

# Implementing a Virtual Server Environment:

## Getting Started



Chapter 1: Introduction .....	3
Chapter 2: Assessing and Planning.....	4
Long-term roadmap .....	4
Short-term roadmap.....	4
Chapter 3: Understanding the Choices for Virtualization Technologies.....	6
Partitioning Solutions: .....	6
Why choose nPartitions (nPars)? .....	6
Key Benefits.....	6
Trade-offs .....	6
Sweet Spots.....	6
Why choose Virtual Partitions (vPars)? .....	6
Key Benefits.....	6
Trade-offs .....	7
Sweet Spots.....	7
Why choose Integrity Virtual Machines (VMs)? .....	7
Key Benefits.....	7
Trade-offs .....	7
Sweet Spots.....	8
Why choose Resource Partitions or Secure Resource Partitions (SRPs)?.....	8
Key Benefits.....	8
Trade-offs .....	8
Sweet Spots.....	8
HP Utility Pricing Solutions:.....	9
Why choose Instant Capacity (iCAP)? .....	9
Key Benefits.....	9
Trade-offs .....	9
Sweet Spots.....	9
Why choose Temporary Instant Capacity (TiCAP)?.....	9
Key Benefits.....	9
Trade-offs .....	9
Sweet Spots.....	9
Why choose Global Instant Capacity (GiCAP).....	9
Key Benefits.....	9

Trade-offs .....	10
Sweet Spots.....	10
<b>Automation Solutions:</b> .....	10
Why choose Workload Manager (WLM) or Global Workload Manager (gWLM)?.....	10
Key Benefits.....	10
Trade-offs .....	10
Sweet Spots.....	10
<b>Chapter 4: HP VSE Reference Architectures .....</b>	<b>11</b>
<b>Chapter 5: Identifying a Pilot Project .....</b>	<b>12</b>
General guidelines.....	12
<b>Chapter 6: Making Your Choices .....</b>	<b>13</b>
Determine CPU processing requirements for workloads.....	13
Capacity Advisor .....	14
Choose the type of partitioning.....	14
Choose the type of Integrity server .....	15
Utility Pricing .....	16
Organizing partitions and resource sharing .....	16
Automation .....	17
High Availability and Disaster Tolerance.....	17
Implement the Pilot Project.....	18
Repeat the Process .....	18
<b>For more information.....</b>	<b>19</b>

# Chapter 1: Introduction

Many customers have said that they like the concept of the HP Virtual Server Environment (VSE), but that it seems complex and disruptive to their environment, and they don't know where to begin. Achieving a virtualized environment and a truly Adaptive Infrastructure is a journey, and not an all-or-nothing goal that must be reached in the first implementation. You do not even need a completely defined VSE architecture in place in order to start taking advantage of the many benefits offered by HP VSE. You can start benefiting from virtualization technologies without introducing disruptive and time-consuming implementations. One great thing about HP's 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 – both now and in the future – and be confident that they will work together.

---

**Note:**

This paper is intended for readers who are familiar with the HP virtualization technologies that comprise the HP Virtual Server Environment. References to more information, including a white paper entitled, "An Introduction to the HP Virtual Server Environment," can be found on the last page of this paper.

---

## Chapter 2: Assessing and Planning

Every organization is unique and has different challenges and problems to solve, but all must use a similar process. This chapter is intended to help you organize your thoughts and plans.

### Long-term roadmap

The first step is to identify your organization's business drivers and long-term goals. For most customers, the ultimate goal is to create an agile and responsive IT organization, one that aligns with the needs of the business and, therefore, is both a service to the business and a competitive advantage. For this reason, it is important to understand the key business drivers and to think about what you want your environment to look like in the long term. The following are some pertinent questions and considerations:

