CHAPTER 4

REALLOCATING
Unused Server and Storage Hardware

Virtualization's abstraction of resources offers some compelling solutions for keeping around old equipment past its usefulness for primary workloads.

data center servers & storage

CHAPTER 3

Reallocating
Unused Server and Storage Hardware

Virtualization's abstraction of resources offers some compelling solutions for keeping around old equipment past its usefulness for primary workloads.

By Greg Shields

If you've made the jump to virtualization, you've likely also made the jump to hardware virtualization extensions like Intel-VT or AMD-V.

Intel and AMD first-generation processors with virtualization extensions brought a set of useful enhancements to virtual environments. Adding processor assistance to the hypervisor enabled significant improvements in performance and scalability. Some hypervisors, such as Microsoft's Hyper-V and Citrix Systems Inc.'s XenServer, actually require processor assistance in order to accomplish their mission.

But technology has never rested for long on its laurels. In late 2008, a second generation of virtualization-enabled processors entered the market. Buzzwords like Nehalem and Opteron have been quickly associated with the newest hardware you can purchase from major server vendors. These second-generation virtualization extensions, with codenames Intel-EPT and AMD-RVI, took a new and even greater leap into increasing raw performance in virtual environments.

Adding processor assistance to the hypervisor enabled significant improvements in performance and scalability.

Intel-EPT and AMD-RVI add a dra-

matic performance boost to your server workloads. In one example, called Project Virtual Reality Check, rigorous testing by the LogIn Consultants group discovered that Intel's

SERVER VIRTUALIZATION, ABSTRACTION AND DRIVER INDIFFERENCE

TACTICS IN REALLOCATING OLD SERVERS

TACTICS FOR REALLOCATING STORAGE

Although the privilege of designing, purchasing and installing new servers is still a fun task for most of us, the process of throwing away the old to make way for the new often represents a sunk cost.

Nehalem class of processors—which include support for Intel-EPT—nearly doubled the number of concurrent users that a single server could support.

These continual improvements to hardware are important in their own right. As time passes, Moore's Law suggests that hardware will grow ever more powerful. In fact, improvements to raw power tend to happen at a rate that is even faster than a business' average server refresh cycle. By refreshing your servers every three to five years—which is a relatively average period of time—you're sure to discover that each new refresh brings with it new technologies, new hardware architectures and a whole new set of drivers to manage.

That's exactly the story that virtualization's early adopters are discovering today. Companies that invested in first-generation hardware probably look longingly at what's new and exciting in its second generation. Although the privilege of designing, purchasing and installing new servers is still a fun task for most of us, the process of throwing away the old to make way for the new very often represents a sunk cost.

There are smarter ways to deal with reallocating your unused server and storage hardware. You no longer need to destroy the hard drives as you provision the chassis to the garbage can. Every server refresh doesn't have to involve a resulting fire sale on eBay. With a bit of planning, a specific tiering-out of your system priorities and a dash of virtualization, you can breathe new life into your old servers—even as you bring in replacements with all of their powerful new capabilities.

SERVER VIRTUALIZATION, ABSTRACTION AND DRIVER INDIFFERENCE

Recall, first, that virtualization is at its core all about the abstraction of resources. In a virtualized server environment, every virtual machine (VM) is presented with the exact same surface of drivers and hardware inter-

faces. So, one job of the hypervisor is to translate resource requests into a language that the hardware can understand where it is installed.

You can visualize this like a set of puzzle pieces. **FIGURE 1** shows two different pieces of physical hardware. Perhaps the first is a Dell server, while the second is from HP. Or, maybe these two servers are from different generations of the same hardware vendor. Each server's resources require a set of specialized driver code to operate, as represented by the black puzzle pieces that plug into the green hypervisor layer.

When they plug into the hypervisor, the first server's puzzle pieces are somewhat different from the second server. But the hypervisor layer handles this "translation" as one of its functions. No matter which server it's installed on, its translation presents a uniform set of green puzzle pieces to the VMs on top.

NOTE: By abstracting hardware, resources like processor power and RAM become commodities that you can distribute to virtual machines as you see fit.

Virtualization's resource abstraction grants you a measure of driver indifference with your servers. VMs that you create today can seamlessly operate on top of today's servers just as well as yesterday's servers. The only remaining difference is the level of resources that the server provides.

