An Introduction to the HP Virtual Server Environment

Chapter 1: HP's Adaptive Infrastructure

HP has been promoting an Adaptive Infrastructure for several years. The goal of this architecture is to synchronize business and IT in order to capitalize on change.

In today's business world, it seems that almost every business decision results in an IT event. These events can be anything from changing a product's price or running a marketing promotion, to implementing a large merger. Any of these might require:

- Modifications to applications
- Deployment of a new application
- A new instance of an existing application
- An increase of resources available to those applications
- A wholesale re-architecture of the data center

With this as the backdrop, it is critical that IT organizations become more agile – more able to respond to changes very quickly.

Virtualization in an Adaptive Infrastructure

Although everyone agrees that virtualization technologies are a key component of an adaptive infrastructure, HP takes a somewhat different view of virtualization than that of our competitors. The HP Virtual Server Environment (VSE) encompasses technologies that many people wouldn't consider to be strictly "virtualization technologies." HP takes the view that virtualization is more than just the ability to pool and share resources. Rather, HP provides software to optimize resource utilization and to apply those resources automatically to the workloads that need them. From this perspective, some non-traditional virtualization technologies, such as utility pricing, workload migration, and high availability, are included in HP's VSE. In addition, the HP Virtual Server Environment software provides an integrated set of tools to help plan, configure, and manage the VSE from a central workstation, (running either HP-UX or Windows), using HP Systems Insight Manager (HP SIM).

HP VSE now includes a breakthrough technology, called logical servers, which brings the flexibility of virtualization to physical Integrity blade servers. Logical servers are a new class of abstracted servers that allow administrators to manage both physical and virtual machines in the same way. This simplifies infrastructure change and increases business agility.

At HP, the most important goal is to make sure that the entire solution is integrated. It is not sufficient to just have products in each of these areas. Customers cannot take full advantage of such products if they do not work well together. HP has invested heavily in providing a cohesive solution that can be either implemented as components or integrated with each other.

VSE Design Principles

The VSE is a suite of HP virtualization products and solutions that run on HP 9000 and HP Integrity servers. As with the Adaptive Infrastructure, HP's VSE is focused on the solution it provides for customers. Some characteristics of VSE are:

- Workloads are encapsulated in isolated "compartments."
- Workloads are installed once and can run anywhere.

- Resources are automatically applied to workloads as they are needed according to business rules defined by the customer.
- If necessary, workloads are moved to provide better resource allocation or high availability.
- Groups of systems are managed as a single entity, from a workload service and physical perspective.

This requires a number of technologies. It requires the ability to run workloads in flexible partitions that can be dynamically sized based on the load of the applications running in them. It requires the ability to move workloads between systems transparently. It requires an integrated set of management tools that provide a view into the environment so that administrators can know where workloads are running, what resources they are using, what other workloads are running there, how to optimize and load balance the resources, and how to plan for future growth.

HP's VSE offers these technologies.

Chapter 2: Why Should an IT Organization Care About Virtualization?

Virtualization has become a hot topic in the last few years. In the late 1990s, many companies were growing very fast and it was very common for IT organizations to deploy separate servers for each application that was required by the business. The end result was a sea of servers in most companies' data centers – most of which were seriously underutilized. Let's explore how and why this overprovisioning occurs and how virtualization can help to significantly reduce computing costs.

Overprovisioning

Servers are overprovisioned for a number of reasons. Some of these include:

- Responsibility of application groups for their own IT resources
- Variations in application load
- Difficulties with capacity planning
- Severe penalty for underprovisioning

Responsibility of Application Groups for their Own IT Resources

This is probably the primary reason why servers are overprovisioned. It is common for each major application or business unit to have a separate application group that is responsible for its own IT resources. This situation leads to overprovisioned servers because each server must be sized to handle peak loads without regard to how much time the servers are underutilized. In addition, the resources are usually dedicated and cannot be shared.

Variations in Application Load

Most data centers have two common types of workloads: applications that users interact with each day, and the batch processing that is used to extract and process information from large amounts of data, which typically occurs during off hours.

Both of these workload types differ in their load profiles. Interactive workloads tend to vary in short windows of time. A workload might consume 15% of a system's resources at one time and then 75% a minute later. Batch processes typically have longer peak periods, but they tend to go from 0% to near 100% and then back to 0% as the job starts, runs to completion, and then finishes.

Because servers normally are sized for the peak load (and sometimes for additional growth), the average load is usually low in relation to the size of the server.

Difficulties with Capacity Planning

Another reason for overprovisioning is that capacity planning is an inexact science. Customers are attempting to predict the future using, in many cases, no real information about the resource requirements of the workload. Even when real data is available, customers speculate as to the amount of resources required on a new server, as well as expected growth of the workload over the life of the system. The result is that, just to be on the safe side, business units typically overestimate their requirements.

Severe Penalty for Underprovisioning

A major reason for overestimating resource requirements is that the impact of undersizing a server is severe. If a customer purchases a new system that has insufficient resources to meet the demand for the workload, customers must then:

- Purchase another, larger server, thereby increasing the cost of the project significantly.
- Wait for the new server to arrive.
- Initiate another installation, test, and migration project to move the workload to the new server.

An even more significant impact is that the organization that designed the original solution loses credibility, causing future recommendations to be met with skepticism, and resulting in the need to do more work to justify these recommendations.

Anatomy of a Workload

This section explores the two most common types of workloads – interactive and batch – in more detail.

User-Interactive Workloads

The load on user-interactive workloads is normally very spiky, as illustrated in Figure 1.