- Analyze your application environment and some of your business processes as they relate to the IT organization. Do you have more than one line of business, and does each line of business or application group own their own servers for testing and production? Application "silos" such as this are a leading cause of underutilized servers and server sprawl. One long-term goal might be to consolidate these applications and allow IT to provision the necessary resources as they are needed. This approach can allow the application groups to focus on meeting the needs of the businesses they support instead of worrying about the IT infrastructure.
- Are you expecting significant growth in the business or growth through acquisitions or mergers? If so, is it important to be able to deploy new applications or new instances of existing applications quickly? Some HP VSE customers find that a virtualized environment allows them to deploy a new application in days rather than months.
- Do you need to achieve a 24x7 environment or add a disaster-recovery capability? HP virtualization technologies are well integrated with the HP Serviceguard suite of high-availability and disaster-tolerant solutions.
- Do you have to support and maintain many different versions of the operating system and application software? Does your organization already have certain standardized processes, or is each application group completely independent? If you consolidate and share resources, what kinds of internal or "political" issues will you need to address? These issues might not be trivial to solve, but it is important to understand them. You can address many problems and pain points using the virtualization technologies within these "application silos," but in the long run, customers who can also solve the business process issues will benefit the most from a Virtual Serve Environment and can achieve a more robust Adaptive Infrastructure.

### Short-term roadmap

As previously described, defining a long-term architecture for a virtualized environment requires significant thought and planning, and perhaps even some significant changes in your IT infrastructure and business processes. The good news is that the VSE architecture does not have to be completely defined before you can take advantage of HP virtualization technologies. You may need to address some pain points and solve some short term problems first.

It is clear that most companies need to reduce their infrastructure costs. The problems of overprovisioning, underutilized servers, server sprawl, paying too much for software licenses and support, increasing cost of power and cooling, and simply running out of space in the data center are

common. Which of these problems apply to your organization, and which ones need to be addressed first?

For most customers, any change in the IT infrastructure or application environment poses some level of complexity or risk. This paper discusses the various ways that you can implement changes incrementally, from those that are least disruptive to your current environment to those that represent the greatest change and therefore provide the greatest benefit.

# Chapter 3: Understanding the Choices for Virtualization Technologies

This chapter describes the key benefits, trade-offs, and sweet spots for some of the HP virtualization technologies. Understanding these will help you determine which technology is most appropriate for solving a specific problem or achieving a certain benefit. Use this information as you decide on a pilot project and assess the level of complexity and the amount of change that may be required to implement them.

---

**Note:**

Most of the information in this chapter is taken from a book entitled *The HP Virtual Server Environment*. For information about how to obtain this book, see the last page of this paper.

---

## Partitioning Solutions:

### Why choose nPartitions (nPars)?

**Key Benefits**

- Hardware fault isolation (electrical).
- Operating system isolation.
- Choice of OS (HP-UX, Linux, Windows®, OpenVMS).
- No negative performance impact (in some cases you might see improved performance resulting from less SMP overhead).
- Easy implementation.
- Dynamic cell OLAR (Online Addition and Removal) with HP-UX 11i v3.

**Trade-offs**

- Requires cell-based system.
- Granularity for a partition is at the cell level.
- No resource sharing across nPars (unless being flexed with Instant Capacity cores).

**Sweet Spots**

- Mission-critical applications that require fault isolation and dedicated resources.
- Need to run multiple operating systems on the same physical server.
- nPars are supported on both PA-RISC and HP Integrity servers.
- nPars can be a mix of PA-RISC and HP Integrity on Superdome servers (excellent for mixed environments or during transitions).

### Why choose Virtual Partitions (vPars)?

**Key Benefits**

- Dedicated hardware resources (processor, memory, I/O).
- Partition size can be scaled in increments of 1 processor core.

- Operating system isolation (each vPar is a unique instance of HP-UX).
- CPU and memory resources can be shared or moved dynamically between vPars that are within the same nPar.
- Good choice for I/O-intensive applications as compared with HP Integrity Virtual Machines.