FIGURE 1: *Virtualization abstracts physical resources into a uniform surface for virtual machines.*

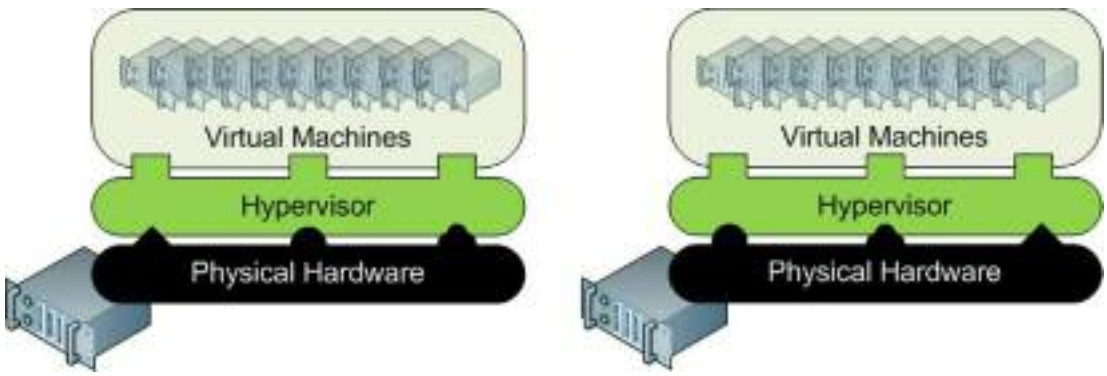

TACTICS IN REALLOCATING OLD SERVERS

Abstraction of resources means that you don't have to throw your "old" servers in the garbage as you replace them with newer and more powerful hardware. Your old servers will need to support the minimum requirements of your chosen hypervisor. Those old servers may not be able to support the same scale of VMs, but virtualization means that they'll retain the capacity of supporting any VM, should you need them.

This being said, there are some common tactics in making decisions about what to do with your old hardware. Many data centers are at or approaching their maximum capacity, making the disposal of old equipment a priority from the perspective of space management. Other enterprises simply don't want the added administrative headaches that are associated with keeping old equipment around past its warranty lifecycle.

Yet, today's use-what-you've-already-got financial climate makes a compelling reason for keeping some of this equipment around for other needs. To that end, consider a few of the following tactics as potential ways to save money by breathing new life into your old servers:

- **Add high availability to critical services.** Your environment probably already has a sense of which servers, and their installed workloads, are high priority versus low priority. When email systems, production databases or critical file servers go down, it creates a hardship on your company. As a result, you probably take extra care when working with these high-priority systems, double-checking your work before performing common maintenance tasks or installing patches and updates.

But bad things do happen. Even the most careful of systems administrators occasionally make the mistake that causes high-priority failures. Even the most well tested patches sometimes find a conflict on that high-priority server once it's installed.

That's why many software vendors today are focusing their efforts on high-availability offerings. Today's computing services have become so critical to business operations that even a few minutes or hours of downtime can break the budget.

The problem, however, with high-availability offerings is their need for extra hardware. Every HA solution for a service requires at least one extra server "at the other end" that can take over when the primary server goes down—and it's expensive to host that other end. Its backup server does little until the outage occurs because the backup servers at that other end have very little to do until then.

For many, secondary servers in high-availability pairs represent an excellent use for old servers. In gener-

SERVER VIRTUALIZATION, ABSTRACTION AND DRIVER INDIFFERENCE

TACTICS IN REALLOCATING OLD SERVERS

TACTICS FOR REALLOCATING STORAGE

al, the secondary server in an HA pair often requires substantially less processing power than the primary—until the primary goes down.

When a virtual server atop your "old" hardware sits atop the same surface as another atop your "new" hardware, it means that the virtual secondary in your HA pair can reside successfully atop last-generation hardware that you would otherwise throw away. As it performs its duties in remaining synchronized with the primary, it may not need the powerful level of processing that your first-tier servers have. But when the primary dies, that secondary server will find itself servicing clients.

That's when virtualization's live migration features come in handy. Virtual secondary servers needn't remain on aging hardware once they've been promoted to primary. Your hypervisor platform's live migration capabilities can easily reprovision that VM to better hardware when the need occurs. After the problem is fixed, that server can be relocated back to its original place on your hardware.

■ **Provision low-priority servers to no-cost hypervisors.** Secondary servers in HA pairs aren't your only low-priority options. Every IT environment has its own set of servers that need to stay up but that don't cause a major problem if they happen to go down. Update servers, antivirus servers, test and development servers, even the server hosting your IT file share—any of these can represent a low-priority server that doesn't need substantial processing power and won't affect the business if a problem occurs.

