

# System Sizing Guidelines for Integrity Virtual Machines Deployment

*Hardware Consolidation with Integrity Virtual Machines*  
Isom Crawford, HP BCS

A Technical White Paper from HP



Introduction.....	3
Using Integrity Virtual Machines for Hardware Consolidation.....	3
Identifying Workloads Suitable for Hardware Consolidation .....	4
Preparing for Consolidation - Utilization Reconnaissance.....	5
Peaks and Averages.....	5
Data Collection Tools.....	6
HP Operations Center's Performance Agent and GlancePlus .....	6
HP Capacity Advisor .....	6
Calculating Platform Utilization.....	7
Expected Platform Utilization .....	7
Expected Utilization Example .....	8
Expected CPU Utilization.....	8
Expected I/O Utilization .....	8
Extending the Utilization Calculations .....	8
Consolidation Considerations and Utilization Goals .....	9
Utilization Thresholds.....	9
Aggressively Increasing Hardware Utilization.....	10
Analyzing the Capacity of the Target System .....	11
Mass Storage Requirements for Integrity VM.....	11
VM Host System Configuration for Optimal Memory Utilization .....	11
Virtual Memory Base Page Size .....	11
File Cache Size.....	12
Inode Caches.....	12
Memory Requirements for the VM Host .....	12
Additional Memory Requirement for I/O Adapters on Integrity i2 Servers .....	12
Storage Requirements for each Virtual Machine .....	13
Calculating the Virtual Machine Capacity of an Existing Server .....	13
VM Host Memory and Storage Sizing Example.....	13

Identifying Memory Available for Use by Virtual Machines .....	14
Identifying the Capacity Requirements for a New VM Host .....	14
Computing the Memory Capacity for the VM Deployment .....	14
Computing the Storage Capacity for the VM Deployment.....	15
Consolidating Workloads on the VM Host .....	15
Stacking Workloads - Example 1 .....	15
Stacking Workloads - Example 2 .....	16
Stacking Workloads - Example 3 .....	17
Consolidating Symmetric Multiprocessing (SMP) Systems .....	17
Deploying Consolidated Workloads .....	18
Qualitative Utilization Information and Consolidation .....	18
Configuring I/O Based on Resource Utilization .....	19
The Benefits of Consolidating with Integrity VM .....	21
For More Information .....	21

## Introduction

Integrity servers help to meet your business needs by allowing you to consolidate multiple applications onto a single server using a combination of HP-UX 11i Partitioning Continuum and Virtual Server Environment (VSE) technologies. The partitioning techniques, which you can use alone or in combination, include hard partitions, virtual partitions, virtual machines, and resource partitions.

- **Hard partitions** (nPartitions) offer electrical isolation and cell board granularity, allowing you to service one partition while others are online. Integrity VM supports multiple operating systems, including HP-UX, Windows Server®, and Linux®.
- **Virtual partitions** (vPars) are separate operating system instances on the same nPartition or server, with O/S, application, and resource isolation. HP-UX 11i Virtual Partitions enable you to dynamically reallocate CPU resources between vPars as workload requirements change. Virtual Partitions also offer single CPU core granularity.
- **Virtual machines** have their own separate operating system instances (guests), on the same nPartition or server, with different OS versions, applications, and users, in a fully isolated environment. HP Integrity Virtual Machines software provides shared CPU (with sub-CPU granularity), shared I/O, as well as dynamic CPU and memory resource allocation based on demand and entitlement.
- **Resource partitions** allocate resources to specific applications and users within an operating system. They offer fully dynamic allocation of resources, including CPU or sub-CPU and percent memory granularity.

Successful IT consolidation efforts regularly use HP partitioning to isolate production environments from test and development environments. HP partitioning and VSE can:

- Improve your RoIT through optimized server utilization with minimal overhead.
- Increase server flexibility through easy resizing of partitions.
- Reduce risks by offering a variety of isolation choices, including hard partitions, virtual partitions, virtual machines, and resource partitions.
- Improve service levels, such as application response time, through the intelligent policy engine (HP-UX Global Workload Manager), which provides automatic provisioning of resources based on predefined service level objectives.

## Using Integrity Virtual Machines for Hardware Consolidation

Virtual machines share hardware while providing complete operating system instance isolation, so the Integrity Virtual Machines (Integrity VM) product is the perfect solution for hardware consolidation. Using virtualization, you can lower overall hardware costs by:

- Reducing the complexity and number of hardware components required to run multiple operating system (OS) instances.
- Increasing overall hardware utilization.

Virtualization makes it easier to share hardware because one Integrity Server can host multiple instances of an operating system. However, combining multiple OS instances on a single physical machine requires planning.