#### **Trade-offs**

- No hardware fault isolation for vPars running in the same nPar (that is, a hardware failure within an nPar will affect all vPars in that nPar).
- HP-UX is the only operating system supported.
- Only supported on cell-based systems if using Integrity servers (some older, non-cell-based systems are supported on PA-RISC).
- Requires dedicated hardware resources, which might result in overprovisioning for small workloads.

#### **Sweet Spots**

- A good choice if you require finer granularity than an nPar but still need dedicated hardware, have I/O-intensive applications, or need a unique instance or version of HP-UX.
- A good choice if you want to dynamically share processor or memory resources between vPars (within the same nPar).
- Instant Capacity or Temporary Instant Capacity resources can be activated for any vPar within the same nPar.
- Easy for deploying a new application or a new instance of an existing application by simply creating a new vPar (instead of deploying a new server).

## Why choose Integrity Virtual Machines (VMs)?

#### **Key Benefits**

- Granularity is sub-CPU (as little as 5%).
- CPU entitlements and memory can be changed dynamically.
- Dedicated hardware not required; CPU and I/O resources are shared.
- Supported on all Integrity systems (running HP-UX 11i v2 or later) and on both cell-based and non-cell-based servers.
- Operating system isolation and flexibility (that is, each guest OS is a unique instance).
- Multiple OS guests are supported (HP-UX, Windows, and Linux; OpenVMS is planned).
- OS guests can run without modification.
- Any virtual storage device (disk, CD, DVD) can be implemented as a file.
- An ISO image of a preferred software image can be implemented as a virtual DVD, and can be used to quickly deploy common software, operating system updates, or patch bundles to multiple virtual machines.

#### **Trade-offs**

- No hardware fault isolation.
- Integrity VMs and vPars cannot be used within the same nPar.
- Each Integrity VM currently limited to 4 cores.
- Hardware resources are shared, so not a good choice if dedicated hardware is required.

- Not supported on PA-RISC systems.
- There is a slight decrease in performance for I/O, so not the best choice for I/O-intensive applications.

#### **Sweet Spots**

- Good choice for applications that do not need dedicated hardware (or an entire CPU) but do need OS isolation, different OS versions, different OS types, or a unique version of the application stack.
- Good choice for non-cell-based systems that need a partitioning solution (if they are not I/O intensive).
- Applications with spiky workloads can often get more than their entitlement of CPU cycles if the other virtual machines are not demanding those cycles.
- Easy for deploying a new application or a new instance of an existing application by creating a new virtual machine.

## Why choose Resource Partitions or Secure Resource Partitions (SRPs)?

#### **Key Benefits**

- HP Process Resource Manager (PRM) product can be used to manage system resources (CPU, memory, and disk I/O bandwidth) according to a user-defined priority by placing processes in processor sets (PSETs) or Fair Share Scheduler (fss) groups.
- The granularity of resource allocation for PSETs is at the whole-CPU or core level.
- The granularity for resource allocation when using fss groups is sub-CPU (as little as 1%)
- Does not require a separate instance of the OS, as do vPars or Integrity VMs.
- Memory and I/O can be shared; memory entitlements can be reallocated on line.
- Supported on both HP 9000 and HP Integrity server systems; runs on both cell-based and non-cell-based systems.
- Can save significant amount of money on software licenses when application stacking by reducing the number of OS instances required.
- Workload Manager (WLM) can be used to add goal-based workload management and automation of iCAP resource usage.
- By using the Security Containment feature of HP-UX, you can place one or more secure compartments in PRM groups to create a Secure Resource Partition (SRP). Processes in each SRP are isolated and cannot communicate with or access the resources of processes in other SRPs.

#### **Trade-offs**

- No hardware isolation (same as vPars and Integrity VMs).
- No OS isolation or flexibility because Resource Partitions or SRPs are in the same OS.
- Supported only on HP-UX.
- Requires same patch levels and kernel tunables because all partitions in the same OS.

#### **Sweet Spots**

- Excellent choice for application stacking. Can run multiple instances of the same application on the same OS while maintaining application isolation and resource-level guarantees.
- Allows for resource-level control and application isolation without requiring a separate instance of the OS.