Figure 1. CPU usage over time

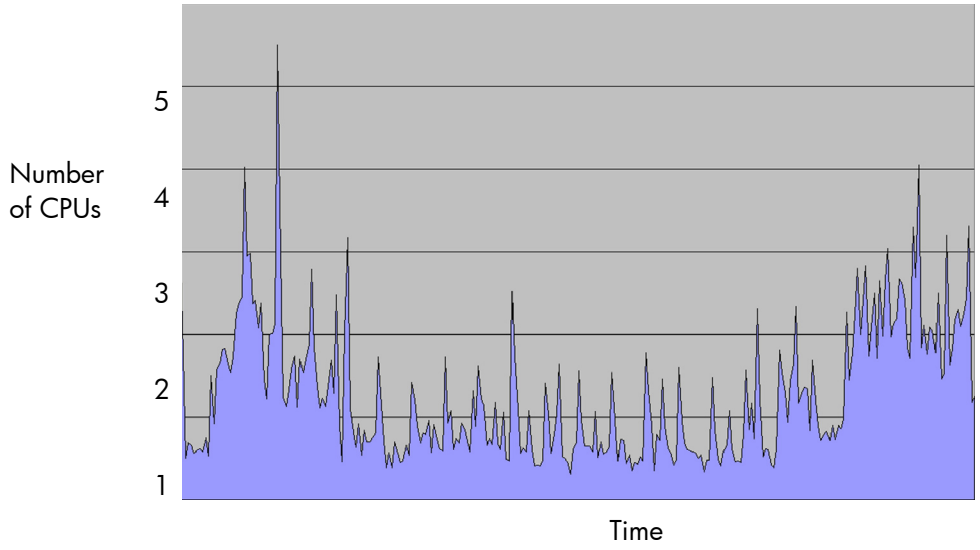

The graph in Figure 1 shows a real workload in the HP IT infrastructure and represents a web-based application on the www.hp.com portal. Because this is a revenue-generating workload, it is critical that it does not experience performance problems when it is under heavy load. As a result, the server that this workload runs on must be able to handle the peak load. If this workload were running on a fixed-size server or partition, it would need a minimum of 6 CPUs. However, the *average* load on this workload is only about 1.5 CPUs. This results in a 25% average utilization.

Batch Workloads

Batch-processing workloads have a very different resource profile. They tend to consume as much resource as is available for the period of time they are running, and then drop back to zero. These workloads also tend to result in a tremendous amount of unused resources because they require resources to be allocated for the peak, even if the resources are not used while the batch job is not running. The peaks are longer in duration, but the disparity between peak load and minimum load is enormous. This situation is made worse if the server is dedicated to the batch application and sits idle much of the time.

Other Costs Associated with Overprovisioning

Clearly, overprovisioning causes more money to be spent on server hardware, processors, or memory. However, overprovisioning also results in other costs that are not obvious at first glance.

Direct Costs

Customers might be spending a significantly greater amount than is necessary for software licenses. Most software licenses are based on the number of processors, so it stands to reason that if you can reduce the number of processors, you can reduce your software license costs. Having more servers or processors also results in higher maintenance and support costs, as well as additional floor space, power, and cooling requirements.

Indirect Costs

Having more servers means that you need additional resources to manage them. Having fewer servers and operating systems to manage allows you to free up those resources for more productive tasks.

Key Components of an Adaptive Infrastructure

The traditional approach for dealing with varying resource demands was to size the server for the peak demand. Sizing could then be optimized by controlling when certain workloads were running to limit the possibility of resource contention. For example, interactive workloads were commonly run on a system during the day, and batch processing was run during evening and night-time hours, when users were typically not using the system.

Virtualization allows you to look at this from a different perspective. Rather than manage the workloads to fit the resources, virtualization lets you manage the resources to fit the workloads.

Two key features are required to provide this different approach. The first is flexibility. Workloads must be put into compartments that are flexible so that utility resources can be applied to the workload when they are needed, and then are sent back to a free pool (or reallocated to another workload) when they are no longer needed. The second is automation. It is not realistic to expect administrators to recognize when resources are needed, identify where to take the available utility resources from, and then reallocate those resources before performance problems impact workloads – particularly with spiky interactive workloads.

Flexibility

Workloads must be placed into a compartment that can flex from the minimum fixed entitlement to a maximum entitlement by acquiring additional utility resources. As discussed in later sections,

the compartment can be anything from a server with utility components, to hardware- or software-based partitions that can flex to share resources or to take advantage of the underlying utility components.

Automation

It is not realistic to expect a human being to watch a screen such as that shown in Figure 1, see a spike in the workload, find idle utility resources available somewhere else on the system, deallocate those resources, and allocate them to this workload before the spike subsides. The only way utility resources can be applied optimally to workloads like this is to use automation. For VSE, this means workload management.

Utility Resources

The ultimate goal of a virtualized environment is to create a compute utility in which resources are used by workloads when they are needed, and then are released when not needed. Ideally, the customer pays for resources only when those resources are being used.

Each workload is typically set up with a minimum amount of resources that are always active. These are fixed resources. Additional resources can be applied to the workload when they are needed; these are utility resources. It does not matter where those resources come from; for this workload, they are consumed as needed and then released when not needed. This definition of utility resources includes the reallocation of idle resources from one partition to another.

Chapter 3: Virtualization Helps Reduce Overprovisioning

Virtualization can help reduce overprovisioning in many ways. It is no longer necessary to size a server (or workload) for peak capacity or for unexpected future growth. The size of the various types of partitions can be flexed with HP Instant Capacity (iCAP) or Temporary Instant Capacity (TiCAP) processors. These processors can be dynamically reallocated between Virtual Partitions or even shared with Integrity Virtual Machines (Integrity VM) or Secure Resource Partitions. New partitions can be added or existing ones permanently resized (if necessary) much more quickly than bringing a new server on line. The following sections describe these features in more detail.

Workloads Can Share Resources

The flexibility of virtualized environments means that resources that are not needed by one workload, can be used by others. The graph in Figure 1 shows that the length of time that this workload needs 6 CPUs is very short, and that most of the time, the peak sustained load can be handled by 3 CPUs. The other 3 CPUs can be put into a shared pool and can be used by whatever workload is peaking.

Even with this model, spare capacity is still needed. However, the amount of spare capacity required is much lower because it can be shared by multiple workloads. Engineers in HP analyzed workloads running on over 1300 servers and found that they were able to satisfy the peak loads of all workloads in a consolidated environment using 40% fewer CPUs. For example, if 10 workloads shared resources, as few as 36 CPUs, rather than 60, might satisfy their requirements.

Utility Resources Can Address Peaks

Many customers are uncomfortable running their environments without sufficient (guaranteed) spare resources. To address this issue, HP offers Utility Pricing Solutions that can provide for the additional capacity. In this way, in the unlikely case that several workloads become busy at the same time; the customer can activate additional capacity if idle capacity on the system is insufficient. The following paragraphs provide a brief overview of these HP solutions. Later sections discuss specific, related HP technologies.

The iCAP solution allows customers to purchase a server with more CPUs than they need. The iCAP CPUs are physically installed in the box but do not have usage rights; therefore, they are inactive and are not used by the OS running on the system. At a later date, they can be activated with full usage rights, and the bulk of the CPU cost is paid at activation time rather than at acquisition time. The TiCAP solution allows activation of the iCAP CPUs (without full usage rights) for short periods of time to handle peak loads.

