

# 6

## A Most-Significant Step—“Virtualizing” Your IT Systems

“Virtualization can increase hardware utilization by five to 20 times and allows organizations to reduce the number of power-consuming servers.”

—Gartner Data Center Conference  
November 2007

The most-significant step most companies can make in their quest for green IT is in IT virtualization, as briefly mentioned in previous chapters. This chapter describes the significant concepts of virtual servers and virtual data storage for energy-efficient data centers. The descriptions include VMware and other server virtualization considerations. In addition, the virtual IT world of the future, via grid computing and cloud computing, is discussed. Although the use of grid computing and cloud computing in your company’s data center for mainstream computing might be in the future, some steps toward that technology for mainstream computing within your company are here now. Server clusters via VMware’s VMotion and IBM’s PowerVM partition mobility are here now and used in many company data centers. Both of those technologies are described in this chapter.

Based on my experience with data centers for more than ten years, I believe the most important reason to use virtualization is for IT flexibility. The cost- and energy-savings due to consolidating hardware and software are also significant benefits and nicely complement the flexibility benefits.

There are many aspects to IT virtualization. This chapter structure covers the rational, server virtualization, storage virtualization, client virtualization, grid and cloud concepts, cluster architecture for virtual systems, and conclusions.

Over the past 30 or more years, data centers have gone from housing exclusively large mainframe computers to housing hundreds of smaller servers running versions of the Windows operating system or Unix® or Linux® operating systems. Often the smaller servers were originally distributed throughout the company, with small Windows servers available for each department in a company. During the past few years, for reasons of support, security, and more-efficient operations, most of these distributed servers have moved back to the central data center. The advent of ubiquitous high-speed networks has eliminated the need for a server in the same building. These days, network access even to our homes through high-speed networks such as DSL and cable allows network performance from our homes or distributed offices to the central data center to be about equivalent to performance when your office is in the same building as the data center. The Internet was and remains the most-significant driving force behind the availability of high-speed networks everywhere in the world—including to homes in most of the developed world. When we access a Web site from our home, from the airport with a wireless connection, or from the countryside using a PDA or an air card with our laptop, we have a high-speed connection to a server in some data center. If the Web site is a popular site such as Google, the connection might be routed to any one of many large data centers.

When the distributed servers that had been in office buildings were moved in the past ten years to centralized data centers, operations and maintenance became greatly simplified. With a company server at a centralized data center, you could now call the help desk on Sunday morning and find out why you had no access, and central operations could have a technician “reboot” the server if it had gone down. So, the centralized data center provides many advantages—especially with high-speed networks that eliminate network performance concerns. However, with the rapid growth in servers used in business, entertainment, and communications, the typical data center grew from dozens of separate physical servers to hundreds of servers, and sometimes to thousands. Purchasing, operating, and maintaining hundreds of separate physical servers became expensive. The innovative solution was to consolidate perhaps ten of the separate servers into one bigger physical server, but make it appear as if there were still ten separate servers. Each of the ten

virtual servers could retain its own server name, its own Internet address (IP address) and appear—even to web developers—to be a separate physical machine (as it had been before becoming a virtual server). Costs go way down because one large physical box is much less expensive to buy than ten smaller physical boxes. Also, it’s significantly less expensive to maintain and operate (“take care of”) one big server than ten smaller servers. The analogy might be exaggerated—but it’s a bit like taking care of one big house rather than ten separate smaller houses.

In simple terms, server virtualization offers a way to help consolidate a large number of individual small machines on one larger server, easing manageability and more efficiently using system resources by allowing them to be prioritized and allocated to the workloads needing them most at any given point in time. Thus, you can reduce the need to over-provision for individual workload spikes.

In general, virtualization at the data center is applied broadly—not just to server virtualization. It provides the capability to simulate the availability of hardware that might not be present in a sufficient amount—or at all! Virtualization uses the available physical resources as a shared pool to emulate missing physical resources. Virtualization is capable of fine control over how and to what extent a physical resource is used by a specific virtual machine or server. Thus, we have the concept of virtual computer memory (which is not real memory but appears to be real) and virtual data storage.