The steps to consolidating workloads with Integrity VM are:

1. Identify hardware utilization characteristics of workloads that are candidates for consolidation.
2. Determine utilization levels for the workloads based on averages or peaks – carefully consider the latter if the workloads all peak at the same time.
3. Convert utilization levels to expected utilization on the physical server's hardware.
4. Identify an aggregate (physical hardware) utilization target/threshold you are comfortable with. It is usually prudent to start low, to allow for workload growth.
5. Calculate the size of the physical server(s) required for consolidating the workloads, taking into consideration the memory and mass storage requirements of the VM Host as well as those for each workload's VM.

## Identifying Workloads Suitable for Hardware Consolidation

A key consideration during the planning stage is to determine whether the workloads are suited for hardware sharing or dedicated hardware. Dedicated hardware solutions include standalone servers, hard partitions, or virtual partitions.

The hardware requirements often determine whether the workload is suitable for hardware sharing or dedicated hardware solutions. Workloads that may be better suited for dedicated hardware include:

- Extremely I/O or OS (system call) intensive
- High utilization, i.e., resource utilization exceeding 50% of target hardware configuration

Workloads that are ideal candidates for hardware sharing using Integrity VM include:

- Production workloads with low (less than 20% of target) hardware utilization
- Software development
- Software validation and quality assurance testing
- Functional test of software, including proof-of-concept functionality

Software development and test of tiered applications (such as client-server) are particularly well suited for Integrity Virtual Machines. The flexibility provided by unique configuration of the software stack in each VM lends itself to powerful consolidation development and test scenarios.

## Preparing for Consolidation - Utilization Reconnaissance

When considering a workload for hardware consolidation, collect the utilization data that is available, note its characteristics, and determine what information is important and what is not. If you are preparing to consolidate several workloads, collect the utilization information (both quantitative and qualitative) for every workload – not just for a few. You will need specific data to determine which workloads can be located on the same physical server.

Where possible, collect utilization and capacity information for all basic hardware components, including CPU, memory, network, and mass storage. Statistical utilization information for CPU is usually available; detailed data for utilization of other resources is not always available. If that is the case, collect qualitative information for the other resource utilization, especially I/O (network and mass storage). Categorizing utilization of these resources as “high,” “medium,” or “low” will be very useful for consolidation planning.

### Peaks and Averages

Start consolidation planning with average utilization of the workloads. Then you can look at the workload usage peaks to avoid consolidating workloads with coincident usage spikes on the same hardware.

Consider spike or peak utilization characteristics of workloads that are being consolidating on the same physical server. Random or out of phase peaks in utilization are not a consideration. Look for workloads that are in phase with other workloads, for which peak utilization occurs simultaneously. For example, if all the workloads are idle until the last five days of the month, then average utilization is not useful for sizing considerations. For this kind of utilization pattern among multiple workloads, use the peak utilization over the five peak utilization days in place of the average utilization.

## Data Collection Tools

Many system management products provide excellent utilization analysis tools. HP Operations Center products can help in identifying average utilization as well as time, duration, and extent of peak utilization (or 'spikes'). The Integrity Essentials Capacity Advisor also features workload and hardware utilization tools as well as capabilities for analyzing workload aggregation on various server configuration scenarios.

### **HP Operations Center's Performance Agent and GlancePlus**

Several tools are available for use in collecting utilization data. One common data collection tool is Operation Center's Performance Agent (OVPA). Operation Center products can collect information for multiple hardware components for various platforms and vendors. The Operation Center products can generate executive summary reports indicating average and peak utilization for various hardware components.

### **HP Capacity Advisor**

Capacity Advisor (CapAd), part of the Virtual Server Environment suite, collects and reports utilization data for systems across the data center. Collected data may be from either OVPA or the various VSE agents. HP Integrity Essentials Capacity Advisor allows you to collect and compare average and peak utilization for CPU, memory, networking, and disk resources.

Capacity Advisor can analyze multiple systems, predicting their potential as candidates for consolidation onto a single physical system. CapAd makes the analysis of utilization data easier, effectively eliminating the manual stacking and fitting of workloads. It combines and aggregates utilization data (from VSE agents or HP Operations Center) for a clear picture of the total utilization curve for the consolidated workloads. After you select multiple systems for aggregate analysis of their utilization data, you can define a scenario using the hypothetical (or existing) system on which you want to consolidate.

CapAd is an excellent tool for planning consolidation. It provides insights into workload behavior that might otherwise be difficult, if not impossible to predict. You can estimate peak utilization, average utilization and statistical percentiles (for example, utilization levels for 90% of a given time period).