Referring to the earlier 10-workload example, a customer can purchase a server with 64 CPUs in order to consolidate the 10 workloads, but 36 of those CPUs can be permanent capacity and the other 28 can be iCAP. This configuration costs much less than if all 64 CPUs were purchased permanently. If you add to this some TiCAP licenses, you now have the ability to activate as many as 64 CPUs if all the workloads are busy at the same time. Although the situation where all workloads peak at the same time is unlikely, it does provide a measure of safety that ensures that any combination of spikes can be handled by the server. In addition, the cost is reduced because the customer pays for only the CPUs that are actually turned on.

Flexibility Simplifies Capacity Planning

The flexibility inherent in these HP solutions drastically reduces the risk of underprovisioning. If one of the workloads turns out to be larger than expected, or if it grows faster than expected, the system has enough spare resources to satisfy the application almost instantly. You no longer need to purchase a larger server and migrate the workload; you simply activate some additional resources and assign them to the workload.

Virtualization Reduces Deployment Time

Another advantage of virtualization is that new workloads or applications can be deployed very quickly on the existing spare capacity. Customers have been able to reduce the time it takes to roll out new services by as much as a factor of 10.

The current process of acquiring a new server looks something like this:

1. Perform capacity planning to determine the size of the server.
2. Get funding approval for the new server.
3. Order the server.
4. Wait for the server to be delivered.
5. Install the server, the OS, any support software, and the applications that will run.
6. Test the new server/software stack to ensure that it will function correctly in the production environment.
7. Roll the server into production.

A virtualized environment still requires most of these steps, but many of the most time-consuming ones either take much less time or are unnecessary. For instance:

1. Capacity planning is still required, but it can be less exacting and, therefore, less time consuming.
2. Funding approval takes less time because the resources are already available. Little or no additional outlay of cash is needed to add the new workload.
3. You do not need to order a new server because you can use existing server resources. This requirement is replaced by the creation a new virtual environment to run the application – a process that typically can be completed in minutes.
4. You do not need to wait for the new server to arrive. This is the most significant time savings in the process.
5. You still need to install the OS, any support software, and the applications that will run in the virtualized environment, but you do not need to find space in the data center for a new server and bring it on line. Also, if the new application can run on an existing server or partition that already has the software installed, this process can be very short.
6. You still need to test the new server/software stack to ensure it functions correctly in the production environment.
7. You still need to roll the new service into production.

Because new virtualized servers can be deployed quickly, used for a short project, and then "destroyed" or "recycled," they can also be useful in development or test environments. Resources can be used for these short projects and then released for use in a subsequent project. It is also much easier to test a number of different OS, application, and patch versions with the same set of resources.

Chapter 4: HP Partitioning Solutions

This chapter discusses the partitioning solutions available for HP Integrity servers, with a particular emphasis on how each of these solutions can be made dynamic.

nPartitions

HP nPartitions are implemented in hardware on HP cell-based servers. Each nPartition consists of one or more cells, each with a physical group of processors, memory, and I/O cards. This is implemented via a state-of-the-art chipset that allows any combination of cells to be electrically isolated from cells in other partitions. Figure 2 shows how the cells are connected to crossbars (labeled "xbars") on the backplane of a Superdome server using the sx2000 chipset. Note that all the connections between the cells and the crossbars are triple redundant for high availability. Each of the links can also be programmatically configured to drop electrical signals if the source and destination cells are not in the same partition.

HP nPartitions can flex dynamically in 2 ways:

1) Using Dynamic nPartitions. This allows you to add or remove a cell board (and its associated CPU and memory, and any attached I/O) to or from an nPartition without shutdown or reboot.

2) Using iCAP (described in the next chapter). This feature allows the number of active CPUs in each partition to be dynamically adjusted while the partitions remain running.

Figure 2. Connectivity among cells and crossbars (sx2000 chipset)

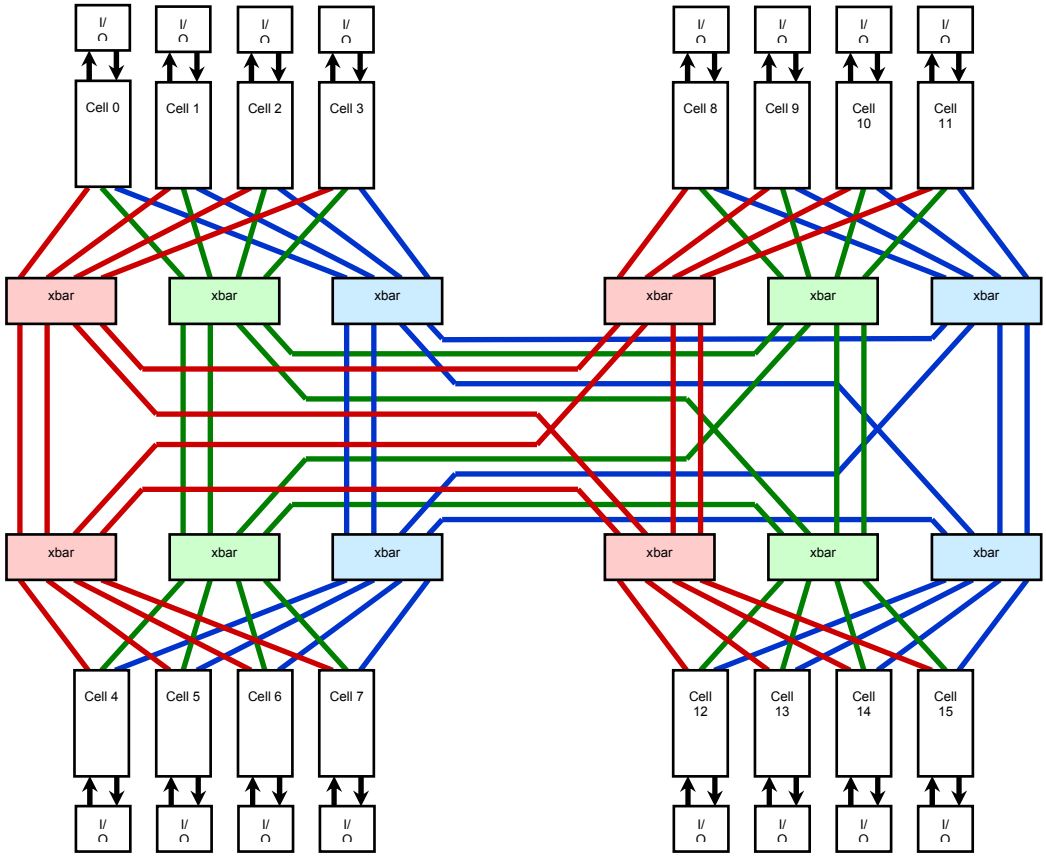

Virtual Partitions

Virtual Partitions allow multiple instances of HP-UX to run on a single system or nPartition. This capability provides finer granularity than nPartitions through the dynamic reallocation of CPUs and memory between Virtual Partitions. Virtual Partitions are implemented by assigning dedicated hardware components such as CPU, memory, and I/O to each partition.