## HP Utility Pricing Solutions:

### Why choose Instant Capacity (iCAP)?

#### Key Benefits

- Allows you to defer the cost of system components (processors, memory, cell boards) until you need the capacity.
- Simplifies capacity planning; reduces the need to overprovision.
- nPars can be flexed with iCAP cores on the same server or across servers with Global Instant Capacity (GiCAP).
- Cell-board iCAP allows you to activate a complete cell. Starting with HP-UX 11i v3, you can dynamically add a cell to an nPar without rebooting the OS.

#### Trade-offs

- Available only on cell-based systems.
- To activate the usage rights of an iCAP core for a partition, the core must physically be in the same nPar.
- Prior to HP-UX 11i v3, rebooting is required to activate an iCAP cell board.

#### Sweet Spots

- Provides a cushion or safety net for capacity planning.
- Provides inexpensive spare capacity that can be activated online.
- Allows you to flex the size of partitions in a failover situation.
- Provides the ability to flex nPars.

### Why choose Temporary Instant Capacity (TiCAP)?

#### Key Benefits

- Allows iCAP cores to be activated for limited periods of time.
- When purchasing TiCAP (in 30-day increments), you can activate one or more iCAP cores against the balance (measured and tracked in 30-minute intervals).

#### Trade-offs

- Available only on cell-based systems.

#### Sweet Spots

- Excellent choice for responding to short-term spikes in the workload; preferable to use with WLM or gWLM.
- Cost-efficient way to provide high-capacity test environments for short periods of time.
- Allows you to lower the cost of a failover server because the cores need to be activated only during the failover period.

### Why choose Global Instant Capacity (GiCAP)

#### Key Benefits

- Allows iCAP resources to be shared across multiple servers (usage rights are deactivated on one server and activated on another).
- TiCAP resources can also be pooled and shared across multiple servers.

**Trade-offs**

- Available only on cell-based systems.

**Sweet Spots**

- Can use iCAP resources for load balancing across servers.
- Can be useful in failover situations. (To flex the resources of the failover server, GiCAP usage rights can be transferred to the failover server so that TiCAP resources don't need to be consumed.)

**Automation Solutions:****Why choose Workload Manager (WLM) or Global Workload Manager (gWLM)?****Key Benefits**

- Allows control over how shared resources are allocated between workloads.
- Automates the reallocation of resources between partitions that can share resources.
- Automates and controls the use of utility pricing to manage costs.
- Automates the activation and deactivation of TiCAP resources, ensuring they are active only when the load requires them (resulting in lower costs).
- Allows application performance to remain at consistent levels as load varies.
- A single gWLM policy can be applied to many servers.

**Trade-offs**

- Both WLM and gWLM support HP-UX. In addition, gWLM supports Windows, Linux, and OpenVMS if they are running as a guest OS on an HP Integrity Virtual Machine.
- WLM supports Resource Partitions and Secure Resource Partitions; gWLM does not.
- WLM must be configured on each server.
- WLM does not support Integrity VM.
- gWLM can manage multiple servers from a single management server.

**Sweet Spots**

- Automates sharing of processors between vPars.
- Automates flexing of nPars with iCAP processors.
- Minimizes cost of TiCAP solutions, and automatically deactivates TiCAP processors when they are not required.
- Automates reallocation of resources in a failover situation.

## Chapter 4: HP VSE Reference Architectures

HP VSE Reference Architectures (RAs) are documented best practices for solutions based on the VSE components and key industry applications. The VSE RAs may provide blueprints or guidelines for doing just what you are intending to do. HP VSE RAs are based on proven, real-world IT deployments, and might help reduce the time it takes you to implement similar solutions. They can also help you optimize your design time by providing examples of proven designs that you can apply or customize to fit your specific requirements. Some of the VSE RAs that currently exist are:

- Databases: Oracle® and Oracle RAC
- Enterprise resource planning applications: SAP R/3 and mySAP Business Suite
- Application Servers: BEA Weblogic Server, Oracle Application Server, and IBM Websphere
- Business Intelligence: SAP Enterprise BI Server

---

**Note:**

For information about how to access the HP VSE RAs, see the last page of this paper.