## Calculating Platform Utilization

In many cases, candidate workloads for consolidation will move from other, probably older, servers. This immediately creates a challenge in addressing how to accommodate the difference in hardware performance. The problem of comparing different hardware platforms for purposes of workload performance expectation is not new and certainly not unique to Integrity VM. You must consider issues such as server performance capability and application performance metrics when consolidating workloads onto a new system.

### Expected Platform Utilization

Generally speaking, one needs a 'multiplier' to convert utilization on the "source" hardware resource to an expected utilization on the "destination" hardware. For example:

$$\text{old\_utilization} \times \text{multiplier} = \text{new\_utilization}$$

The multiplier is the fraction of the new hardware resource required to deliver one unit of the old hardware's performance. Typically, this will be a number less than 1 because the target hardware platform is an Integrity Server and the workloads are running on older hardware.

Every workload has a performance (and capacity) footprint that is unique. Therefore, you should use a different multiplier for each workload you are planning to consolidate. The best multiplier is the ratio of the workload's performance on the old hardware to that on the new hardware configuration. If you do not have a multiplier based on the actual workload's performance, you must use another approach to identify a multiplier.

For CPU utilization, the processor clock speed may be an adequate measure. This is especially effective when the processors are in the same architecture family, but it is of limited value when comparing processors from different architecture families (for example, PA-RISC and Itanium) due to the various CPU architecture features that impact performance, such as local cache size, cache line size, data latency, pipelining, number of functional units, and so on.

Standard benchmarks are sometimes good indicators of a given application's performance. They also provide the benefit of being available for various server configurations – across multiple vendors. SPECint® and TPC-C® are examples of such benchmarks.

Do not use peak performance specifications when deriving multipliers for network or mass storage utilization. If possible, use realizable bandwidth or latency performance (whichever is appropriate for the workload) of the I/O hardware to predict utilization.

## Expected Utilization Example

After you identify the multipliers, you can generate the utilization information in a very straightforward way. Consider a workload with the following utilization footprint:

CPU – 15%

Disk – 5%

Network – 2%

This workload currently resides on an HP 9000 Model A500 single-CPU server with one network interface card (NIC). The server used for consolidation is an Integrity rx4640 (1.5GHz CPU, 4MB cache Itanium 2 processor). The new server will be connected to the same SAN that the A500 is currently using. Some enhancements will be made to the network configuration for the new server.

### Expected CPU Utilization

From experience, the workload's performance is known to be very similar to the SPEC CINT2000 (SPECint) benchmark. The CINT2000 results for both platforms are available – 422 for the A500 and 1372 for the rx4640. Hence, calculate the multiplier as follows:

$$\frac{\text{A500 CINT2000}}{\text{rx4640 CINT2000}} = \frac{422}{1372} \approx 0.307$$

As a result, the expected CPU utilization is:

$$15\% \times 0.307 = 4.6\%$$

That is, the expected utilization will be 4.6% of a single CPU on the rx4640.

### Expected I/O Utilization

The mass storage configuration is roughly the same, so the multiplier used will be 1.0. The expected utilization will be simply 5%, as it was before.

The workload is much more sensitive to network bandwidth than latency. The new network configuration has 80% more bandwidth than the old network. The formula to generate the multiplier is:

$$\frac{1}{1.8} \approx 0.556$$

Hence, the expected network utilization is:

$$2\% \times 0.556 \approx 1.1\%$$

Hence, the expected utilization of one physical NIC on the rx4640 is 1.1%.

## Extending the Utilization Calculations

The previous examples explain the basic arithmetic for calculating expected utilizations, but the impact of workload consolidation also depends on the number of CPUs and I/O cards to which the utilization applies. For simplicity, these examples use one of each, but a real scenario may have, for example, four CPUs and two NICs on the existing server. In this case, multiply the utilizations by the number of hardware units (for example, a four-CPU configuration results in a utilization of 4.6%).

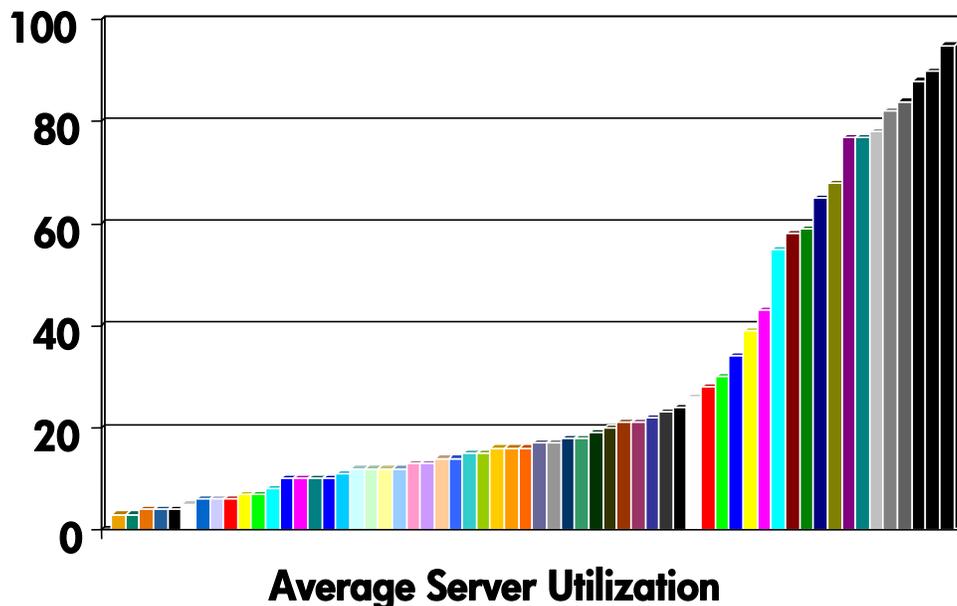
## Consolidation Considerations and Utilization Goals

