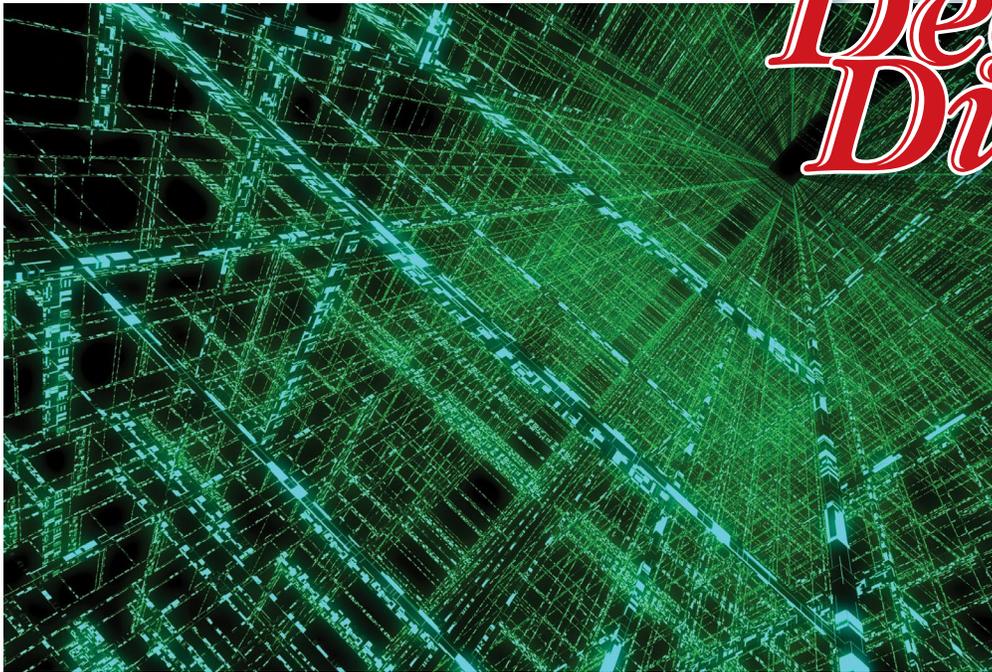


Server Virtualization

*Deep
Dive*



Best practices for a virtual infrastructure

Building a solid virtualization foundation

Dramatic cost and agility benefits make server virtualization an essential part of any data center. Here's how to plan, deploy, and maintain a sound virtual infrastructure

By Paul Venezia

FEW TECHNOLOGIES have become a fundamental part of the data center as quickly as server virtualization. That's because the basic value proposition is so easy to grasp: When you run many logical servers on a single physical server, you get a lot more out of your hardware, so you can invest in fewer physical servers to handle the same set of workloads. It sounds almost like found money.

The details, of course, are more complicated. The hypervisor, a thin layer of software upon which you deploy virtual servers, is generally wrapped into a complete software solution that incurs some combination of licensing, support, and/or maintenance costs (depending on which virtualization software you chose). And very likely you will need to upgrade to server processors that support virtualization.

On the other hand, reducing the number of servers yields indirect cost savings — less space to rent, less cooling to pay for, and of course lower power consumption. Even more compelling is virtualization's inherent agility. As workloads shift, you can spin up and spin down virtual servers with ease, scaling to meet new application demands on the fly.

The path to rolling out a virtualized infrastructure has its share of pitfalls. You need to justify the initial cost and disruption in a way that does not create unrealistic expectations. And you need to know how to proceed with your rollout, to minimize risk and ensure performance stays at acceptable levels.

MAKING THE CASE FOR SERVER VIRTUALIZATION

It's pretty easy to sell server virtualization. Who

doesn't want to get the most possible use out of server hardware? In fact, the basic idea is so compelling, you need to be careful not to oversell. Make sure you account for the likely capital equipment, deployment, training, and maintenance costs. The real savings achieved by virtualization, as with so many other new technologies, tend to accrue over time.

