

TECH TRACKER

For IT By IT

A Virtual Standstill?

Drive deep into server virtualization without considering its impact on your storage infrastructure and you could hit a painful bump in the road

By Michael Healey



IT'S NOT A MATTER OF IF you're going to virtualize your way out of server sprawl, but when. Thing is, if you don't bring your storage infrastructure up to par in tandem, it's akin to building a server superhighway that bottoms out on a dirt road.

Specifying a SAN as part of your server virtualization program is more than just a complementary add-on. It brings critical performance options that don't exist with other storage strategies, as well as opening the door to new space allocation and disaster recovery benefits built into virtualization platforms from VMware, Citrix, Virtual Iron, and others.

This may not be an easy sell—the server group is typically a separate entity from the storage and disaster-recovery teams, so it's easy to see how a company's virtualization and SAN initiatives can branch off in separate directions. And, in lean economic times, getting both new technologies funded is a challenge; we help with an ROI analysis (see story, p. 45). But this is one political and budget battle worth fighting.

NEED MY SPACE

Virtualization is a great way to increase hardware utilization, in terms of both CPU and memory. However, the same can't be said for virtualization's effect on con-

ventional storage systems. For example, VMware and other virtual machine players demand large initial space allocations on the storage array per virtual server, oftentimes more than for a physical server, to provide room to write the server state to the storage disk if the physical server on which the VM resides were to have problems. This tendency toward overallocation makes for both low storage utilization rates and higher costs, drastically reducing the benefits of a virtualized server deployment.

The answer is the intelligent, or thin, provisioning, embraced by SAN manufacturers including Hitachi, EqualLogic, Xiotech, and NetApp, that lets IT present a logical volume of any size to the virtual server, while committing capacity only when data is written. Say your database team swears it needs 500 GB for an application. You can provision the full 500 GB, but only actual usage will be committed, saving critical disk space. This even gives you the ability to overcommit beyond current capacity.

What happens when your overcommits catch up? You shouldn't get caught unaware because SAN manufacturers provide management software that sets thresholds, raises alarms, and can even automatically reprovision space as real utilization grows. However, it's not like overbooking a flight, then some sending folks home if too many show up—if your storage

grows to the full capacity allocated, you'll have to settle up on your promises. That may mean adding more storage to that SAN or grouping in another storage array.

WHAT IF ...

Virtualization has an odd way of waking organizations up to the challenges they face related to disaster recovery. A large municipality, for example, recently began a server virtualization project but became concerned about the risk associated with consolidating 10 to 15 servers to three or four physical boxes. It planned for additional failover equipment and added a SAN to better protect itself.

Funny thing is, prior to starting the virtualization initiative, the group had no spare hardware for any of its servers, nor did it have the ability to restore data anywhere on the network. This risk was even more critical because the municipality had older servers running key systems like payroll and billing. The situation was just easier to ig-

nore when the servers and data were spread across the network.

Virtualizing brought home the risk.

A SAN lets you address disaster-recovery plans effectively. Take the newest crop of servers from Dell, IBM, and Hewlett-Packard that sport embedded hypervisors, meaning the devices can boot off flash memory, then quickly become part of a server farm. The functionality is available only for VMware 3i. Time—and sales—will tell if this goodness spills over to virtualization software from XenSource, Microsoft, and earlier versions of VMware.

Whatever the method, VMs are all stored on the SAN, giving IT a higher level of RAID protection and automatic replication and host/application recovery.

And, if you don't virtualize using a SAN as part of your design, you can't take advantage of features specifically designed for centralized storage. For example, if you use VMware ESX without a SAN,

you give up live migration of virtual servers from one physical box to another, easy bare-metal restores, and dynamic sharing and allocation of server resources, like memory and network cards, as well as:

» Automated failover, the high-availability option within VMware that recognizes that a physical server is down. VMware knows the VMs that were running, and it will automatically restart those VMs on one of the other hosts within the resource pool.

» Consolidated backup, which lets a proxy server back up a snapshot of a virtual machine without significantly impacting its live operations.

» Access to upcoming features, like Site Recovery Manager, which can help automate a complete data center failover.

These same SAN benefits hold true for other vendors' disaster-recovery options, including those from Citrix XenSource. Deploying server virtualization and migrating physical servers to virtual machines



ROI Analysis: Virtualization

CIOs ARE ALWAYS SKEPTICAL of ROI-driven initiatives that don't have hard cost savings associated—productivity gains that never materialize can haunt you for years. Virtualization does have measurable savings in power, space, and cooling, but it's difficult to pinpoint an exact tipping point where hardware and software costs are offset. So we asked Brad Krick, manager of technical support configurations at GreenPages, to provide some cost analysis.

Krick's assumptions: A company must replace 10 aging servers and their associated data storage. We specified standard Hewlett-Packard servers and VMware ESX enterprise licensing. We didn't add in the common components of all three: tape backup, Microsoft licenses, or installation. You would probably see lower initial installation costs to convert 10 physical servers to virtual versus migrating apps to new physical boxes.

Storage requirement: 500 GB of Tier 1 data, 500 GB of Tier 2 storage for disk backup (near-line archiving).