---

## Chapter 5: Identifying a Pilot Project

As stated earlier, moving to a virtualized environment is a journey. It's completely up to you whether to start with a project that is highly visible and very important to your business, or with a smaller project just to get oriented to a virtualized environment. HP has worked with customers who have done both. Starting small allows you to always add new projects and new functionality later because these technologies are designed to work well together.

Do you want to address some of your short term problems first? Is there some problem that can be solved very simply and quickly? Do you want to replace legacy hardware or implement a completely new project? Do you want to create an application service or utility so that new applications can be deployed more quickly in the future?

Regardless of what project you pick, here are a just a few things to consider:

- What version of the operating system and application software will you need? If you are moving to a new hardware platform or a new version of the OS, HP recommends that you do a complete software-stack assessment to ensure that all of the software you need is available and supported.
- It is wise to get buy-in from your application software providers about your new strategy or architecture.
- Be sure that you understand the performance requirements for sizing partitions or for application stacking. (The next chapter addresses this point.)
- Be sure to obtain management support within your own organization or from the line of business that owns the application.

### General guidelines

Here are some general guidelines and thoughts to keep in mind regarding the technology choices that you must make for your virtualized environment.

As an alternative to having a separate server for each application, you can divide a larger server into partitions. Partitioning solutions are great for isolating applications that have security or availability concerns. They are also useful if applications need a different version of the OS, different kernel tunables, or different versions of the application software. Finally, creating a new partition on existing resources is much easier, quicker, and cost effective than provisioning a new server. Partitioning is a relatively low-risk option because it allows you to maintain application isolation even though the application may not be on its own physical server. In addition, if there is a concern about availability due to multiple partitions on one server, nPars or the HP Serviceguard product might be a good solution.

Avoid sizing for unexpected growth or for peak processing. Instead, use iCAP or TiCAP. Partitioning solutions also allow you "right size" and move cores between partitions for handling peak loads. Although this is not risky, the solution can become slightly more complex when using TiCAP, since you will most likely want to add automation with WLM or gWLM to manage its usage effectively.

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

## Chapter 6: Making Your Choices

This chapter provides an ordered set of steps that you can follow to help select your particular implementation of HP virtualization technologies. The following steps consider the possibilities incrementally, from those that are least disruptive to your current environment to those that represent the greatest change and, therefore, the greatest benefit.

### Determine CPU processing requirements for workloads

Before you decide on the type or types of partitioning solutions you think you want, you must first determine the size of the Integrity server (or partition) you need for each application – just as though you were doing one-for-one server replacement. Specifically, determine how many CPU cores are needed on the new Integrity server to provide the performance equivalent to that of your older server. This exercise is no different than any other server upgrade. Also determine how much memory and what type of I/O resources are needed (for example, what kind of I/O cards and how many of each type). In addition, identify the I/O characteristics in order to determine later in the process whether Integrity VMs are a viable option.

At this point in the process, do not concern yourself with resource sharing. This is an additional step in refining the solution and will be considered later. If possible, determine the base level of resource usage for “normal” processing for your environment (whether that’s 50% or 80% utilization, or whatever number is right for your organization). Be sure you understand the resources necessary for peak load processing. Finally, factor in a growth forecast for 3 to 5 years.

Table 1 and Table 2 provide worksheets that can help you organize your resource usage and workload sizing information, respectively. Fill in the existing resource usage information for each server or application workload in Table 1. After you determine what type of Integrity processors you will be using, fill in the information in Table 2 for each server or application workload.