After you collect utilization information for all of the workloads, consider the capacity of the target physical server in the context of these workloads. Specifically, you need to identify the maximum aggregate hardware utilization level that is acceptable. This level ultimately determines how many workloads you should consolidate on the target server.

### Utilization Thresholds

HP Consulting has found that average utilization in a large sample of data centers is near 50%. Other studies indicate an average utilization nearer 30%. Many data centers have a large number of servers with utilization well below 20%; for these sites, hardware consolidation achieving 50% utilization provides substantial savings. Figure 1 illustrates the typical utilization curve for servers in major data centers.

**Figure 1 – Representative curve of server utilization across a data center**



Consolidating for utilization levels of 80% to 100% is probably not realistic for several reasons. Extremely high utilization targets, or thresholds, allow no room for error and do not accommodate spikes in utilization. In late 2005, Gartner<sup>1</sup> recommended that their clients “target 40% to 50% final average utilization” as a best practice in deploying and managing server virtualization.

<sup>1</sup> *Server Virtualization: Taking Charge of Your Servers*, Tom Bittman, Gartner 24<sup>th</sup> Annual Data Center Conference, December 2005

Depending on the types of workloads and the Integrity server you are using for consolidation, you can set your target higher (or lower).

Use the high end of your target utilization range for CPU consolidation and the low end for I/O (both network and mass storage) consolidation. Consider higher thresholds (beyond the upper range) for workloads that are computationally intensive and have very little interaction with the operating system or I/O systems. Integrity VM uses the HP-UX fair-share scheduler (FSS) technology to allocate CPU resources. This technology has been in production for several years as a core component of HP's Process Resource Manager (PRM) and Workload Manager (WLM) products. Its efficiency has steadily improved over the years and adds very little overhead to CPU-intensive workloads.

Virtualization and sharing of I/O hardware by Integrity VM (or any other virtualization solution) will have some impact on performance. Rerouting transactions from the virtual device to the physical device may require more time and more resources. Because of this, the utilization threshold for I/O utilization (both mass storage and network) must be more conservative and a utilization threshold on the low end of your target range is recommended. If extensive resource sharing is acceptable (e.g., in disaster recovery), then a higher threshold may be considered. Even so, that threshold should not exceed the high end of your target range.

## Aggressively Increasing Hardware Utilization

Integrity servers with higher physical CPU counts may be candidates for higher CPU utilization thresholds. Marginal increases are appropriate for servers with four to eight physical CPUs; they can be higher for servers with more than eight physical CPUs. For example, using 50% as a baseline, consider thresholds of 60% for servers with four to eight CPUs and a threshold approaching 70% for servers with more than eight physical CPUs.

If you choose to use aggressive hardware utilization ranges, consider using HP Global Workload Manager (gWLM) to create resource allocation policies – including performance targets relative to priorities – to workloads executing in the virtual machines. The gWLM can automatically reallocate CPU resources to the virtual machines in order to meet service-level objectives defined in the workloads' policies. Thus, even during periods of extremely high utilization, gWLM guarantees high priority workloads access to the necessary resources to perform at the desired performance level.

The gWLM can also help address periods of higher CPU resource utilization by managing HP's Instant Capacity products. The gWLM product allows you to assign a policy to the VM Host server to activate Instant Capacity CPU resources whenever hardware utilization exceeds the specified level.

# Analyzing the Capacity of the Target System

After identifying the resources required for the workloads being consolidated, one is ready to begin calculating the capacity required by the Integrity VM Host system. To do so, the storage requirements for the Integrity VM product must also be accounted for. Subsequently, you are ready to either

- Calculate the virtual machine capacity of an existing Integrity server, or
- Identify the capacity requirements for a new Integrity VM Host server

In this section, we discuss how to calculate the capacity required for both the Integrity VM system along with the workloads to be deployed there.

