PLANNING ASSUMPTION

June 16, 2003

Recentralizing Server Sprawl Through VMware: From Best Strategies to Cost Savings

Brad Day

Giga Position

Giga estimates that more than 30 percent of all midrange (up to four-way) and high-end **Intel** server (eight-way and beyond) symmetric multiprocessing (SMP)-based server platforms are being shipped for a variety of systems consolidation projects. Most IT organizations have realized that the continued growth of server population inherent in distributed computing environments can be extremely costly, resulting in the inefficient use of floor space (and with it increased cooling requirements) — both factors leading to increasing the ongoing operating costs associated with the proliferation of one-, two- and four-way SMP servers. Additionally, most applications software vendors did not encourage the deployment of different applications or multiple copies of the same application running on a single operating system (OS) instance. Consequently, IT departments previously tended to follow the unwritten technology rule of "one application per operating system instance," which resulted in servers being underused. Giga estimates that (as of the end of the 2002) more than 60 percent of the low-end and midrange Intel server population (e.g., four-and eight-way SMP-based systems) are running at only 15 percent to 30 percent of their total system capacities. More recently, however, through the use of **VMware**'s virtual machine GSX and ESX Server technology, collapsing the operating system images on a fewer number of servers has provided some economic benefits through operating system and application consolidation.

Recommendations

Giga advises companies to consider the choice for virtual machines and/or partitioning technology within any server platform environment (Unix, Windows, Linux or legacy systems) as the best fit technology solutions in the following four IT deployment scenarios:

1. Applications dependencies:
   - Different versions (or fix levels) of an operating system
   - Separate change management procedures for the operating system

2. Security requirements:
   - Different owners/administrators
   - Strong separation of sensitive data
   - Distributed applications with network firewalls between them

3. Recovery procedures:
   - High-availability (HA) clustering and application failover
   - Different disaster recovery procedures
   - When strong separation of performance is needed

4. If the performance characteristics must not interfere with each other:
   - When strong isolation of failures is needed
   - If the application or operating system failures must not interfere with each other

Within each of the major categories of reasons to deploy partitioned servers, a variety of customer

requirements have been addressed through the use, specifically of both VMware's GSX and ESX virtual machine systems software. Its use is advised in the following practical operational scenarios:

Providing the best approach for managing test and migration issues:

- Avoiding the need for dedicated, underused test servers
- Testing on the real production hardware, while concurrently running production workloads
- Allowing deployment of new operating systems and applications into existing infrastructure

Optimizing/reacting dynamic capacity demands:

- Allowing changeable partition capacity for applications with rapidly changing workloads
- Allowing resource redeployment for workloads that peak (or simply run) at different times
- Enabling rapid provisioning and deployment of new logical servers and applications
- Setting up "recovery/takeover" servers with excess server hardware capacity
- Providing either horizontal (clustered) or vertical (large single image) growth in a single server
- Growing systems capacities at a more predictable rate by smoothing out the bumps

Managing operational costs by the improved sharing of compute resources:

- Resource sharing is implicit, with the additionally ability to reassign resources as needed
- Consolidation onto a single power/packaging footprint
- Less inventory to constantly manage

ESX supports larger amounts of system RAM per virtual machine, as well as supporting a larger number of virtual machines per processor than GSX server. These two factors — coupled with the workload management functionality of ESX 1.5, which gives a system administrator the additional opportunity to set goals or metrics in terms of how the virtual machine consumes memory, CPU, disk and network resources — make ESX the preferred product over GSX when the customer objective is server workload consolidation.

Proof/Notes

From a high-level perspective, partitions are physical or logical mechanisms for isolating operational environments within single or multiple servers to offer the flexibility of dynamic resizing, while ensuring that applications are protected from disruptive unrelated events, such as interruption or performance degradation. When an application is isolated within a partition, it is able to (1) make dedicated use of server resources, (2) be managed separately or written with a group and (3) take advantage of unique requirements. Partitioning and resource management can be done statically or dynamically. Static partitions can be changed only after a systems reboot; dynamic partitions can be changed at runtime. Dynamic partitioning is a more sophisticated implementation of the operating system, middleware and applications so that software can properly deallocate resources before giving the resources to a new entity. From an introductory point of view, there are many approaches to increasing the number of applications running on a single server and better using systems resources, such as:

1. *No partitioning*: Clearly, running multiple applications without any partitioning or resource management is risky because of single-instance operating system failures. For example, an operating system failure due to resource conflicts in a consolidated applications environment can adversely impact the productivity of the entire user base. While an operating system crash on a server running a single office productivity application is costly and time consuming, if that same server is collectively running a procurement database management system (DBMS), Web site application and e-mail application, the impact of the systems outage has a more significant impact.
2. *Hardware partitioning*: In hard partitioning solutions, the hardware is actually electrically isolated so that the operating system has no indication that other resources are available. Hard partitions are designed to provide complete electrical and software isolation. This guarantees that any fault within one partition cannot impact any other partition. Applications running with hard partitions are not

subject to hardware or software events in other partitions. Hard partitions can usually support a user-defined partitioning of resources and isolate hardware and software errors.

3. *Logical (or sometimes referred to as virtual) partitions*: Logical partitions provide greater flexibility and granularity, while also providing greater fault isolation. Greater flexibility in logical partitions is achieved with the ability — using simple software commands — to add and delete CPUs dynamically from one virtual partition to another. In addition, the CPU granularity in which virtual partitions can be set up is one CPU or, in some cases, fractions of a single CPU.
4. *Resource (or sometimes referred to as workload) management*: Resource management allows for the ability to isolate specific system resources for use by an application or applications. It often assumes a single instance of an operating system.

While all three server workload consolidation options have their pros and cons, clearly the first option of running multiple applications without any type of partitioning or resource management solution should be avoided for the reasons already explained. An operating system failure due to resource conflicts in a consolidated environment can affect a great number of users and ultimately have an adverse effect on the productivity of an entire organization. When running multiple applications on a single instance of **Microsoft** Windows 2000, the potential for memory leaks within those applications increases, conceivably leading to operating system crashes.

With logical partitioning or virtual machine technology, the amount of fault tolerance and reliability depends on how well the software can truly isolate the operating system. The development of software to isolate multiple instances of the operating system within a server demands a large investment in engineering resources. Resource or workload management can bring many of the benefits associated with logical partitions. It can provide the stability needed to run multiple applications on a single server without the potential of additional hardware capital costs of a server using hardware partitioning. Resource management also reduces the complexity of a logical partition by managing only specific resources, such as memory and processors, allocating these resources to specific applications.

IT organizations considering any form of workload consolidation within a Wintel server environment on fewer larger SMP machines should understand the benefits of each workload consolidation technology before deciding which product, or combination of products, is best. Consider:

1. *Hardware partitioning*: Running multiple OS images along completely isolated hardware boundaries (this is currently not available on industry-standard Intel server platforms)
2. *Logical partitioning through virtual machines*: Running multiple OS instances where the guest operating system — in this case, Windows — owns its own applications, CPU, memory and I/O resources
3. *Dynamic resource or workload management*: Isolating specific resources (CPU and memory) to either single- or multi-instance application within a single OS instance

An important current restriction of VMware virtual server technologies (based on production-class customer reference site information) is that its virtual machines currently scale as a fractional portion of a single CPU workload. VMware has announced support for two- and four-way SMP configuration recently, but prior to interviewing customer betas, we would not recommend its deployment until the first half of 2004. Based on intelligence gained from Giga client inquiries, the following applications consolidation scenarios have been popular within VMware GSX and ESX Server deployments:

- **Citrix** Metaframe
- Lotus Notes
- Lotus Domino

- File and print services
- Select Java Application Servers (JAS) (predominantly based on older Java virtual machines (JVMs))
- Apache Web servers
- Microsoft IIS Web servers
- Active Directory
- Windows NT domain controllers
- DNS/DHCP/firewall servers
- In-house-developed applications
- SMTP-based servers (e.g., Exchange/Sendmail)

VMware GSX and ESX — Taxonomy and Product Comparison

On the Unix side, the partitioning technology solutions vary from hard or soft (virtual) partitioning, as well as the use of resource (workload management) partitioning. In all of these cases, the majority of the best-of-breed solutions are captive to the solutions provided by the same vendor that owns the design center for both the systems architecture and operating systems (e.g., **IBM**'s LPARs, **HP**'s nPARs and vPARs and **Sun**'s Dynamic System Domains). On the Wintel and Lintel server side, the only solution that has garnered any historical traction is VMware's virtual machine technology. As companies start to audit application portfolios of distributed server environments, IT operations managers are discovering that many distributed entry and midrange systems are typically running at ranges between 10 percent to 25 percent of full performance capacity.