Figure 3 illustrates how Virtual Partitions work. Rather than boot HP-UX directly off the hardware, the vPar Monitor is booted. The vPar Monitor is configured so that it can identify all the partitions as well as the physical components that belong to each partition. The vPar Monitor loads each Virtual Partition, which then boots its OS. When the OS queries the kernel to determine what hardware components are available, the kernel responds with only the subset of the resources that are allocated to that partition. An advantage to this model is that the Virtual Partition is able to communicate directly with the hardware resources, resulting in a very low overhead.

Virtual Partitions can adjust the amount of CPUs and memory allocated to each partition while they are running. This allows CPUs and memory to be moved between Virtual Partitions by deallocating them from one Virtual Partition and subsequently allocating those same CPUs to another one.

Figure 3. HP-UX Virtual Partitions

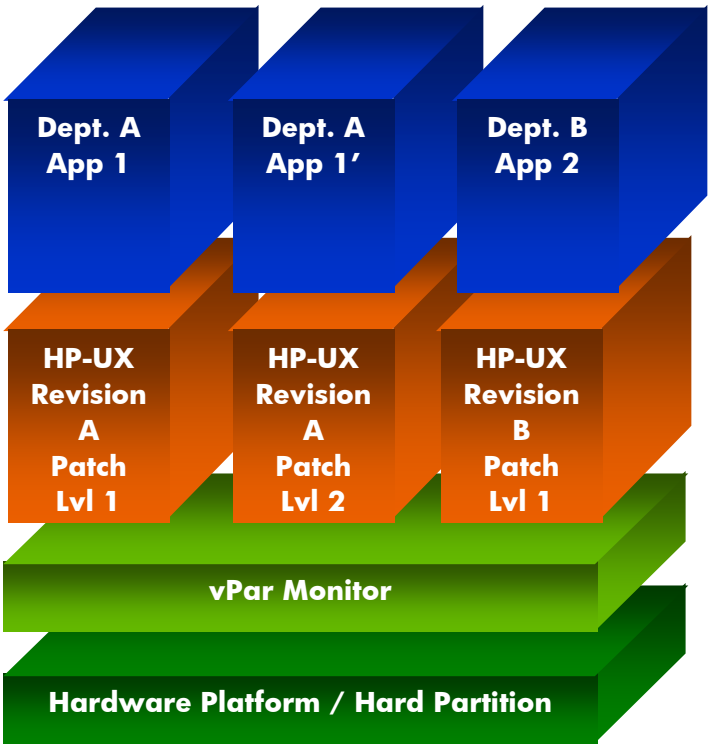

Integrity Virtual Machines

Another solution in HP's partitioning continuum is called HP Integrity Virtual Machines. This is a fully virtualized environment for running applications. Figure 4 illustrates how Integrity Virtual Machines are architected. The VM Host runs on any HP Integrity system or nPartition, which means that you do not need a cell-based system to use Integrity Virtual Machines. On top of the VM Host you run virtual machines, which present themselves to the VM Host as physical servers. When an OS is installed on a virtual machine, the virtual machine becomes a guest. In an Integrity Virtual Machines environment, all of the resources in a guest are virtualized. The physical CPUs, memory, and I/O devices are managed by the VM Host. What the OS inside the virtual machine sees is a virtual resource that is mapped on top of the physical devices in the system. This allows the physical resources to be shared by multiple guests.

The virtualization that is provided by Integrity Virtual Machines is so complete that the operating systems running inside the virtual machines run without modification. This means that the same Integrity version of each OS (HP-UX, Windows, and Linux) that runs on a standalone system or nPartition will run on a virtual machine.

Because of this level of isolation, software running inside a virtual machine cannot determine that it is not running on a physical system. All of the resources that are presented are virtual, and the virtual machine might be sharing those resources with multiple other virtual machines. As a result, the resource allocation to virtual machines is very dynamic by default, but is generally transparent to applications running inside the virtual machine.

Integrity Virtual Machines utilize Accelerated Virtual I/O (AVIO) drivers for networks and/or storage. These VM-aware drivers streamline the I/O path and related performance of Integrity Virtual Machines.

HP also provides a capability called Integrity Online VM Migration. Integrity Online VM Migration enables an Integrity virtual machine, its guest O/S and its applications to be dynamically moved from one VM Host to another without service disruption. No shutdown or reboot is required. All I/O connections to networks and storage remain active throughout the migration. This capability is particularly helpful for proactive maintenance and workload balancing.