Most virtualization deployments require new hardware, mainly because hypervisors require newer processors that support virtualization. So the best time to roll out virtualization is when you need to add servers to your existing infrastructure or when it's time to replace aging hardware.

The superior efficiency of newer servers will help make your case. Begin by calculating the power consumption and cooling levels the current infrastructure requires. (Ideally, this should be conducted on a server-by-server basis, which can be time consuming, but will result in far more accurate numbers.) Then check the same specs for the hardware you plan to buy to get an idea of any power and cooling cost savings.

Add the fact that you will be using fewer physical servers for the same workloads, and your proposed virtualized infrastructure will look very, very good compared to the existing one. If the new hardware is sufficiently powerful, you may be able to run many logical servers on each physical unit.

Unfortunately, determining how many virtual servers will fit on a physical host is never an exact science. But there are tools that can help. Some server consolidation tools will even allow you to specify the make and model of your current and planned hardware, and will monitor your existing infrastructure for a period of time.

Armed with all that data, you can run reports that show exactly how many virtualization hosts you'll need,



what type, and your expected ratio of virtual servers to physical hosts. Some will even calculate the expected power consumption and cooling capacity for the new infrastructure. Investigate the options available from VMware, Microsoft, and others in order to get the most accurate data before you leap into any virtualization project.

But again, don't oversell. It's important for everyone to realize that reducing the number of physical servers does not mean reducing logical servers — and does not necessarily lead to reducing IT staff. In fact, it's generally beneficial to hire a competent consultant to help plan any virtualization endeavor. Although the basic concepts are simple, the planning, design, and implementation stages can be quite tricky without proper knowledge and experience.

TRAIN BEFORE YOU FIRE IT UP

It's also important to take into account training for existing staff. Virtualizing an existing IT infrastructure means changing the structural foundation of the entire computing platform; in a sense, you're collecting many eggs into a few baskets. It's vitally important that IT admins are well versed in managing this infrastructure when it goes live, as virtualization introduces a number of hazards that must be avoided.

If at all possible, make sure your staff is trained before you embark on a full-blown virtualization implementation. Your chosen vendor should provide many options for specific training, or online classes at the very least. In addition, take advantage of the evaluation periods that many virtualization platforms offer. For example, VMware's enterprise framework can be downloaded, installed, and run for 60 days without purchase, and that time can prove invaluable to familiarize admins with the tools and function of the proposed environment. There's no substitute for this type of hands-on experience.

Don't make the rookie mistake, however, of letting your sandbox training implementation turn into your production platform. When it's time to fire up a production virtualization foundation for the first time, make sure it's with a clean install of all components, not a migration from a training tool.

It's also essential to ensure that training isn't limited to the software. Hardware considerations are crucial to

a virtualization implementation, from the number of Ethernet interfaces, to CPU choices, RAM counts, local and shared storage — the whole works. It's vitally important that your admins are well versed in the day-to-day operation and functions of supporting tools like SAN array management interfaces, Ethernet, or Fibre Channel switches. In a virtualized environment, a mistake that affects a single port on a single server can affect all the virtual servers running on that particular host.

OUT WITH THE OLD

One major benefit of embarking on a virtualization project is that it gives IT the opportunity to jettison old hardware and old frameworks. There's never a better time to inspect the whole infrastructure and identify components that have fallen through the cracks, aren't necessary anymore, or can easily be folded into other tools or projects.

As you step through the planning stages of virtualization, you should pay close attention to anything that can be culled from the back-room herd without too much pain. It will ease the transition and cut down on the number of servers that need to be migrated or rebuilt on the virtualized foundation.

It's also a good time to inspect the network requirements of the proposed solution. Ethernet trunking to the physical hosts is generally a must in any reasonably sized infrastructure. By trunking, you enable the virtual machines to participate on any trunked network, rather than just the layer-2 network that the host is directly connected to. You can also switch hosts between networks on the fly. It's a very easy way to bring a substantial amount of flexibility into the mix.