Today's hypervisor product ecosystem breaks down into two distinct camps. On one side are hypervisor products that come with a cost and include high-end capabilities that are critical for your high-value workloads.

On the other side are hypervisors that you can implement for no added cost. Although these almost-free alternatives don't often provide the same level of management functionality and flexibility as for-cost products, they are a successful platform for general-purpose virtualization.

Like 'em or not, you have to admit that no-cost hypervisors still successfully *virtualize.*

The lack of cost for this second class of hypervisors makes them a compelling option for your low-priority servers. Once you've made the business decision to rate these servers at a lower tier of priority, why should they require the same high-priority capabilities as their mission-critical brethren? If they don't, and you relocate their processing to a no-cost hypervisor atop a server you would have thrown away anyway, it can mean a massive cost savings.

At this point, you may be asking: "If I'm using two different hypervisor

SERVER VIRTUALIZATION, ABSTRACTION AND DRIVER INDIFFERENCE

TACTICS IN REALLOCATING OLD SERVERS

TACTICS FOR REALLOCATING STORAGE

suites for my virtualizing, wouldn't it complicate my management?" At first blush, it can. Two different hypervisors will indeed come with two different hypervisor management tools from the vendor.

However, third-party management toolset vendors are addressing the problem of multi-hypervisor management . These vendors have created an entirely new set of administrative solutions that layer atop all of the major hypervisors. That position atop all hypervisors enables them to manage each at the same time, all beneath a single pane of glass.

■ **Leverage old equipment for hypervisor cluster failover.** Odds are that your hypervisor of choice also comes with a solution for failing over VMs when a host fails. Whether that's XenServer's or VMware's HA capability or Microsoft's Windows Failover Clustering, virtualization and clusters have been two complementary technologies for a long time.

But the problem with clustering technology isn't in its protective capacity. A properly configured cluster from any vendor will successfully relocate a VM to a surviving cluster member after a failure. Clustering's biggest problem is in designing it with enough resources to actually survive that failover.

Consider the situation where you've created a cluster out of four virtual hosts. Into that cluster, you've added your needed VMs to fill the cluster to some percentage of its total capacity. The issue here arrives when filling that cluster goes beyond the capacity of that cluster minus one of its hosts.

For example, for a four-node cluster, you will generally want to load the cluster to somewhere around three-quarters of its total capacity. This "cluster reserve" is necessary so that when any of your cluster nodes fail, your relocated virtual environment will retain the same level of service.

Reallocating otherwise unused servers—or those you would have discarded—for that cluster reserve is also effective. Previous-generation servers may not have the level of performance you require for normal operations, but their ability to receive VMs during the occasional failover incident makes them invaluable when that incident occurs.

REMINDER: Some virtualization technologies today have the ability to oversubscribe resources such as RAM. This is particularly useful to ensure virtual machine availability during a cluster failover situation where your cluster is overloaded. But to maintain best performance, it's still a good practice to reserve some cluster resources for failover.

■ **Enhance your disaster recovery posture.** The final issue is the potential to extend your virtual infrastructure to a completely different site. Virtualization flexibility in hosting any VM atop any physical server has brought about a resurgence in disaster recovery planning and implementation. Today, even the smallest business can inexpensively implement a fully realized disaster recovery solution.

When creating a DR solution, it's important to recognize that many businesses will not require full functionality of every service in the hours or days after a disaster. Of course, services such as email, core databases and file access are critical to operations, so they must be quickly restored when disaster hits.

But, an even larger set of services don't need to be given immediate priority. And others can be given priority toward making them available, but they don't have to operate at the same level of performance until normal operations resume.

This final tactic for reallocating your old servers involves repurposing them for use only when a disaster occurs. With the right replication technology in place, you can reliably copy VMs and their disk files to an alternate site.

Depending on your business recovery time objective, those servers can be just sitting in that remote site. Or, they can be fully connected, powered on and receiving up-to-the-second updates from their VM primaries. Their relocation to the secondary site gives them a second career in maintaining your business operations during a disaster.

TACTICS FOR REALLOCATING STORAGE

Servers aren't all you need to be aware of as you go through the evolutionary process of upgrading your hardware. Today's servers include their fair share of storage. Even servers of a few years ago are likely to be configured with hundreds of gigabytes of storage space, which can be used for new purposes.

Each tactic for reallocating storage requires some element of storage space for hosting VM disk files. Whether you're looking for a secondary server for an HA solution, the repositioning of a low-priority server to a no-cost hypervisor, one more machine to add to your cluster reserve or a set of them for disaster recovery, each solution requires storage space for VM processing.