Figure 4. HP Integrity Virtual Machines

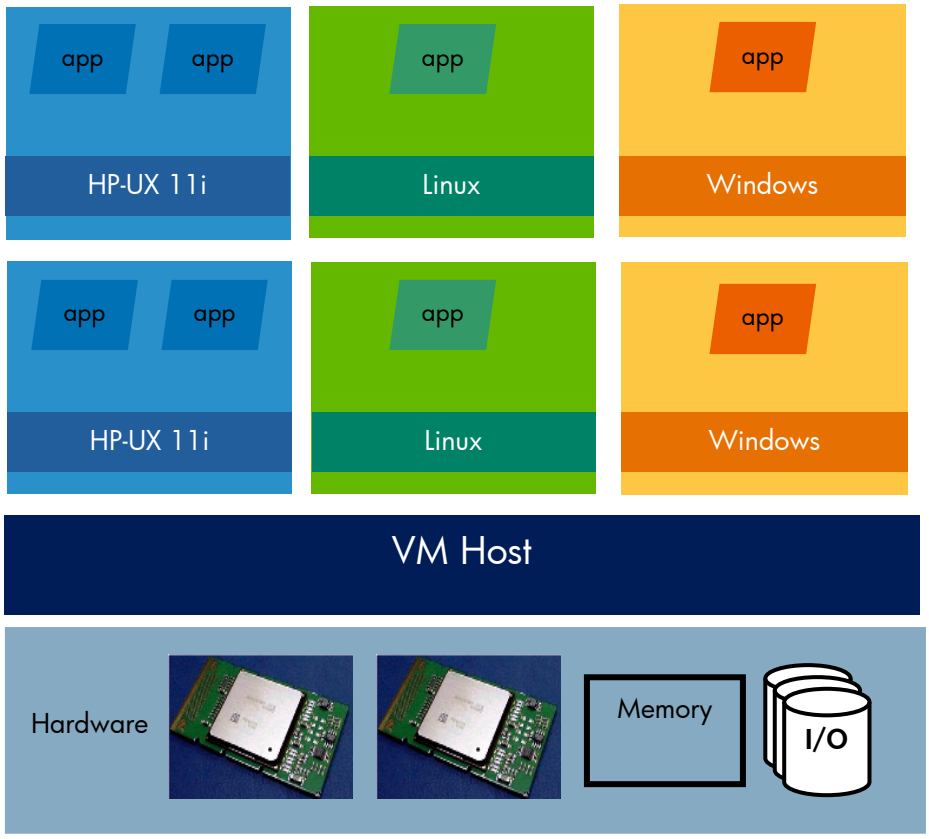

Secure Resource Partitions

Secure Resource Partitions is a set of technologies that allow you to run multiple separate instances of applications within a single instance of HP-UX while maintaining resource and security isolation between the instances. Some of these technologies have existed in HP-UX for over 10 years. The capabilities of this combination, in conjunction with the Process Resource Manager product (PRM), Workload Manager (WLM), or Global Workload Manager (gWLM), is usually referred to as "application stacking."

The underlying technologies include:

- Fair Share Scheduler – for sub-CPU granularity CPU control
- Processor sets – for whole-CPU granularity CPU control
- Memory Resource Groups – for control of real and shared memory
- LVM and VxVM – for control of disk I/O bandwidth within each volume group
- Security Compartments – for security isolation between groups of processes
- PRM Application Manager – to manage how processes are assigned to groups for each of the technologies listed here

Figure 5. Secure Resource Partitions

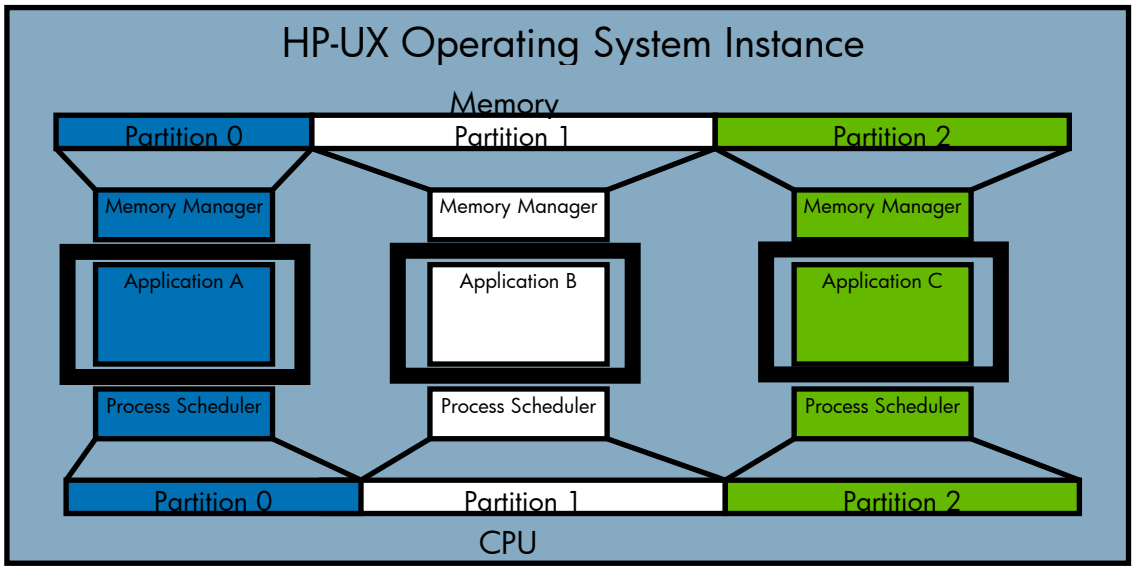

Chapter 5: VSE Utility Pricing Solutions

Utility Pricing Solutions can be applied to any cell-based HP server. The server does not have to be partitioned. However, the Utility Pricing Solutions are integrated with partitioning, so it is possible to use both if needed.

This section briefly describes the Utility Pricing Solutions available as part of the VSE suite of technologies. The focus here is in describing how these solutions can be used as part of the flexibility requirements for Adaptive Infrastructure.

Instant Capacity

Instant Capacity (iCAP) enables you to purchase and install additional processing power through the use of a two-step purchase model. Initially, you purchase system components (processors, cell boards, or memory) at a fraction of the regular purchase price because the usage rights are not included. These iCAP components are inactive, but installed and ready for use. When extra capacity is needed, you pay the remainder of the regular purchase price for the usage rights to activate the component or components. This greatly simplifies capacity planning and helps avoid overprovisioning for unexpected capacity growth. If and when the additional capacity is needed, it can be activated dynamically.

Also, and perhaps even more important than the two-step purchase model, is the fact that the iCAP software keeps track of only the number of usage rights you purchase for a particular server; it does not track which physical components you purchase without usage rights. This feature allows for great flexibility because the usage rights can essentially flow from partition to partition as needed, depending on capacity demands. When one partition needs more capacity, cores can be deactivated in other partitions in order to release usage rights which can be used to activate cores in the partition that needs extra capacity. The only requirement is that the number of activated cores across the entire server does not exceed the number of purchased usage rights or purchased temporary capacity. This means that iCAP can be used to flex nPartitions.

Figure 6 shows how this flexibility works. The diagram shows a cell-based system configured as two dual-cell nPartitions. Each nPartition has a total of 8 physical (single-core) processors, but the complex has 4 iCAP processors. This means that 12 cores on the server are active and 4 are inactive. For iCAP, it does not matter which partition contains the inactive cores. In the diagram, the partitions can each have 6 active cores, or nPar1 can have 8 cores while the nPar2 has 4, or nPar2 can have 8 cores, while nPar1 has 4. As long as the number of active cores does not exceed 12 across all the partitions, these reconfigurations can be done in any combination. This allows each partition to contain from 4 cores to 8 cores, in single-core increments, while both partitions are up and running and under load, and while maintaining full electrical isolation between the partitions.