## Mass Storage Requirements for Integrity VM

Installation of Integrity VM on a physical Integrity server includes the Integrity VM software and the HP-UX operating system. You must consider the mass storage and memory required for their operation when you are determining the capacity of the physical server. The following summarizes the capacity requirements for Integrity Servers running Integrity Virtual Machines version 4.0 and later.

Mass storage required for the Integrity VM software installation (independent of the individual virtual machines used by the workloads) includes disk space sufficient for the HP-UX OE installation (minimally 20 GB) and disk space required for swap. No additional swap space is required for virtual machines with Integrity VM version 4 and later. In general, refer to the “HP-UX 11i v3 Installation and Update Guide” and “HP-UX 11i v3 Read Before Installing or Updating” documentation for swap space recommendations with HP-UX 11iv3. As a result, the storage required on the Integrity VM Host system is:

$$\text{VM Host disk space} = (\text{HP-UX OE installation}) + (\text{swap space})$$

### Storage Requirements for each Virtual Machine

In addition to the above requirements for the VM Host, each virtual machine (VM) requires sufficient disk space for the workload (OS and application). These disk requirements are the same as those of the workload on a physical server.

## VM Host System Configuration for Optimal Memory Utilization

The Integrity VM Host system is a dedicated, special purpose server. As a result, tuning the VM Host system is required for optimal efficiency of VM operation. To do so, HP recommends the following tunable settings for the VM Host system only. Configure and tune Individual VMs for the workload they are executing using the tunable settings recommended by HP for Integrity servers.

### Virtual Memory Base Page Size

Increasing the `base_pagesize` tunable to 64 will provide significant memory savings on the VM Host system. Some software may be impacted and rendered inoperable by setting the `base_pagesize` to anything other than the default value. For more information, consult the white paper *Tunable Base Page Size*, available from HP’s documentation website. To set the `base_pagesize` tunable to the recommended value of 64K:

```
kctune base_pagesize=64
```

Note that changing this tunable does require a reboot of the system.

## File and Buffer Cache Sizes

By reducing the size of the file cache on the VM Host, the memory overhead of Integrity VM is reduced significantly, especially for systems with more than 32GB of physical memory. HP recommends the file cache tunables be set to use less than 128MB of memory:

```
kctune filecache_max=128000000 filecache_min=128000000
```

Avoid this configuration on deployments with virtual machines that use a significant number of virtual disks mapped to files (file backing stores) as it may have significant performance impact. The `filecache_min` and `filecache_max` tunable values must not use alphabetic characters (such as `m` or `g` suffixes) as that will prevent Integrity VM from starting up gracefully.

Another cache that can be reduced on Integrity VM Host system is that of the VxFS metadata buffer cache. By default, this cache can vary in size dynamically and consume up to 1GB or more of memory. HP recommends changing this tunable, limiting the amount of memory used by the VxFS metadata buffer cache to 64MB:

```
kctune vxfs_bc_bufhwm=64000
```

It is important to note that the value of this tunable is specified in kilobytes.

## Inode Caches

Inode caches on the VM Host system need not be very large due to its using a relatively smaller number of files. To reduce the size of the HFS inode cache:

```
kctune ninode=1024
```

Note that Integrity VM does not support virtual disks mapped to HFS files. Moreover, VxFS is the default file system for HP-UX and the VxFS inode cache is not controlled by this tunable.

HP also recommends reducing the size of the VxFS inode cache through the following tunable settings:

```
kctune vx_ninode=32768 vxfs_ifree_timelag=-1
```

The VM Host system typically uses very few files when compared to general-purpose HP-UX systems. As a result, the inode cache for the VM Host system should be significantly smaller than typical HP-UX systems. The tunable settings above will reduce the VxFS inode cache to the minimum recommended by HP.

## Memory Requirements for the VM Host

Based on the recommendation that the `base_pagesize` tunable is set to 64 (for 64K page size), the memory required by the physical VM Host - independent of that required by each virtual machine deployed - includes 1.2GB plus 8.5% of physical memory. That is:

$$\text{VM Host memory} = 1.2\text{GB} + 8.5\% \times (\text{Total Physical Memory})$$

Configuring the VM Host system to use base page size other than 64K will cause the VM Host memory requirement to increase significantly.