In the larger acceptance of partitioning and/or virtual machines within Wintel installations, all of these functions are equally critical in a CTO's selection process. Historically, VMware's engineering roots started in the development of virtual machine software for workstations. VMware GSX should be considered as a solution for environments where running virtualized servers on a host operating system and where ease and speed of deployment is paramount. Subsequently, while VMware ESX Server 1.1 was well received, using virtual machines for a more cost-effective and efficient software test and development environment, VMware ESX Server 1.5 holds the most promise as the software virtualization technology to enable both applications and data-centric server consolidation projects (see Table 1).

Generally speaking, VMware's GSX solution is not as suitable for large-scale application server workload consolidation as VMware's ESX l.5 Server solution. ESX 1.5 is a solution chiefly designed to provide the opportunity to create multiple virtual machines per uni-processor in SMP systems, with the goal of consolidating multiple single applications running on multiple distributed servers onto a larger single server (ranging from one- to eight-processor systems).

The key benefits of ESX Server are multi-faceted, depending on the customer installed nuances of what is currently installed (e.g., hardware, OS and applications) and where the IT shop is focusing its recentralization efforts and human resources (often a part of a server consolidation planning exercise). Generally, however, a VMware ESX (vs. GSX) solution should be considered the best fit when:

- Workloads run within separate OS images, providing complete security isolation.
- Fatal workload errors can be contained within that workload.
- Workload consumption of processor, memory, disk I/O and network are all controllable.
- Low-priority workloads experience graceful degradation under congestion.

- Workloads with varying system DLLs or OS requirements can collocate.
- Workloads are portable among systems.

As noted in Table 1, ESX supports larger amounts of system RAM per virtual machine, as well as supporting a larger number of virtual machines per processor than GSX server. These two factors — coupled with the workload management functionality of ESX 1.5, which gives a system administrator the additional opportunity to set goals or metrics in terms of how the virtual machine consumes memory, CPU, disk and network resources — make ESX the preferred product over GSX when the customer objective is server workload consolidation.

Table 1: Comparing VMware's Server Solutions

Feature/Function	GSX Server 2.0	GSX Server 2.5	ESX Server 1.5
Infrastructure setting	Field office, department, test/development lab	Field office, department, test/development lab	Department, larger data center, test/development lab
Critical functionality	Shrink-wrap environment: quick installation, varied levels or hardware/OS platform and driver support	Shrink-wrap environment: quick installation, varied levels or hardware/OS	Advanced environment: higher levels of scalability, more dynamic control of hardware resources
Recommended system configuration	One to four CPU configurations	One to eight CPU configurations	Two to eight CPU configurations
Server hardware support	Compatibility is inherited from the host OS	Compatibility is inherited from the host OS	VMware certification is recommended
New device support	Compatibility is inherited from the host OS	Compatibility is inherited from the host OS	Drivers available through VMware
VM resource controls	Can preallocate memory, CPU and disk	Can preallocate memory, CPU and disk	Dynamic memory, CPU, disk and network
Skill level required	Minimal systems administration required	Minimal systems administration required	Higher systems-level administration skillset
Recommended number of VMs per processor	Four per CPU	Four per CPU	Eight per CPU
System RAM size	Up to 8GB	Up to 64GB	Up to 64GB

Source: Giga Research, a wholly owned subsidiary of Forrester Research, Inc.

While VMware's ESX Server 1.5 could be appropriate for some of the following cases, we strongly advise IT managers to alternatively consider GSX Server 2.0 and 2.5 more appropriate for the following customer scenarios:

- *Software development and test groups*: Where encapsulated, hardware-independent, instantly configured machine images let a systems administrator handle the proliferation of test environments more efficiently.
- *Departmental server administrators*: While ESX is the more optimal solution in larger application server consolidation projects, GSX can still be effective in increasing utilization and speed of applications deployment in departmental or replicated field site server consolidation projects.
- *Training center managers*: GSX increases classroom use and reduces setup times by letting the student work in isolated, centrally-hosted virtual machine environments.
- *Software demo center managers*: GSX allows for instant switching systems between multi-tiered server demonstrations as encapsulated in virtual machines and provides full remote access to those virtual machines.
- *Help desk managers*: GSX provides technical support staff with access to a server full of virtual machines that could be set up to duplicate every possible end-user (or customer) configuration.

