with IOV in virtualized servers," Neterion whitepaper, April 2008
12. PCI-SIG I/O Virtualization (IOV) Specifications. <http://www.pcisig.com/specifications/iov/>.
13. VMware, "Server consolidation and containment with virtual infrastructure," VMware solutions brief. http://www.vmware.com/pdf/server_consolidation.pdf.
14. D. Robb, "So much for simplicity," in: Getting the most from virtualization: an Internet.com Networking eBook, 2008, pp. 11-13.