### Additional Memory Requirement for I/O Adapters on Integrity i2 Servers

The VM Host memory formula above is the minimum required for the HP-UX operating system and virtualization layer. It does not account for memory needs of certain I/O adapters. For that reason, one must adjust the VM Host memory accordingly. This results in the actual VM Host memory requirement being:

$$\text{VM Host memory} = 1.2\text{GB} + 8.5\% \times (\text{Total Physical Memory}) + (\text{Memory for I/O Adapters})$$

For most Integrity systems, the memory required for I/O adapters is negligible. One exception is the HP Integrity BL8x0c i2 Server Blade system – based on the Intel Itanium processor 9300 series. The 10GbE (10GigE) network ports on BL8x0c i2 server blades require an additional 500MB of memory **per 10GbE network port** when they are configured for Flex10 or jumbo frames. If these ports have both Flex10 *and* jumbo frames enabled, then the additional memory requirement is 1GB **per 10GbE network port**. Refer to latest memory subsystem documentation for changes to these requirements.

For example, consider a BL860c i2 server blade with 64GB of memory with the default number of 10GbE network ports (4). If those ports are configured for both Flex10 and jumbo frames, then the VM Host memory requirement for that system is:

$$1.2\text{GB} + 8.5\% \times (64\text{GB}) + (4 \text{ adapters} * 1\text{GB}/\text{port}) = 1.2\text{GB} + 5.44\text{GB} + 4\text{GB} = 10.64\text{GB}$$

### Memory Requirements for each Virtual Machine

When defining a virtual machine to host a given workload, allocate the amount of memory as recommended for a physical Integrity server hosting that same workload. Defining a virtual machine with less memory than is recommended for the workload will lead to performance problems – just as it would on a physical system.

The physical memory required for VMs is the sum of sufficient memory for its operating system and workload plus an additional 8.0 to 8.3% of the VM's memory (as specified in its configuration) for VM Host data structures.

---

### Best Practice

To simplify sizing exercises, HP recommends you always allow for the maximum memory overhead for VM memory: 8.3%.

---

## Calculating the Virtual Machine Capacity of an Existing Server

This section illustrates the use of the requirements discussed above to size an Integrity system for use as a VM Host.

### VM Host Memory and Storage Sizing Example

On a physical Integrity server with 16 GB of physical memory and assuming a swap space size of 4 GB, calculate the disk storage requirements as follows:

$$\begin{aligned} & (\text{Host OE installation}) + (\text{swap space}) \\ & \approx 20 \text{ GB} + 4 \text{ GB} = 24 \text{ GB Mass Storage} \end{aligned}$$

Calculate the memory required for the VM Host (independent of VM requirements) as follows:

$$\begin{aligned} & 1.2 \text{ GB} + (8.5\% \text{ of physical memory}) \\ & = 1.2 \text{ GB} + 8.5\% \times 16 \text{ GB} \approx 2.56 \text{ GB Physical Memory} \end{aligned}$$

As a result, the total amount of memory available for use by virtual machines on this system is:

$$(16 \text{ GB} - 2.56 \text{ GB}) = 13.44 \text{ GB}$$

## Identifying Memory Available for Use by Virtual Machines

An additional 8.3% of memory is required by Integrity VM for memory to be used by a VM. Therefore, if you define a VM to have 10 GB of memory, then Integrity VM requires 10.83 GB of memory to instantiate that VM. Similarly, a VM defined with 4GB of memory requires approximately 4.33 GB of memory to instantiate it. Note that these two VMs cannot be instantiated concurrently on the physical server configuration described in the previous example because the total required memory by the VMs is 15.16 GB of memory, but only 12.94 GB is available.

You may calculate the amount of memory to be used in configuring VMs by dividing the total memory available (after the VM Host requirements) by 1.083.

For the VM Host with 16GB in our previous example, the effective amount of memory available to VM configurations is:

$$13.44 \text{ GB} / 1.083 \approx 12.4 \text{ GB}$$

As a result, we may configure two virtual machines with 5GB of memory each on this system (and one of them may actually have 7.4 GB of memory).

When configuring virtual machines to use dynamic memory allocation or Integrity VM's Automatic Memory Reallocation (AMR), the amount of memory defined for a VM may be more than the amount summarized here. Consult the Integrity VM product documentation for more details.

## Identifying the Capacity Requirements for a New VM Host

The previous section is based on the assumption that an *existing* server is to be used as a VM Host. When consolidating multiple physical systems onto virtual machines – each with its own individual capacity needs – one may need to size a new server. This section addresses the scenario in which a new Integrity server will host a set of virtual machines.

Consider the deployment of several VMs on a single physical server. Suppose the total aggregate VM mass storage (disk) capacity requirement is 550 GB and VM memory total is 28 GB.