Figure 6. Flexing nPartitions with iCAP

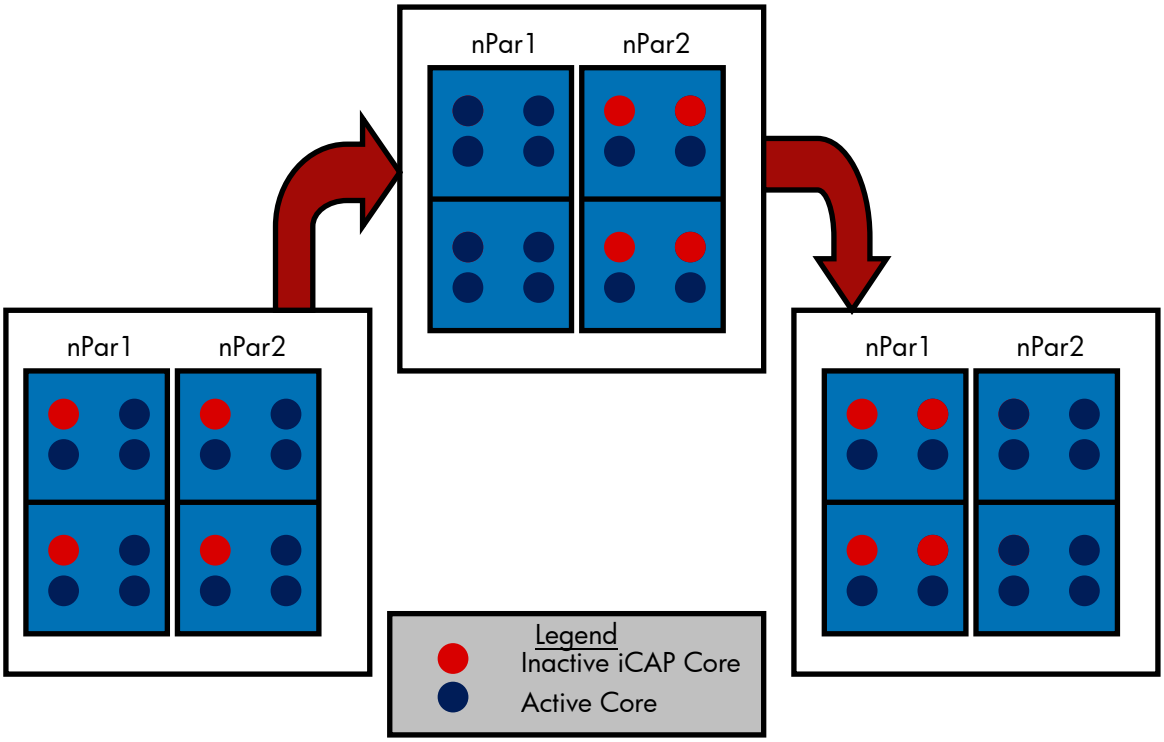

Temporary Instant Capacity

Temporary Instant Capacity (TiCAP) gives you the option of purchasing a prepaid block of processing time that enables you to activate and deactivate iCAP cores without having to purchase the usage rights for permanent activation. The process involves purchasing a block of processing time (in 30-day increments) that enables iCAP cores to be activated for a limited period of time (measured and tracked in 30-minute intervals). In fact, you can use this block of processing time to activate more than one iCAP core at the same time, thereby enabling you to activate sufficient capacity in response to short-term spikes in processing requirements.

Global Instant Capacity

With Global Instant Capacity (GiCAP), the concept of shared usage rights extends to a group of servers. Cores can be deactivated on one server in the group in order to facilitate activation on another server. Again, the main requirement is that the number of activated cores across the GiCAP group does not exceed either the number of purchased usage rights for the group or the amount of purchased temporary capacity. This effectively extends the iCAP and TiCAP capabilities to partitions that are running on separate physical servers that can be located in separate cities or even separate countries.

Pay-per-use

Pay-per-use (PPU) differs from iCAP in that it is a leasing option. When you lease a system from HP, you can choose a flexible lease payment that is based on the actual CPU utilization of the server each month. The utilization is measured every 5 minutes for each core on each partition of the server. This data is then averaged each day across all the cores on the system (or complex). This daily average is

used to calculate a monthly average, which is the value used for the PPU lease-payment calculation. One feature of PPU is that it gives you the option of deactivating individual cores to ensure that they contribute a net zero utilization to the average utilization calculation for as long as they are deactivated. This option can also be automated using either HP-UX Workload Manager or HP-UX Global Workload Manager (gWLM), which activate the cores when they are needed and then deactivate them when they become idle.

Chapter 6: VSE Automation Solutions

Two tools in the VSE suite are designed to help automate the migration of resources between workloads that share an infrastructure. HP-UX Workload Manager (WLM) was the industry's first goal-based workload manager product for any Enterprise UNIX platform. The more recent HP Global Workload Manager (gWLM) is an evolution of WLM and was specifically designed for the "IT as a service provider" market.

HP Global Workload Manager

The gWLM product was designed to take advantage of the experience gathered from customers who used the WLM product. It was also designed for a new business model that is becoming increasingly prevalent in IT organizations: IT as a service provider. Some of the key architectural differences between WLM and gWLM include:

- The gWLM product is managed from a single Central Management Server (CMS), making it easier to manage larger numbers of servers and workloads in a virtualized environment.
- With gWLM, a single policy can be applied to more than one workload and on more than one server.
- A primary goal of gWLM is ease of use. Most IT organizations do not have extensive knowledge of the applications in their environments, so gWLM was developed with intuitive-use models targeted for IT users rather than application owners.
- The gWLM product offers a number of standardized reports, both historical and real time. These reports provide the service provider with information to manage the infrastructure, as well as document how resources are being shared and consumed by client business units.
- As with HP-UX WLM, gWLM supports the allocation and reallocation of iCAP, PPU, and TiCAP resources, but in addition, is also integrated with GiCAP.

HP-UX Workload Manager

The HP-UX WLM product provides the ability to monitor workloads that run on an HP-UX platform and to reallocate resources as the loads vary over time. It supports workloads running in nPartitions, Virtual Partitions, and Secure Resource Partitions. With WLM, resource allocation can be continually adjusted in response to demand, and is governed by a set of rules based on business priorities. It also supports the allocation and reallocation of iCAP, PPU, and TiCAP resources among the workloads on the server.

HP-UX WLM has an easy-to-use configuration wizard and a remote graphical user interface to simplify the monitoring and reconfiguration of WLM instances.

Choosing Between gWLM and WLM

HP's strategic workload management product is gWLM. HP gWLM is ideal for both a large-scale deployment with multiple servers, as well as for workloads on a single system or in a high-availability environment. A typical example is a shared IT deployment where a single IT department manages servers for multiple business units and many applications run a large number of servers – each with several partitions. HP gWLM's ease-of-use features include centralized policy administration, a library of policies, and reporting features that enable IT to easily set up, manage, and track resource usage.