This chapter gives details on virtualization technologies at the data center and explains how those technologies are usually the first and most-important step we can take in creating energy-efficient and green data centers.

## The Concepts of Consolidation and Virtualization

In Chapter 2, “The Basics of Green IT,” Figure 2.2, shows an overall strategy for moving to a green data center. The IT infrastructure energy-efficiency strategy consists of centralizing data centers, consolidating IT resources at those data centers, virtualizing the physical IT resources, and integrating applications. Server consolidation and server virtualization both reduce energy use by reducing the number of physical servers, but they use different methods. Server virtualization enables you to keep all your servers, but they become virtual servers when many physical servers share the same physical machine. The diagrams and descriptions of the concepts of consolidation and virtualization were based on the descriptions in the IBM red paper, “The

Green Data Center: Steps for the Journey.” (See the Bibliography for author Mike Ebbers for additional information.) These diagrams and descriptions should clarify the difference and some of the pros and cons of the two methods: consolidation and virtualization.

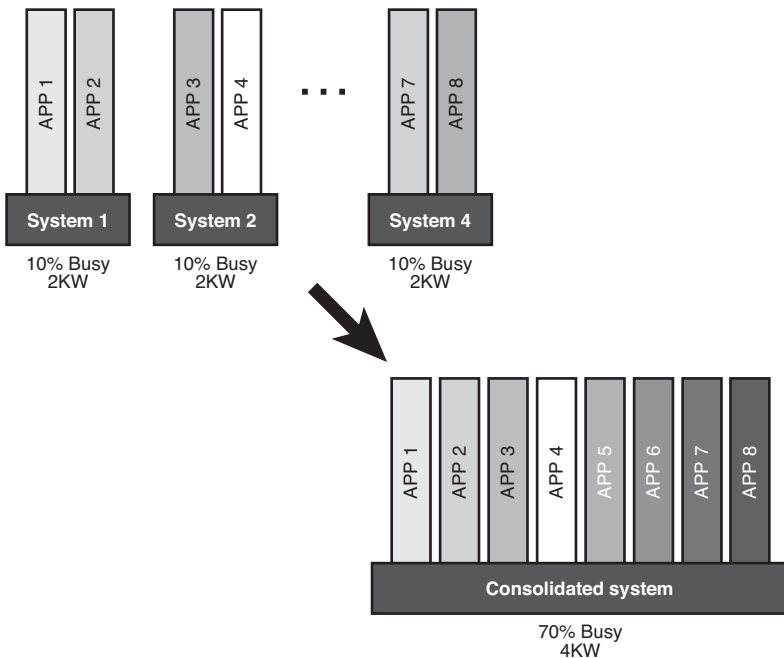
### **Consolidation: A Key in Energy Efficiency**

A common server consolidation example that I’ve seen with many projects over the past few years is the consolidation of e-mail servers. As discussed at the beginning of this chapter, for reasons of cost reduction and server management efficiency, there are significant advantages to moving servers to a central data center. As part of the distributed computing architecture where smaller servers were distributed throughout the company, we had e-mail servers that were distributed, often one for each corporate facility with often only a couple hundred users for each server. When the e-mail servers were centralized, dozens of smaller servers could be consolidated onto one or two large mail servers. This was more than consolidating the physical servers onto one large physical server; the large centralized e-mail servers only had one copy of the e-mail application. So, server consolidation refers to both consolidating physical servers and consolidating the application.

Figure 6.1 illustrates this idea of consolidation and the energy efficiencies to be gained. Let’s assume we have four systems, each running two applications (APP). Also, each machine consumes 2 kW power, 8 kW in total. However, as is often the case for small x86 servers, they are utilized at only 10 percent. If we can consolidate these eight applications to a single, more powerful server and run their operation at a utilization of 70 percent with a power usage of 4 kW, this single server can operate more energy efficiently. In addition, if we perform a simple power management technique of switching off the previous four systems, the result is a total power consumption of 4 kW and a 70 percent utilized system.