### Computing the Memory Capacity for the VM Deployment

Let us consider memory requirements first. Using the relations outlined above, sufficient physical memory relates to the aggregate VM memory size as follows:

$$\text{Physical memory} \geq 1.18 \times (\text{aggregate VM memory size}) + 1.3 \text{ GB}$$

For our example, this is:

$$\text{Physical memory} \geq 1.18 \times (28 \text{ GB}) + 1.3 \text{ GB} \approx 34.34 \text{ GB}$$

*Note that memory requirements of I/O ports are not included in these inequalities and any such requirements will increase the physical memory requirement.*

## Computing the Storage Capacity for the VM Deployment

Storage capacity for each VM has no appreciable overhead. Therefore the VM needs exactly the storage needed for the application stack (including the operating system) to be run on that VM. In this regard, storage sizing is identical to that of physical servers.

As mentioned above, the VM Host requires a minimum of approximately 24GB of disk space.

Hence, the total storage is the sum of the aggregate required by the virtual machines and that of the VM Host:

$$\text{Physical Mass Storage} \geq (\text{aggregate VM mass storage}) + (\text{VM Host mass storage})$$

Using the example outlined above, we have:

$$\text{Physical Mass Storage} \geq 550 \text{ GB} + 24 \text{ GB} = 574 \text{ GB}$$

## Consolidating Workloads on the VM Host

The following examples show how to stack the workloads so that they do not exceed any particular utilization threshold.

Note that these examples assume the workloads are sensitive to I/O bandwidth, not necessarily to response time requirements. If response times are critical, then one will need to accommodate I/O operations per second rather than bandwidth. In doing so, the number of I/O adapters will need to be sized for that requirement.

### Stacking Workloads - Example 1

Consider an Integrity server with a single CPU, single NIC, and a single disk available to the virtual machines. The three candidates for consolidation have random spikes in utilization. The average expected utilizations of the new server's hardware components are given in Table 1.

**Table 1- Average expected utilizations workloads in Example 1.**

<b>Workload</b>	<b>CPU Utilization</b>	<b>I/O Storage Utilization</b>	<b>I/O Network Utilization</b>
A	10	5	20
B	20	10	15
C	30	30	20
<b>TOTAL</b>	<b>60</b>	<b>45</b>	<b>55</b>

Not all three of these workloads can be consolidated on the new server without adding other hardware components (another CPU, disk, or NIC). However, any two of them can be located on the same physical server because none of the possible aggregate utilizations would exceed the suggested thresholds.

## Stacking Workloads - Example 2

In this example, an Integrity Server hosts ten workloads, all with the expected utilizations as shown in Table 2.

**Table 2 – Average expected utilizations of workloads in Example 2.**

	<b>CPU Utilization</b>	<b>I/O Storage Utilization</b>	<b>I/O Network Utilization</b>
Each Workload	10%	10%	4%
<b><i>TOTAL</i></b>	<b><i>100%</i></b>	<b><i>100%</i></b>	<b><i>40%</i></b>

You can determine the integral number of physical components necessary to host these workloads by dividing the total aggregate utilization by the threshold utilizations and rounding up. For example:

$$\text{Number of CPUs} = 100\% / 50\% = 2$$

$$\text{Number of Disks} = 100\% / 40\% = 2.5 \approx 3$$

$$\text{Number of NICs} = 40\% / 40\% = 1$$

Therefore, the server will need two CPUs, three disks available to the virtual machines, and one NIC. Of course, the actual amount of available disk space must be well beyond what is required for the VM Host's system disk.

### Stacking Workloads - Example 3

For an Integrity server to host ten similar workloads that all experience peak load during the months of September through October, the workloads have the expected utilization characteristics described in Table 3.

**Table 3 - Average expected utilizationa of workloads in Example 3.**

	<b>CPU Utilization</b>	<b>I/O Storage Utilization</b>	<b>I/O Network Utilization</b>
Peak Utilization	22%	12%	8%
Average Utilization	5%	2%	1%

Because the workloads are in phase (that is, all peak during the same period), the averages are disregarded. Hence, you should use peak utilizations for sizing purposes. Calculate the hardware requirements as follows:

$$\text{Number of CPUs} = (10 * 22\%) / 50\% = 220\% / 50\% = 4.4 \approx 5$$

$$\text{Number of Disks} = (10 * 12\%) / 40\% = 120\% / 40\% = 3$$

$$\text{Number of NICs} = (10 * 8\%) / 40\% = 80\% / 40\% = 2$$

The Integrity server needs five CPUs, three disks, and two NICs to host the workloads in this example.

### Consolidating Symmetric Multiprocessing (SMP) Systems

Be sure to consider the number of CPUs on the existing hardware system when you estimate workload utilization. For example, the overall sizing calculation is 20% for an expected utilization of 5% on four CPUs,

Consider the motivation for the current system's multiple-CPU configuration when planning to consolidate that workload. Perhaps peak utilization requires the application to have several CPUs at its disposal or the workload will not perform well. You should reduce CPU count only after you carefully consider the workload requirements, and when the new hardware is considerably faster. If fewer virtual CPUs can handle the peak utilization, then the workload can be transferred to a virtual SMP system with that number of virtual CPUs. If not, consolidate the workload using a virtual machine with a CPU count equal to the one that is currently running the workload.