These benefits also make gWLM a suitable solution for many common types of IT environments where fine-tuned policies and specialized resource management are not applicable.

HP-UX Workload Manager is designed to manage workloads on a single system or in a high availability cluster. HP-UX WLM is a suitable solution for a line-of-business (LOB) consolidation in which the LOB owns its servers but relies on an IT team to manage them. WLM also gives you the ability to fine-tune policies to support specialized resource management needs.

HP-UX WLM is intended for installations where it is already in use. HP gWLM should be considered for new installations as well as installations where growth, in terms of scale or complexity, are anticipated.

Table 1. Differences between WLM and gWLM

WLM	gWLM
Line of Business (LOB) project	Centralized IT project
Few servers	Many servers
No CMS required	CMS required
Custom policies for each workload	Standardized policies across multiple workloads
HP-UX 11i servers only	HP-UX 11i on OpenVMS (Windows and Linux within Integrity VM)
In-depth application integration	Optimal HP Integrity VM management

Chapter 7: Combining Virtualization Solutions

A unique feature of HP VSE is that the different partitioning solutions can be combined on a single system. This provides the user with the flexibility to consolidate varied workloads with different isolation and flexibility requirements onto the same platform comfortably.

You can use nPartitions, Virtual Partitions, Integrity Virtual Machines, and Secure Resource Partitions on the same system at the same time. The key advantage is that you can decide which is the right technology for any particular workload and then put it on the system using the most appropriate partitioning technology. You can get the advantages of hardware fault isolation *and* the flexibility of Virtual Partitions or Integrity Virtual Machines, as well as the application-stacking capabilities of Secure Resource Partitions without having separate systems to run them. In addition, you can use HP utility pricing solutions, such as iCAP, to augment the existing resources. A further advantage is that the flexibility of all these technologies can be combined, and workload management tools can provide the ability to automatically move resources wherever they are needed; from Integrity Virtual Machines in one nPartition over to Virtual Partitions in another nPartition, as required.

Combining Partitioning Solutions

The following combinations of partitions are supported:

- nPartitions – You can run Virtual Partitions, Integrity Virtual Machines, or Secure Resource Partitions inside of an nPartition. This provides the benefit of the hardware fault isolation between the nPartitions and the granularity of any of the other partitioning options. Note, however, that Integrity Virtual Machines and Virtual Partitions cannot run in the same nPartition.
- Virtual Partitions – You can run Secure Resource Partitions inside of Virtual Partitions for application stacking, which may be running inside an nPartition.
- Integrity Virtual Machines – You can run Secure Resource Partitions inside of an HP-UX virtual machine guest, which may be running on a standalone server or within an nPartition. However, because Integrity Virtual Machines already provides sub-CPU granularity and shared I/O, using Secure Resource Partitions within an HP-UX virtual machine guest is not likely to be commonly used.

Adding iCAP to Partitioning

Instant Capacity is supported with all of the partitioning solutions, but it is supported only on cell-based systems. The only other thing to note is that iCAP needs to be managed on the VM Host for Integrity Virtual Machines. This is because the CPUs inside a VM guest are virtual CPUs, and iCAP operates only on physical CPUs. So the mechanism for applying iCAP to a VM guest is to activate the CPU on the host and then apply the newly activated CPU cycles to the guest by increasing its entitlement.

Chapter 8: Managing the Virtual Server Environment

HP is not the only vendor with partitioning solutions or instant capacity solutions. However, HP does offer more choices in these areas, allowing customers to use the right technology to do the job. What really differentiates HP partitioning solutions and sets them above our competitors are HP's capabilities to automate, plan, and manage the virtualized environment.

HP Virtualization Manager

As a plug-in to HP Systems Insight Manager (HP SIM), the Virtualization Manager software provides a central point of control for managing all of the resources in your VSE. Virtualization Manager leverages functionality within HP SIM to provide auto-discovery and role based security for more efficient system administration. Virtualization Manager has an easy-to-use interface that lets you see all of your available virtual resources, how they're being used, and how they relate to your physical infrastructure in a contextual graphical view. It is seamlessly integrated with the management tools necessary to manage both your physical and virtual infrastructure for the HP Integrity, BladeSystem, and ProLiant platforms.

Virtualization Manager is also integrated with Capacity Advisor and gWLM to help you plan and automate resource sharing in your VSE. And as of January, 2009 Virtualization Manager is integrated with the new technology called logical servers, which we will explain later in this chapter.

HP Capacity Advisor

Capacity Advisor is an easy-to-use capacity-planning tool. It captures performance data and then allows you to graphically view the historical utilization data. Unlike many other capacity planning tools, Capacity Advisor captures more than just CPU and memory utilization data. Capacity Advisor also captures networking I/O, disk I/O, and power consumption data. It collects more than a thousand data points per server per day for up to 4 years, which improves the quality of analysis and allows you to see and understand trends over time.

For example, the graph in Figure 1 was taken from Capacity Advisor. By using the information in this graph, you can easily determine that only 3 CPUs are required to handle the sustained peak load the majority of the time. The remainder of the CPU resources can be shared by either moving CPU resources from another partition or by activating TiCAP cores. If you have several different workloads, you can create scenarios for combining them into the same nPartition by using Virtual Partitions or Integrity Virtual Machines. In this way, you can determine how many resources are required based on the historical performance data captured by Capacity Advisor. This allows you to size for the "peak of the sums" and not the "sum of the peaks." That is, when you combine the workloads, you look for the total CPU demand of the combined workloads. Since the workloads probably won't peak at the same time, you size for the peak of the combined workloads.

Another significant benefit of Capacity Advisor is that you can use it to create "what if" scenarios that allow you to match workloads with various configurations in order to choose the one that works the best in your environment. In the Scenario Editor, you can manually create hypothetical servers and move workloads onto them. An innovative 5-star rating system will help you determine the optimum workload placement. Or, you can use the built in "Smart Solver" technology to automatically consolidate and load balance your workloads to servers or virtual machines. Comparison reports can be generated to compare the results of different consolidation scenarios, including power

consumption and the associated costs. (You can specify the type of currency and cost per kilowatt hour).

Logical Servers

Logical servers are a breakthrough technology that bring the freedom and flexibility of virtualization to physical servers. They are a new class of abstracted servers that allow administrators to manage both physical and virtual machines in the same way. A logical server is defined by a server profile that is easily created and freely moved across physical or virtual machines. Logical servers bring many of the attributes of virtual servers to physical servers in terms of flexibility, and the ease of creating, moving, and retiring them.