---

Table 1. Sizing worksheet: resource usage information

Legacy server or application workload	Number of CPU cores	Type and speed of processor	Memory required	Disk I/O rates	Networking I/O rates

---

Table 2. Sizing worksheet: workload information

Legacy server or workload	Number of Integrity CPU cores	Memory required
Normal Processing		
Peak Processing		
Expected Growth		
Total for Workload		

## Capacity Advisor

One tool that you can use for gathering performance utilization data is Capacity Advisor, a management tool provided in the HP VSE suite of products. You can use Capacity Advisor to collect performance utilization data from your existing servers and then to create what-if scenarios to determine the best way to consolidate those servers. Capacity Advisor allows you to test configurations with the size characteristics that you define, and to determine whether you can combine various workloads onto this new server (existing or hypothetical) before you make the changes or purchase the new server. It even allows you to factor in forecasted growth percentages for each workload or server.

In order to use Capacity Advisor, however, you must first install the VSE Management Software on a Central Management Server (CMS). This step might seem illogical because you are in the process of defining your virtualized architecture. However, it might be worthwhile to install this software on a CMS and on any managed nodes that you plan to consolidate into your virtualized infrastructure so that you can use the capabilities of Capacity Advisor in planning your new configuration. Once installed, you can collect performance utilization data over a period of time by using the HP Utilization WBEM Provider collection agent for each of the managed nodes that you choose. You can also import OpenView Performance Agent (OVPA) data if that data is available. After your new virtualized environment is set up, Capacity Advisor is valuable for monitoring and evaluating your workloads to make the most of the available systems, and for evaluating the possible effects of moving workloads around for optimal utilization. For more information about using Capacity Advisor, see the link at the end of this paper to the *HP Integrity Essentials Capacity Advisor User's Guide Version A.03.00.00*.

## Choose the type of partitioning

The partitioning solutions are the basic building blocks or foundation of the Virtual Server Environment. Choose one of the following partitioning options for each application or workload.

Table 3. Selection criteria for choosing type of partitioning

Type of partition	Selection criteria for each application or workload.
nPartitions	If hardware or electrical isolation required. If work load needs dedicated hw resources .

Virtual Partitions	<p>If hardware or electrical isolation not required .</p> <p>If work load needs dedicated hw resources.</p> <p>If workload is not suitable for a VM: (e.g. Needs dedicated hardware resources, is I/O-intensive, requires more than 4 cores, or can take advantage of more than 4 cores).</p>
Integrity Virtual Machines	<p>If hardware or electrical isolation not required.</p> <p>If work load does not need dedicated hw resources.</p> <p>If workload requirements are often sub-CPU (smaller than 1 core).</p> <p>If workload runs with 4 cores or fewer.</p> <p>If workload not I/O intensive.</p>
Resource Partitions or Secure Resource Partitions	<p>If hw/electrical isolation not required.</p> <p>If applications can run on same instance of HP-UX, with same patch levels and same kernel tunables.</p>

---

**Note:**

If you are seriously considering Integrity Virtual Machines, see the very useful white paper entitled, "Hardware Consolidation with Integrity Virtual Machines." It can help you determine which workloads are good candidates for Integrity VMs. It also provides recommendations for assessing the performance of your current workloads (both CPU and I/O), and can help you size the target hardware for the VM Host. See the link at the end of this paper.

---

## Choose the type of Integrity server

There are many different choices for Integrity servers, ranging from low-end (non-cell based) to high-end (cell based). Now that you understand the hardware resource requirements of each workload, and you know which workloads are candidates for particular types of partitioning solutions, you can decide what type of Integrity servers you need. The information from the preceding worksheets will help you later in the process when you determine how to combine or group partitions on a server. Additionally, it will help you determine whether or not resources can be shared effectively between those partitions or certain application workloads.

If you need nPars or vPars, you must move to a cell-based server. If you think you will want to take advantage of iCAP or TiCAP resources, then you must also choose a cell-based server. If you only need Integrity Virtual Machines or want to do application stacking with Secure Resource Partitions, you do not need a cell-based server.