It’s important to note that a decrease in overall power consumption is not the only factor. Hand-in-hand with the power reduction goes the same amount of heat load reduction and another add-on for the infrastructure. This double reduction is the reason why consolidation is an enormous lever to moving to a green data center.

However, a particular drawback of consolidation is that none of systems 1 through 4 is allowed to be down during the time that the respective applications are moving to the consolidated system. So, during that migration time, higher demands on resources might occur temporarily.



These materials have been reproduced by Pearson IBM Press, an imprint of Pearson Education, Inc., with the permission of International Business Machines Corporation from IBM Redbooks® publication REDP-4413-00 The Green Data Center: Steps for the Journey (<http://www.redbooks.ibm.com/abstracts/redp4413.html?Open>). COPYRIGHT © 2008 International Business Machines Corporation. ALL RIGHTS RESERVED.

**Figure 6.1** Consolidation of applications from under-utilized servers to a single, more-efficient server

## Virtualization: The Greenest of Technologies

An alternate method to consolidation is virtualization, the concept of dealing with abstract systems. As discussed at the beginning of this chapter, virtualization allows consolidation of physical servers without requiring application consolidation. So, as discussed earlier, with server virtualization we can take ten servers with completely different applications and consolidate them onto one large physical server, where each of the ten stand-alone servers can retain their server name, IP address, and so on. The virtual servers still look to users as if they are separate physical servers, but through virtualization, we can dramatically reduce the amount of IT equipment needed in a data center.

Virtualization eliminates the physical bonds that applications have to servers, storage, or networking equipment. A dedicated server for each

application is inefficient and results in low utilization. Virtualization enables “car pooling” of applications on servers. The physical car (server) might be fixed, but the riders (applications) can change, be diverse (size and type), and come and go as needed.

The example in Figure 6.1 shows how specific applications were moved to another system with a better energy footprint. In the simple case illustrated, we assume all systems are running at the same operating system level. However, what if the applications require different operating system levels or even completely different operating systems? That is where virtualization comes into play.

The term “virtualization” is widely used and has several definitions:

- Can create logical instances of a computer system consisting of CPU, memory, and I/O capabilities
- Can be put together from other virtual components
- Can consist of a virtual CPU or virtual memory and disk
- Can be a virtual network between a virtual computer and the outside world

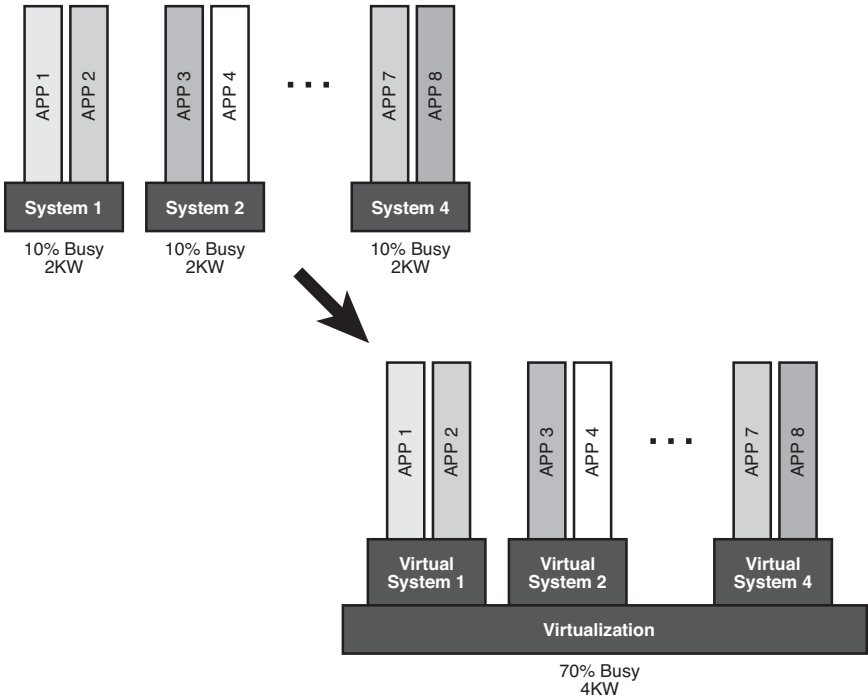
To have real work done by a virtual system, the virtual system must run on a real system. Obviously, additional intelligence is required to do this. There are pure software solutions, or a system’s firmware might offer virtualization features, or such features might be hardwired into the system. Many of the current processor architectures have virtualization features integrated, which can be taken advantage of by software solutions such as the IBM System z<sup>®</sup> and p machines. In the field, various other solutions are available, such as VMware Server, VMware ESX, Microsoft Virtual Server, and Xen.