The VMware ESX Server architecture — consisting of a VMware virtualization layer, resource manager and service console — focuses on four critical areas: (1) hardware independence, (2) fault isolation, (3) performance isolation and (4) encapsulation (see Table 2).

Table 2: VMware ESX Server Architecture — Features/Benefits

Feature	Function	Benefit Realized
Fault isolation	When potential errors or user action	It provides increased applications
Hardware	Each virtual machine presents to its	It delivers more reliable OS, in that the
Encapsulation	A virtual machine is contained within two	Files can be copied, duplicated and/or
Performance	Resource management allows for the	It allows administrators to define, allocate

workloads change.	

Source: Giga Research, a wholly owned subsidiary of Forrester Research, Inc., and VMware

As with any new solution from a software company, in-depth customer reference checks are a must. This is even more critical for system software, which on one hand can improve the overall efficiency of the computing infrastructure, but on the other might be initially perceived as being so complex and intrusive to install, deploy and manage that the functional benefits of the technology outweigh the business advantage. Because many business infrastructures evolved by just adding servers when new application requirements arise, the result over time is considerable complexity and a sizable population of multiple underused servers. Through consolidation of these resources, it makes financial sense to increase the use of the servers; however, it should not be done at the expense of either applications availability or performance. The optimal outcome would be to share the server resources while maintaining the application isolation required by any of the lines of business that IT supports.

Customer Benefits of VMware ESX

After Giga interviewed three VMware customers (a leading clothing retailer, a large financial services company and a leading application service provider), it was clear that the technical and business benefits outlined in Table 1, coupled with addressing the following customer selection criteria, resonated equally with everyone interviewed:

- *Improved manageability*: It is easier to manage a smaller number of centralized systems than it is to control a larger number of distributed systems. The manageability of systems directly impacts the cost and availability of the environment. One customer interviewed cited that cost savings specific to using virtual machines vs. individual and separate physical servers ranged from 20 percent to 70 percent.
- *Higher systems use*: Servers are typically sized for peak workload demands. As a result, when multiple systems are deployed, there is often extra capacity built into every server configuration. As discussed earlier, this can result in servers running at 15 percent to 25 percent capacity. One reference customer cited virtual machine use that increased capacity 70 percent to 85 percent, with operations management productivity gains ranging from 25 percent to 50 percent, predominantly stemming from improved systems administration.
- *Scalability*: As new applications/solutions are deployed, it is easier and faster to set up new partitions or virtual machines within a single CPU, rather than sizing an application solution for an additional fully configured server. Customer interviews suggest that human resources productivity gains alone weigh in at one hour vs. one day when installing and configuring a virtual machine vs. installing and configuring a new server.
- *Total cost savings*: When an organization combines multiple servers into one, it can achieve economies of scale by reducing floor space, power and cooling expenses. In addition, hardware and software licensing and maintenance costs go down by combining what were previously separate physical server environments. Finally, it takes fewer personnel to manage a smaller number of systems. Across a wider group of VMware customers (not all contacted by Giga), the company claims that customers estimate average cost savings ranging from 23 percent to 48 percent from hardware and systems cost reductions and 40 percent to 72 percent in operations cost reductions. Customer case scenarios suggest the new workload management features of ESX Server 1.5 are being adopted incrementally. Disk I/O resource controls are new in ESX Server 1.5 and are the last building block for achieving more complete resource partitioning among virtual machines.

References

Related Giga Research

Planning Assumptions

Data Center and Server Consolidation Better Than Ever, Brad Day and Colin Rankine

Server Workload Consolidation — Evaluation of Unix Systems Partitioning, Brad Day

Server Workload Consolidation, Part 2 — Evaluating Unix Workload Management, Brad Day

Server and Storage Infrastructure: Capacity On-Demand and Utility Computing Models, Adria Ferguson and Colin Rankine

IdeaBytes

AIX 5L Critical to IBM's Workload Management, Brad Day

HP's PRM and the Benefits of Workload Management, Brad Day

IBM's AIX 5L Road Map: What's Required to Maintain Customer Appeal? Brad Day

VMware's ESX Server: Virtual Machines Benefit Intel Server Consolidation, Brad Day

Solaris 9 Debuts, Part 1: What's Significant, What's Still a Work in Progress, Brad Day