If you want to consolidate applications from a lot of servers, you might want to consider the larger cell-based servers, which will give you all of the options for partitioning. It might cost more initially than a one-for-one server-replacement strategy, but when you consider the resource-sharing potential as well as the ability to more rapidly create new partitions and deploy new instances of applications, this might be more cost effective than it appears. This approach might also offer indirect cost savings in reduced floor space, power, and cooling requirements.

At this point, you have essentially completed a server consolidation exercise and have created a plan to reduce the number of physical servers in your data center with the addition of new servers and the use of partitioning solutions. If you chose cell-based system, the next step is to consider adding iCAP or TiCAP resources.

## Utility Pricing

It's easy to see that you should use iCAP resources if you are planning to run on a cell-based system. Perhaps you even chose a cell-based system specifically because of the iCAP capabilities. iCAP is a very cost-effective way to have spare capacity available for growth, whether expected or unexpected. Most of the cost for this additional capacity is deferred until you need it, and it can be activated dynamically with no disruption to your users. It's a very low-risk and simple option.

TiCAP resources are useful if the increases in your workload demand are spiky or cyclical (seasonal, monthly, or weekly). Again, this is not risky to implement, and it will add only a small amount of complexity in order to automate TiCAP resource management with WLM or gWLM so you can better manage its usefulness and cost.

Data from the worksheet in Table 2 can help you determine how many iCAP and TiCAP resources are appropriate for each independent workload. The number of cores required for peak processing is the number of TiCAP cores that you need, and the number of cores for expected growth is the number of iCAP cores that you need. Again, at this point in the process, this data is still independent of resource sharing. That topic is discussed in the next section.

## Organizing partitions and resource sharing

Now that you understand the resource requirements for each application or workload, and now that you have chosen which type of partitioning technology is the best fit, you must determine the best way to organize and combine the partitions. This can get a little complicated because of the numerous variables that factor into this decision. The first consideration is based on the sizing rules and limits. Table 4 describes the rules and limits for each type of partition.

---

Table 4. Partition sizing rules and limits

Type of partition	Sizing rules
nPars	Must be one or more cells.
vPars	One or more vPars can be in the same nPar. Maximum number of cells per nPar is 8 (when vPars are used). Maximum number of vPars per nPar is 8. Maximum number of vPars in one nPar is also a factor of CPU, memory, and I/O needed for each instance of HP-UX (cannot exceed resources of the nPar). Integrity VMs and vPars cannot be in the same nPar.
Integrity VM	Each virtual machine can have a maximum of 4 virtual CPUs. Maximum number of virtual machines per physical processor is 20 (due to 5% entitlement granularity). One VM Host per nPar or server (non-cell based) is allowed. Maximum number of virtual machines per VM Host is a primarily a factor of the CPU, memory, and I/O requirements of each guest OS (or 254, whichever comes first). Integrity VMs and vPars cannot be in the same nPar.

---

The second thing to consider when organizing the partitions is how the resources might be shared. To do this, you must further analyze each workload to understand its resource usage during peak processing periods as well as any expected growth. Try to determine when the peak usage periods

occur and how predictable this behavior is. If possible, combine virtual machines or vPars so that their peak processing periods don't overlap. (Capacity Advisor is very good at helping with this task.)

With Integrity VMs, the VM Host allocates the resources based on demand, so you can significantly increase the utilization of a server when peak processing periods don't overlap. With vPars, cores within the same nPar can be dynamically moved from one vPar to another so that resources can be applied where they are needed most. Memory can also be moved, but you would not want to do this to address short term peaks. A good example of when moving memory resources might be useful is if you had an application in one vPar that was busy during the day, and another application in a separate vPar that ran at night.

If your vPars (or Integrity VMs) are combined in such a way that you have room for extra processor cores in the nPar, then you can have a pool of iCAP cores with TiCAP usage rights. This is a very good way to flex the size of these partitions to respond to short-term processing spikes. For vPars, TiCAP cores can be activated for a specific vPar. For Integrity VMs, TiCAP resources are activated for the VM Host, and then you can increase the entitlement for a specific VM. If you choose to use TiCAP, HP recommends that you automate their activation and deactivation with WLM or gWLM.