Logical servers can either be blade logical servers that utilize HP's Virtual Connect hardware, or they can be VM logical servers. Logical servers can be created by simply importing an existing blade server or VM, or you can create a new one by using an easy to use wizard that guides you through the process. Each logical server that you create or import will include the following types of information in the logical server profile:

- Identity information, including a unique name, description, platform type, architecture (e.g. ProLiant, ESX Virtual Machine, or Integrity), and operating system.
- Compute resources, including number of CPUs (or virtual CPUs for VMs), CPU frequency, and amount of memory.
- Storage configuration, including server WWNs and storage WWNs for blade logical servers or the shared storage information for the VM hypervisor for VM logical servers.
- Network connectivity information, including MAC addresses.

Currently, ProLiant blades with VMware ESX, ESXi, or Microsoft Hyper-V and Integrity blades (running HP-UX, Windows, or Linux) are supported. Support for HP Integrity Virtual machines is planned for early 2010.

Logical servers can be activated on a server blade or as a VM guest. A 5-star rating system is used to let you select the best choice based on performance and headroom from a list of target locations. Logical servers can be created and saved to be used as "templates", which will allow for rapid deployment of a new server. And if the operating system and application software had previously been provisioned, the logical server could be up and running in the amount of time it would take to boot the operating system. Logical servers can also be deactivated or moved very quickly to accommodate the changing requirements in the data center. For example, on ProLiant, an ESX VM logical server can be moved live from one host to another if Vmotion is configured. When a blade logical server is moved or deactivated, the server is shutdown and it's virtual connect profile is unassigned so that the blade can be used for another purpose. However, because all of the necessary information about the logical server is saved in the logical server profile, and because it boots from SAN, it can be activated quickly onto another blade, even one in a different enclosure as long as it's in the same Virtual Connect Domain.

Some of the use cases for logical servers are:

- Reduce Time to Deploy New Servers: Suppose the application group asks the I.T. department to provision a new server for a particular application. They specify how many CPUs, how much memory, storage, etc on the "work order" or maybe these specifications have been previously defined. So, I.T. could either create a new logical server with these characteristics or use a previously defined logical server template, or copy an existing logical server or template. Then, this logical server could be activated, and let the Logical Server Management software in conjunction with the Capacity Advisor software find the best target location using the 5-star rating system.

- Conveniently expand and contract application server pools: Suppose a company has an application consisting of several servers or nodes in a cluster configuration whose load varies over the course of a business cycle. By using logical servers, the operations group can "dial-up or down" more or less blade servers to meet the changing demand.
- Re-use server resources for another purpose: A company may have multiple Q/A or test systems for different applications. If they are not all needed at the same time, logical servers can be associated with each one of them. When one is no longer needed, it can be deactivated, which will free up the blade for use by another logical server. All of the specific configuration information associated with this logical server is maintained in the logical server profile, including it's MAC address and WWNs. So, it can be re-activated at a later time. Another logical server can now be activated to re-use that server blade.
- Move the Logical Server to a different location for maintenance purposes: Another case is that a blade needs to be taken down to upgrade the firmware or something like that. Just move the Logical Server to a different blade, perhaps even one in a different enclosure as long as it's in the same Virtual Connect Domain. (Of course, the operating system does have to shut down and be rebooted).
- Move the Logical Server for load balancing purposes: VM logical servers can easily be moved from one hypervisor to another in order to balance the load between them. The Capacity Advisor software has a load balancing feature that can make recommendations for the best way to place the VMs. There is even a drag and drop capability within Virtualization Manager which can be used to move them, and the 5-star rating system will show you the best target locations.

HP Global Workload Manager

Global Workload Manager, described briefly in an earlier section, is a part of the HP Virtual Server Environment software suite. The gWLM software is designed to manage a large number of servers and workloads in a virtualized environment. To accomplish this, gWLM runs on a CMS and is integrated with HP SIM and Virtualization Manager. From the CMS, you can centrally manage the servers, partitions, and workloads that gWLM manages. You can apply policies to workloads that govern both when and how resources are shared. Also, you can get reports, both real-time and historical, about how the resources were allocated, how many resources each workload used, and how many each workload needed. The gWLM software operates in "advisory mode" as well as in "managed mode." In advisory mode, rather than automatically move CPU resources, gWLM informs you that it would have done so. One of the design goals for gWLM was to allow users who are already using Virtual Partitions to understand, install, and configure gWLM to manage those Virtual Partitions in less than 30 minutes.

Chapter 9: Start the Journey

Getting to a Virtualized Environment and a truly Adaptive Infrastructure is indeed a journey. It is a journey well worth taking, and you can get there. However, the undertaking is not an all-or-nothing effort. You do not have to have a completely defined virtualized architecture in place in order to start taking advantage of the many benefits offered by HP VSE. There are many non-disruptive actions that you can do to start benefiting from virtualization technologies. The greatest advantage of HP virtualization technologies is that they have all been designed to work well together. You can pick and choose which technologies are right for your business needs and be confident that the combinations will succeed. In addition, you can implement new technologies later, when the time is right, and know that they, too, will work together.

Summarizing the Choices

If you want to consolidate servers, you might want to consider the larger, cell-based servers, which afford many options for partitioning. Initially, these servers might cost more than a one-server replacement strategy, but when you consider the resource sharing potential as well as the ability to quickly create new partitions and deploy new instances of applications, that strategy might ultimately be more cost effective. You might also achieve indirect cost savings as a result of reduced floor space, power, and cooling requirements.

Partitioning solutions are great for isolating applications for security or availability concerns. These solutions are useful if applications need a different version of the OS, different kernel tunables or performance tuning, or different versions of the application software. They also allow you to "right size" and move cores between partitions for handling peak loads. Additionally, creating a new partition on existing resources instead on provisioning a new server is much easier, quicker, and more cost-effective.

Application stacking with Secure Resource Partitions might be an excellent choice for applications that work well together and that can coexist on the same version of the OS. This can increase server utilization and, in many cases, can save a significant amount of money in software licenses and support costs.

With HP virtualization technologies, you do not need to size for unexpected growth or for peak processing. Instead, you can use iCAP or TiCAP for these eventualities. If you need a high-availability or disaster-tolerant solution, consider using Serviceguard.

Last but not least, you can automate resource sharing on your systems with WLM or gWLM.

Getting Started

Pick a pilot project and get started. Once you successfully implement it, you can add new functionality if you need to, or you can start a totally new project. HP has designed all its virtualization technologies to be integrated and to work well together.

For more information

For more information about HP VSE and HP virtualization technologies, see the VSE website at www.hp.com/go/vse.

4AA1-4391ENW Rev. 4, March 2009