Are you planning on running any virtual servers that need to be linked to a DMZ network? If so, it's best that they have a dedicated interface for that traffic on each host, although it's possible to trunk those connections as well. Generally speaking, you should maintain physical separation of trusted and untrusted networks; adding another network interface to your hosts is a minimal cost.

PICKING YOUR SPOTS

When planning a virtualization implementation, make sure you start with the obvious candidates and roll



out at a reasonable pace. It's best to start by migrating or rebuilding ancillary servers rather than the core servers like your email or ERP database servers. Depending on their load and the hypervisor and hardware you've chosen, core servers can probably be virtualized, but it's best to begin with the low-hanging fruit.

In many infrastructures, a good place to begin is with the domain controllers, intranet servers, development servers and the like. Once the infrastructure is stable and familiar, look to bringing the bigger servers into the fold.

One thing that's fairly critical is to make sure that you retain some non-virtualized servers that have duplicate roles to a virtualized counterpart. The biggest one here is Microsoft Active Directory domain controllers. It's best to virtualize one of them, but leave another as a stand-alone physical server. In the event of a total power loss or other cataclysmic event, maintaining a domain controller outside the virtualized infrastructure can substantially help reinstatement efforts.

Also, it's generally best to run your virtualization management components outside the virtualized infrastructure unless it's a very small implementation. This protects against the chicken-and-egg issue of trying to bring up a downed virtualized infrastructure when the management components themselves are VMs. Virtualizing these infrastructure pieces can be done, and may even be condoned by the vendors, but exempting them from the virtual infrastructure can greatly simplify repair efforts in the face of a significant problem.

PICKING YOUR PLATFORM

The obvious leader in virtualization platforms is VMware. They've been in the game the longest and easily have the most mature x86-based virtualization solution on the market. They're also expensive -- you pay for all of those accolades.

The other major players are Citrix's XenServer and Microsoft's Hyper-V. Citrix's solution is based on the open-source Xen virtualization project and is a well-rounded offering, with features such as live migration. Microsoft's Hyper-V is arguably the cheapest option, but is well behind the other players in terms of features and robustness.

If you're a purely Microsoft shop and are looking to

virtualize only a handful of Windows servers, Hyper-V may be attractive. If you're looking at deploying a larger virtualization environment that comprises a few different operating systems or more than a dozen or so virtual servers, you'll be better served looking elsewhere. Eventually Microsoft may catch up to the competition in terms of features, management, and stability, but it's not quite there yet.

The other option is to roll your own. This method is certainly not for the faint of heart, and involves the use of several open-source tools, such as the open-source version of Xen, possibly the use of Linux's KVM virtualization tools, or the use of VMware's free ESXi embedded hypervisor. This is the cheapest solution by far, but it also requires highly skilled administrators and precludes many enterprise-level features such as live migrations and centralized management. If you're going to be strictly virtualizing Linux servers and have the requisite skills available, this may be a good option. If not, it's definitely better to choose a commercial solution.

VIRTUAL APPLICATIONS

When doing due diligence for a proposed virtualized infrastructure, don't forget the applications. It may seem like a no-brainer that a given application will function on a virtual server, but you may run into problems with licensing and support. Although it's not as prevalent now as in the past, some ISVs refuse to support their applications if they're run on virtual servers. In many cases this is somewhat of a cop-out, since there's no technical reason their products would have a problem with a virtual server, but that doesn't fix the problem when you're on the phone with product support. Make sure your critical apps are OK to be virtualized before you make the cutover.

It's not just smaller applications that have problems like this. Microsoft doesn't support several of its key applications on some or all virtualization platforms. Also, there may be problems with keyed software that requires the presence of specific USB hardware license keys or other dongles to function. Not all these issues can be sorted out, so make absolutely certain you're in the clear with your vendors before trying to pull them into the virtual pile.