If you have chosen nPars and need to flex their capacity from time to time, you can move the usage rights of an iCAP core from one nPar to another by deactivating an iCAP core in one nPar and activating an iCAP core in another. You can do this manually if the loading is very predictable and doesn't change too often; or you can add automation with WLM or gWLM.

## Automation

To achieve the most effective benefits of resource sharing or TiCAP resources, you need to use either WLM or gWLM. If you aren't currently using WLM, consider gWLM for its ease of use and its ability to define resource sharing policies for multiple systems. gWLM allows you to combine multiple workloads with differing demand patterns on a single server to make use of idle capacity, while guaranteeing that each workload receives its specified entitlement of resources. For multiple workloads that span partitions or servers, gWLM can move the resources to where they are needed most as defined by the policies you set up. TiCAP resources are enabled when required to meet the workload's demand and always disabled when not required which saves money. An operator or administrator would simply not be able to respond quick enough to shift resources around to meet short term demand. Additionally, TiCAP resources can be enabled or disabled at the workload level. The gWLM tool is not as complex as it might sound. For more information about using gWLM, see the link to the *HP Integrity Essentials Global Workload Manager User's Guide Version A.03.00.00* at the end of this paper.

## High Availability and Disaster Tolerance

It is beyond the scope of this paper to discuss HP Serviceguard high-availability and disaster-tolerant solutions in detail. However, you might want to consider adding high availability or disaster tolerance to your virtualized environment as you are planning your virtualized infrastructure. HP Serviceguard is well integrated with the virtualization products. For example, in a failover situation, gWLM can have a policy to adjust the workload entitlements on the failover server to provide appropriate levels of resource usage to the existing workloads, along with the workload from the failed system. Additionally, iCAP, TiCAP, or GiCAP resources can be activated.

## Implement the Pilot Project

Now that you have determined which virtualization technologies to use for your pilot project, you need to set up and test the new environment. As stated earlier, if your pilot project involves an application from an ISV, be sure to let the ISV know your intentions. They might already have done something similar and therefore might have some useful advice to help you.

One word of caution: If you are changing hardware platforms, operating system versions, or application versions, test the transition or upgrade before you begin to install and test the virtualization technologies. Changing too many variables at the same time is never a good idea.

## Repeat the Process

Once you have successfully implemented your pilot project, you can add new functionality to it if you need to, or you can start another pilot project. HP designed all of the virtualization technologies to be integrated and to work well together.

## For more information

For more information about HP VSE and virtualization technologies, see the VSE website at [www.hp.com/go/vse](http://www.hp.com/go/vse)

White paper: "An Introduction to the HP Virtual Server Environment," located at <http://docs.hp.com/en/11011/IntroToVSE.pdf>

Virtual Server Environment Reference Architectures: follow the appropriate links at [www.hp.com/go/vse](http://www.hp.com/go/vse)

Book: *The HP Virtual Server Environment: Making the Adaptive Enterprise Vision a Reality in your Data Center*, by Dan Herington and Bryan Jacquot, 0-13-185522-0  
[http://vig.pearsoned.com/store/product/1,1207,store-1974\\_isbn-0131855220,00.html](http://vig.pearsoned.com/store/product/1,1207,store-1974_isbn-0131855220,00.html)

White paper: "HW Consolidation with Integrity Virtual Machines" located at <http://docs.hp.com/en/9984/Consolidation2.3.pdf>

Technical manual: *HP Integrity Essentials Capacity Advisor User's Guide Version A.03.00.00*, located at <http://docs.hp.com/en/T2784-90026/T2784-90026.pdf>

Technical manual: *HP Integrity Essentials Global Workload Manager User's Guide Version A.03.00.00*, located at <http://docs.hp.com/en/T2762-90057/T2762-90057.pdf>

© Copyright 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Itanium are registered trademarks of Intel Corporation in the U.S. and other countries. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA1-5746ENW, October 2007