Nearly every data center today finds itself moving toward SAN solutions for centralization of storage and storage processing. It's a good move from the perspective of management and elimination of waste. SAN solutions—particularly those that include thin provisioning technologies—can discretely provision just the right amount of storage to servers and vir-

SERVER VIRTUALIZATION, ABSTRACTION AND DRIVER INDIFFERENCE

TACTICS IN REALLOCATING OLD SERVERS

TACTICS FOR REALLOCATING STORAGE

tual servers alike. Segregating storage processing from server processing further protects data while adding improvements to performance.

Yet almost every environment has its share of local storage as well. Much of that local storage remains attached to the very servers that you no longer need or want to use. Without some plan for reuse, replacing these servers with newer models discards a valuable investment in that storage.

Consider the following two tactics to get more value out of the storage in your aging equipment:

- **Upgrade to disk-to-disk backups.** Even with all of the problems they've solved, today's virtualization environments also introduce many complexities with backups. And adding virtualization to data centers introduces multiple locations from which backups can source.

In a virtual environment, you can back up VMs via agents that are installed into the machines. Alternatively, you can install an agent to the host and back up the VM as a single disk file.

Both solutions have benefits and drawbacks. Neither is best in all situations. And as a result, some backups solutions providers have created new solutions—a "third way"—for completing backups. This new approach involves file system filter drivers that monitor your servers' file system for changes while replicating any they find to external storage. Doing so provides all of the benefits of the two other approaches but without their limitations.

This backup method requires online disks to be available for storing the backup. Its approach simply doesn't integrate well with the point-in-time types of backup that are commonly associated with tape drives.

The concern, of course, in these solutions is in having enough online storage space available so you can actually back up everything you need. The space doesn't need to be exceptionally well performing. In fact, it is most economical when high-performing SAN storage isn't its primary target.

Leveraging your old servers and their attached storage for this use can be an excellent extension of their lifecycle. Many of today's disk-to-disk backup solutions also archive backups to tape, making these solutions necessary for daily needs as opposed to long term.

- **Add SAN space and software-based SANs.** Finally, another class of SAN storage is slowly becoming popular. These SAN solutions install to regular Windows or Linux servers as an application, which creates a software-based SAN out of available disk space and allows you to create entirely new islands of exposable SAN storage out of disks that you might other-

wise throw away. The second generation of these solutions is only now beginning to incorporate the same kinds of storage HA that you're used to seeing in their hardware-only counterparts.

Consider looking at some of the options available for software-based SAN storage. You might not necessarily replace your existing production SAN with their technologies just yet. But implementing them in remote offices, for disk-to-disk backups or even for storage in a disaster recovery site is an excellent solution for keeping that storage space around a bit longer.

As you can see, virtualization's abstraction of resources creates some compelling solutions for old equipment past its usefulness for primary workloads. Your aging hardware stands to gain a full and complete second life as the host for lower-priority services or for expanding your virtual infrastructure for just the times you need it. Even your storage can gain new uses for down-level purposes if you consider the options available. ■

Greg Shields, a Microsoft MVP and VMware vExpert, is an independent author, speaker and IT consultant, as well as a founding partner at Concentrated Technology. With nearly 15 years in information technology, he has developed extensive experience in systems administration, engineering and architecture specializing in Microsoft OS, remote application and virtualization technologies. Shields has authored or contributed to 10 books and countless white papers and webcasts.

SearchDataCenter.com

Cathleen Gagne
Editorial Director
cgagne@techtarget.com

Matt Stansberry
Executive Editor
mstansberry@techtarget.com

Christine Casatelli
Editor
ccasatelli@techtarget.com

Marty Moore
Copy Editor
mmoore@techtarget.com

Linda Koury
Art Director of Digital Content
lkoury@techtarget.com

Jonathan Brown
Publisher
jebrown@techtarget.com

Peter Larkin
Senior Director of Sales
plarkin@techtarget.com

TechTarget
275 Grove Street
Newton, MA 02466
www.techtarget.com

- **Server hardware: Making the right choice**
- **Simplifying the SMB IT Experience—Dell Expands Data Center Options**
- **Total cost of ownership (TCO) of Dell PowerEdge M710 and HP Proliant BL685c G5 blade solutions**

About Dell, Inc. and Intel:
Dell and Intel are strategic partners in delivering innovative hardware solutions to solve your most challenging IT problems. Together we are delivering virtualization optimized solutions to maximize the benefits of any virtualization project.