SERVER SELECTION

The first criterion in selecting an appropriate server platform is RAM. Virtualization host servers can chew through RAM like nothing else, but are generally surprisingly light on CPU utilization. It's not uncommon for a dual-socket, 2.66GHz quad-core Intel server with 32GB of RAM to be able to host 40 light- to medium-loaded virtual machines with little CPU impact.

Naturally, those virtual machines won't be core database servers running heavy queries, but general, day-to-day operational and application servers can coexist in large numbers on the right hardware using higher-performance hypervisors such as VMware's ESX.

The second criteria is the number of network interfaces. At minimum, a virtualization host server should have four. Generally, you'll find that six or more is preferred, especially if you want to implement high-availability and bond interfaces together to prevent downtime.

The general rule of thumb is to dedicate a bonded pair of gigabit NICs to the front-end networking tasks, and run 802.1q trunking across them. If you're using iSCSI or NFS as shared storage, then dedicate another pair of interfaces to that task, either with link aggregation or simple link failover. If you're planning on using Fibre Channel, you should have a pair of HBAs in each server for that purpose.

In many cases, this is enough, but if you wish to connect some virtual machines to other networks, such as the aforementioned DMZ, you may want to add one or more NICs to each server to provide for that physical separation. It's also a good idea to dedicate an interface to virtual machine live migrations and management.

At the other end of this discussion we find 10G Ethernet. With 10G, most implementations will only require two 10G links for redundancy, and all storage, host, management, and migration traffic can flow out the same two pipes. The vendors may provide tools to enforce QoS rules on that traffic as well, so you can cut that 10G pipe up into logical interfaces that have bandwidth limits and priorities.

Beyond the internals of the server, the other question is the form factor. The ideal virtualization server is a box with no hard drives, just CPUs, RAM, and I/O. Many hypervisors can be installed on flash cards or

drives and booted directly to RAM, while others require hard-drive-based installation. If you choose the latter, ensure that the host servers support RAID1 mirroring and use two hard drives in each host for redundancy. Alternatively, you can leverage your SAN to boot the servers, though this is relatively more complex than the other solutions.

If you're planning on implementing a significant number of physical hosts, you may want to investigate a blade platform. Modern blade servers are highly efficient and offer a massive amount of computing power and resources in a small package that is generally cheaper to run and cool than the same number of 1U servers. In addition, centralized management tools are available for most blade chassis that make managing a virtualization host farm simpler than a collection of distinct servers.

When requesting quotes for a virtualization farm, ask for the same hardware in both 1U and blade form factors, and see how the costs compare.

STORAGE CONSIDERATIONS

Any reasonably sized virtualization implementation must make use of shared storage. This is due to the farmed nature of enterprise-level virtualization platforms. For virtual servers to be able to shuttle between hypervisors without skipping a beat, all the hosts must be able to address the same storage at the same time.

This storage can come in many forms: iSCSI, NFS, or Fibre Channel. At the moment, the fastest is likely Fibre Channel, but it's also the most costly. NFS and iSCSI are well suited for production virtualization implementations, assuming the chosen hypervisor supports it. iSCSI in particular provides a significant bang for the virtualization buck, and is the mainstay of a large number of mid- to large-sized virtualization builds.

If you're going with Fibre Channel, the only requirement is to choose HBAs that are supported by the virtualization vendor, and obviously to configure all switches and storage appropriately. With iSCSI, you may derive some benefits by using iSCSI accelerators for the storage links. In high-traffic situations, these interfaces can accelerate storage throughput to iSCSI targets, and are not terribly expensive. That said, standard gigabit interfaces should work fine in many cases.



If you're looking for raw performance, the LUNs cut for the virtual servers should be built on RAID10 volumes, which provide the protection of mirroring with the speed of striping. In most midsize environments RAID5 or RAID5 with a hotspare is more than sufficient, but if you're looking at virtualizing highly transactional workloads, speeding up your storage will only help matters.