To continue with our example, using virtualization gives a slightly different picture, as shown in Figure 6.2. Instead of moving the applications to the consolidated server, we now virtualize the existing systems 1 through 4 on our consolidation target. The effect is clear: Not only is the application moving, but also its complete operating environment has moved with it. Taking a closer look, we find other attractive features, as follows:

- Consider the three separate systems. To communicate, they require a network infrastructure such as NICs, cables, and switches. If our virtualization system supports network virtualization, this infrastructure is no

longer needed. The virtualized systems can communicate using the virtualization system's capabilities, often transferring in-memory data at enormous speed. Performance and energy efficiency increase because the network components are dropped. Once again, this method reduces the need for site and facilities resources.

- Each of the separate systems has its own storage system, namely disks. The virtualized systems can now share the disks available to the virtualization system. By virtualizing its storage, the virtualization system can provide optimal disk capacity—in terms of energy efficiency—to the virtualized systems.



These materials have been reproduced by Pearson IBM Press, an imprint of Pearson Education, Inc., with the permission of International Business Machines Corporation from IBM Redbooks® publication REDP-4413-00 The Green Data Center: Steps for the Journey (<http://www.redbooks.ibm.com/abstracts/redp4413.html?Open>). COPYRIGHT © 2008 International Business Machines Corporation. ALL RIGHTS RESERVED.

**Figure 6.2** Virtualization enables us to consolidate systems, keep the same server names, and so on.

## Server Virtualization

This section discusses the techniques that are available for server virtualization, the most attractive approach to consolidation. In many cases, it is the easiest and most-effective way to transfer workload from inefficient, under-utilized systems to efficient, well-utilized equipment.

### Partitioning

**Partitioning** is sometimes confused with virtualization, but the partitioning feature is really a tool that supports virtualization. Partitioning is the capability of a computer system to connect its pool of resources (CPU, memory, and I/O) together to form a single instance of a working computer or **logical partition (LPAR)**. Many of these LPARs can be defined on a single machine, if resources are available. Of course, other restrictions apply, such as the total number of LPARs a machine can support. The power supplied to the existing physical computer system is now used for all these logical systems, yet these logical systems operate completely independently from one another. LPARs have been available on the IBM System z since the late 1980s and on System p® since approximately 2000. Although the System z and System p partitioning features differ in their technical implementations, both provide a way to divide up a physical system into several independent logical systems.

### Other Virtualization Techniques

Many virtualization techniques are available, in addition to partitioning. Popular in the market are the VMware products, Xen and Microsoft Virtual Server. Also, hardware manufacturers extend their products to support virtualization.

VMware ESX Server and Microsoft Virtual Server come with a hypervisor that is transparent to the virtual machine's operating system. These products fall into the full virtualization category. Their advantage is their transparency to the virtualized system. An application stack bound to a certain operating system can easily be virtualized, if the operating system is supported by the product.

VMware offers a technology for moving servers called VMotion. By completely virtualizing servers, storage, and networking, an entire running virtual machine can be moved instantaneously from one server to another. VMware's



VMFS cluster file system allows both the source and the target server to access the virtual machine files concurrently. The memory and execution state of a virtual machine can then be transmitted over a high-speed network. The network is also virtualized by VMware ESX, so the virtual machine retains its network identity and connections, ensuring a seamless migration process. IBM's System p Live Partition Mobility offers a similar concept.