Scenarios under consideration:

1. Replace with 10 traditional servers: Ten HP ML350 boxes with mirrored boot drives and RAID 5 storage: \$62,100
2. Virtualize 10 servers using three physical servers with local storage: Three HP ML350s with more memory and RAID 5 storage running VMware ESX (licenses included): \$49,300
3. Same virtualized servers in a cluster that boots to a 1-TB SAN: Three HP ML350s with more memory, boot drives, and VMware Enterprise, plus a Hitachi SAN for storage: \$52,730

Even companies leery of virtualization would be hard pressed to have IT spend \$12,000 more for a traditional setup. But can you get the CFO to go \$3,400 or so more for Option 3? You should be able to win this point—the SAN addition increases costs, yes, but it opens up options for redundancy, snapshots, and thin provisioning. Don't let role lock-in stop you from broadening your virtualization plans to include a long-term storage strategy. You'll not only better leverage the full range of virtualization features, you'll end up with a higher-performing environment. —MICHAEL HEALEY

creates the perfect opportunity to consolidate storage, centralize management, increase capacity utilization, improve availability, enhance data protection, and reduce backup windows.

THINK FAST

So what do you look for in a SAN to complement your virtual servers? Think wide striping across multiple drives to spread your data across numerous physical disks, resulting in faster response. The actual disk I/O is always the slowest part of system performance, and just like the disaster-recovery realization that occurs when you virtualize, the problem of disk I/O is more pronounced as you virtualize more servers. Try spinning up 20 or 30 database servers on one

physical host pointing to one disk array—you're going to hit a bottleneck at the disk, negating the benefits you got by virtualizing.

DIG DEEPER

EFFICIENCY EXPERT Managing storage resources well once your SAN is in place is vital to meet the needs of a virtualized network. Learn how in this *InformationWeek* Report: informationweek.com/1159/report_storage.htm

See all our Reports at informationweekreports.com

Its not just spindle count you need to consider. For intensive disk applications, you'll need to make sure your SAN includes multiple controllers for better throughput management, as well as more redundancy to avoid a single point of failure. Other factors include virtualization/SAN feature support, server memory, drive speed, and technology plat-

form—iSCSI vs. Fibre Channel.

Not all SANs are equal when it comes to working in the virtualized world. VMware, Microsoft, XenSource, and even Apple are working with SAN manufacturers to ensure that virtualization features are supported. Vendor certifications and formal partnerships are the norm, so make sure your storage and virtualization vendors are talking.

Disk performance can also become a factor when a server doesn't have enough memory to properly run its VM load, as the added reads and writes caused by swapping result in bottlenecks.

Our advice: Pony up for plenty of memory. If you don't have enough, you'll be relying on the disk subsystem's swap file to bring cached

THE LOWDOWN

» **THE PROMISE** Aligning server virtualization with an upgrade to your storage infrastructure, either by adding a storage area network or ensuring that your existing SAN is optimized to work with your virtual servers, is key to making the most of your investment.

» **THE PLAYERS** VM vendors including VMware, Citrix, and Virtual Iron have built space allocation, disaster recovery, and performance benefits into their virtualization platforms, while SAN vendors such as Hitachi, EqualLogic, Xiotech, and NetApp are making it easier link their gear to virtual networks.

» **THE PROSPECTS** Bringing server and storage teams together and convincing the CFO to fund a new SAN take effort. But they're worth doing, and vendors are building in functionality to make this an easier sell.

info from the disk back into memory. More trips to the disk equals slower performance. Also look hard at the speed of the drives themselves; this is definitely part of the performance equation.

The last question, choosing iSCSI or Fibre Channel, is one of those debates that's starting to get tiresome. Before getting into the fray, make sure you've addressed the design elements discussed above. They're more important. Then look at throughput versus interconnect speeds. A 1-Gbps Fibre Channel interconnect can deliver potential throughput of 96 MB per second, while a 1 Gbps iSCSI interconnect delivers 116 MBps. A 4-Gbps Fibre Channel SAN still delivers the highest throughput, but with 10-Gb iSCSI and 8-Gb Fibre Channel on the way, don't expect this debate to be over soon ... and don't get sucked in by it, either. What often gets skipped over in the yelling is the most critical factor in improving performance: The design of your storage array. The more drives you have, the more spindles there are to read and write data from, the faster your data access. Add the optimum speed drives, the correct number of controllers, and ensure that you've got ample memory in your servers and you'll be fine.

And you're not limited to one or the other. Many orgs are using varied technologies and SAN designs based on the performance needs of the virtualized applications.

"When we started three years ago, our game plan was always to tie it to some type of SAN," says Manny Silva, director of IT for Kuehne & Nagel North America, of the logistics company's virtualization initiative. "We already had Fibre Channel SANs in place but actually opted to go with iSCSI for the virtualization. In our testing the performance of both matched up."

The ability to mix and match your SAN environment lets you spend big on critical performance areas but go with a lower-cost option for disk backup or failover. Vendors like Hitachi, EMC, and EqualLogic have even designed "virtual storage" that lets you mix platforms on the back end while presenting a uniform front end for servers and applications. For sites that haven't yet deployed SANs, consider investing some of the virtualization savings realized in the server room. You'll be glad you did. —MICHAEL HEALEY

Michael Healey is CTO of GreenPages Technology Solutions; www.greenpages.com. Write to him at mhealey@greenpages.com.