When inspecting the storage requirements, make sure to factor in data storage as well as virtual server system drive requirements. Unlike physical servers, it's not a good idea for virtual servers to be loaded with large local disks. Instead, it's much more beneficial to map LUNs to the virtual servers, or in a pinch, use iSCSI or NFS on the server itself to link to central storage arrays. Limiting the size of the local disk of each virtual server makes storage migrations much faster, provides much more flexibility, and reduces backup times.

You can also take advantage of storage tiering by placing server data LUNs on other, slower arrays and linking them to the servers either through virtual interfaces, or as raw devices, assuming that your hypervisor supports this method.

MIGRATING FROM PHYSICAL TO VIRTUAL

After all the meetings, decisions, plans, budgets, and purchases are complete, the hard work starts: migrating.

All kinds of tools are available to assist in physical-to-virtual migrations and most of them are sufficiently mature to handle most tasks. However, performing a physical-to-virtual migration should not be the first resort, it should be the last.

If at all possible, rebuild your servers in the virtual environment and migrate the data, applications, and configurations. While physical-to-virtual tools can certainly work wonders, they can also cause odd problems down the line, especially with Windows servers. It should go without saying that you shouldn't migrate Windows domain controllers. They're relatively simple to build and integrate into an Active Directory domain, and simple to demote.

Linux servers should also be rebuilt whenever possible, though physical-to-virtual migrations of Linux servers are generally more successful than Windows

migrations. Migrating applications and functions between Linux servers is generally far simpler than performing the same operations on Windows. In most cases, it's quicker to perform a fresh installation of Linux on a virtual server and then reinstall and reconfigure than it is to run a physical-to-virtual tool. If you're using templated virtual servers, it's even quicker.

Besides the obvious pitfalls of trying to graft an existing OS installation into a virtual server, you're bringing whatever problems may exist on that server into the new virtual server. Building clean and reinstalling will take more time in the short term, but may save huge amounts of time and effort later on.

That said, in some situations a physical-to-virtual migration is the only realistic option. In that case, be sure to run several test migrations first; don't gamble with production systems.

MONITORING AND MANAGEMENT

Your chosen virtualization platform will naturally come with management tools. Take advantage of any notification and alerting features and treat your virtual servers like you would physical servers. If you use server-side monitoring tools to keep tabs on server disk space, CPU load, and so forth, install them on the virtual servers, too.

It's important to realize that just because you can build a new virtual server with only a few clicks doesn't mean that you should. It also doesn't mean that the new server is "free." When planning for new projects that will require additional virtual servers, budget for those servers as physical servers. The additional load on the infrastructure, and the need to manage and back up those new servers, is just as costly as it is for physical boxes.

Keeping this in mind also reduces the propensity for virtualization implementations to grow wildly out of hand, with dozens of virtual servers floating around simply because they can be built in a few minutes.

BACKING UP VIRTUAL SERVERS

Lastly, there's the question of backups. Backing up virtual servers isn't anything like backing up physical servers. There's no point in installing backup agents on the virtual servers themselves. The best method is to invest in virtualization backup software that can be



automated to take a snapshot of the server and then back up the snapshot. This results in essentially no performance hit on the server itself.

The right tools and speedy hardware can back up a virtual server in a matter of minutes. Restoring the server is essentially as simple as bringing back the snapshot and booting it up. Not all virtualization solutions support this type of backup and restore, so be sure to weigh your options.

Virtualizing your server infrastructure may be the single best IT endeavor you undertake. Any

reasonably sized IT infrastructure should leverage virtualization in some form. The ease of administration, the rapid provisioning, the performance, the consolidation, and the power and cooling savings are all real, tangible benefits when you plan and deploy everything correctly. If you're still weighing the pros and cons of virtualization, it's time to get off the fence and start planning today.

Paul Venezia is an InfoWorld senior contributing editor and a veteran IT consultant.