Xen uses either the paravirtualization approach (as the POWER™ architecture does) or full virtualization. In the partial approach (paravirtualization), virtualized operating systems should be virtual-aware. Xen, for example, requires virtual Linux systems to run a modified Linux kernel. Such an approach establishes restrictions to the usable operating systems. However, although they are hypervisor-aware, different operating systems with their application stacks can be active on one machine. In the full approach, the hardware, such as Intel's Vanderpool or AMD's Pacifica technology, must be virtual-aware. In this case, running unmodified guests on top of the Xen hypervisor is possible, gaining the speed of the hardware.

Another technique is operating system level virtualization. One operating system on a machine is capable of making virtual instances of itself available as a virtual system. Solaris containers (or zones) are an example of this technique. In contrast to the other techniques, all virtualized systems run on the same operating system level, which is the only operating system the machine provides. This can become a limiting restriction, especially when consolidating different server generations. Often the application stack is heavily dependent on the particular operating system. We reach a dead end when we want to consolidate servers running different operating systems such as Windows and Linux.


## Storage Virtualization

Computer systems are not the only candidates for virtualizing; storage can be virtualized, too. This section describes IBM SAN Volume Controller, which provides a virtual pool of storage consisting of SAN-attached physical storage devices.

### **IBM SAN Volume Controller**

The **SAN Volume Controller (SVC)** is a hardware device that brings storage devices in a SAN together in a virtual pool. This makes your storage appear as one logical device to manage. To the connected computers, SVC

offers virtual disks as ordinary SCSI devices. On the SAN side, SVC integrates various storage subsystems, even multivendor, and takes care of the correct block mapping between the SAN devices and the virtual disks for the computers. Figure 6.3 illustrates how it works.



These materials have been reproduced by Pearson IBM Press, an imprint of Pearson Education, Inc., with the permission of International Business Machines Corporation from IBM Redbooks® publication REDP-4413-00 The Green Data Center: Steps for the Journey (<http://www.redbooks.ibm.com/abstracts/redp4413.html?Open>). COPYRIGHT © 2008 International Business Machines Corporation. ALL RIGHTS RESERVED.

**Figure 6.3** Storage virtualization: virtual view and physical view

The following points make the SVC an attractive tool for an energy-efficient storage strategy:

- Data migration from older to newer, more efficient systems can happen transparently.
- Tiered storage enables you to use media with a smaller energy footprint while the SVC cache improves its performance.
- Consolidation of the system's individual storage devices to virtual storage has the same effect—increasing storage utilization—as is shown for server virtualization.

Storage virtualization requires more effort than server virtualization, often requiring us to rethink the existing storage landscape. During consolidation, large amounts of data must be moved from the old systems to the consolidated storage system. This can become a long task that requires detailed planning. However, when done, the effect can be enormous because now storage can be assigned to systems in the most flexible way.

## Virtual Tapes

Tapes are the cheapest medium on which to store data. They offer the largest storage volume at the lowest cost, which is the reason that they are the optimal backup medium. Tapes have a long latency compared to disks. This is not always a benefit. Currently, data centers are faced with a time limitation for backing up and probably restoring their data because the time frames for backups shrink, while the amount of data to back up expands. For this reason, many sites prefer large disk-based backup systems instead of tapes.

Tape virtualization might be a solution to this problem. A virtual tape server behaves just like a tape library, but a very fast one. This is made possible with internal disk arrays and a migration strategy to export to and import from real tape libraries.

## Client Virtualization

A great potential in energy savings is client, or desktop, virtualization. Various studies have estimated energy savings of more than 60 percent by using client virtualization. In a typical workplace, installed PCs show very low average usage rates. Except when launching an application, an office PC spends most of its time waiting for the user to press a key or click a mouse. However, the PC continues to need a considerable amount of energy to operate, heats up its surrounding environment, and produces noise. Desktop virtualization can dramatically improve the situation.

The underlying principal of client virtualization is to replace the office workstation with a box having a much smaller energy footprint. The needed computing power is moved into the data center. Today's virtualization techniques make this approach even more attractive. The concept of client virtualization—often called thin-client computing—is not a new concept and goes back at least 15 years. In fact, thin-client computing where the server does all the processing is similar in concept to the terminals we used to connect to the mainframe before the advent of the PC.