The number of physical CPUs on the Integrity VM Host must be at least the same as the highest virtual CPU count on any virtual machine it hosts. Integrity VM supports virtual machines with up to four virtual CPUs, but the VM Host system can have up to 128 physical CPUs.

## Deploying Consolidated Workloads

Response times and bandwidth alone are not enough to represent native performance. Workloads executing on virtual machines share the same hardware. In order to understand the impact of workload consolidation, it is critical to understand the way that the applications use hardware resources.

### Qualitative Utilization Information and Consolidation

If utilization data for resources (other than the CPU) is available only in a qualitative form, avoid consolidating more than one or two workloads with high utilization. It is best to combine poorly quantified workloads with workloads you know have a low utilization for the same resource. For example, if workloads A and B have high network I/O utilization, and workloads C, D, E, and F have low network I/O utilization, consolidate workload A with workloads C and D on one physical server, and consolidate workloads B, E and F on another physical server.

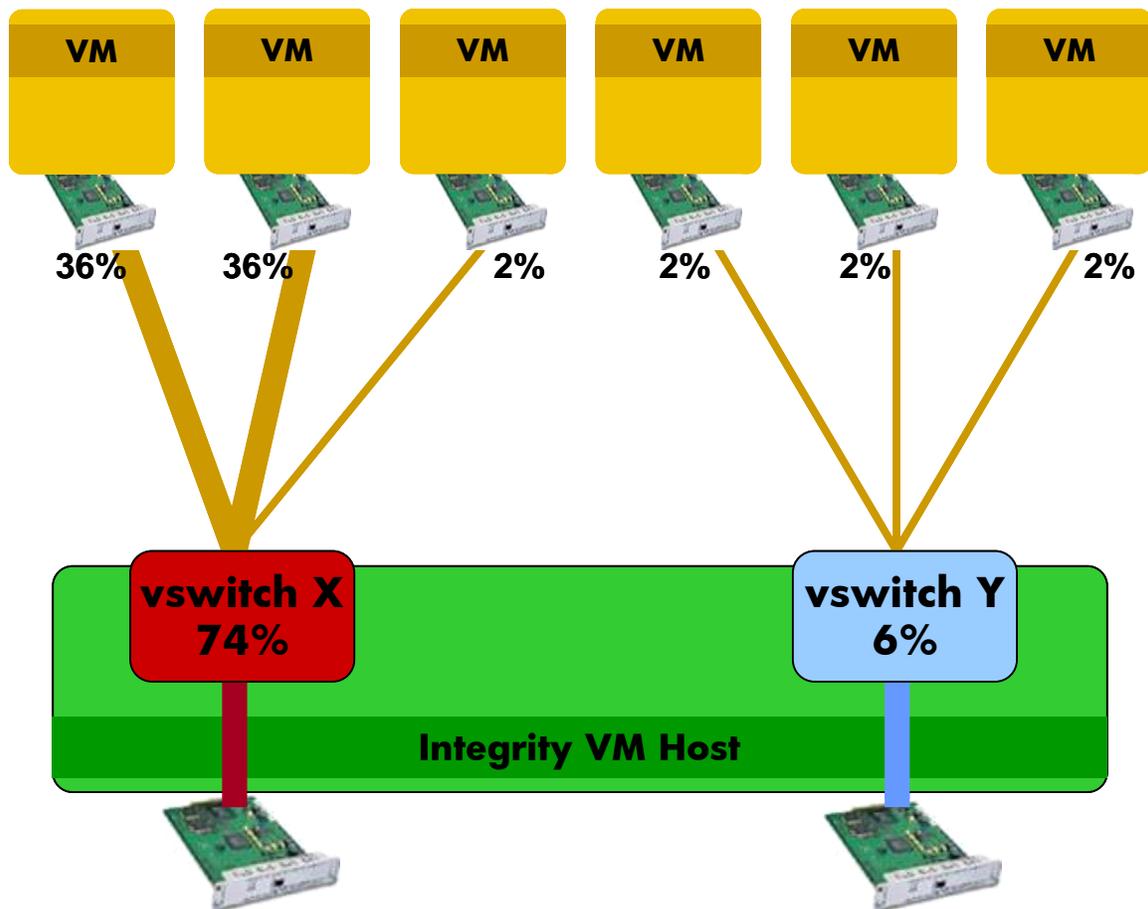
Performance problems are more likely when you consolidate multiple workloads with a medium or high utilization of the same resource.

## Configuring I/O Based on Resource Utilization

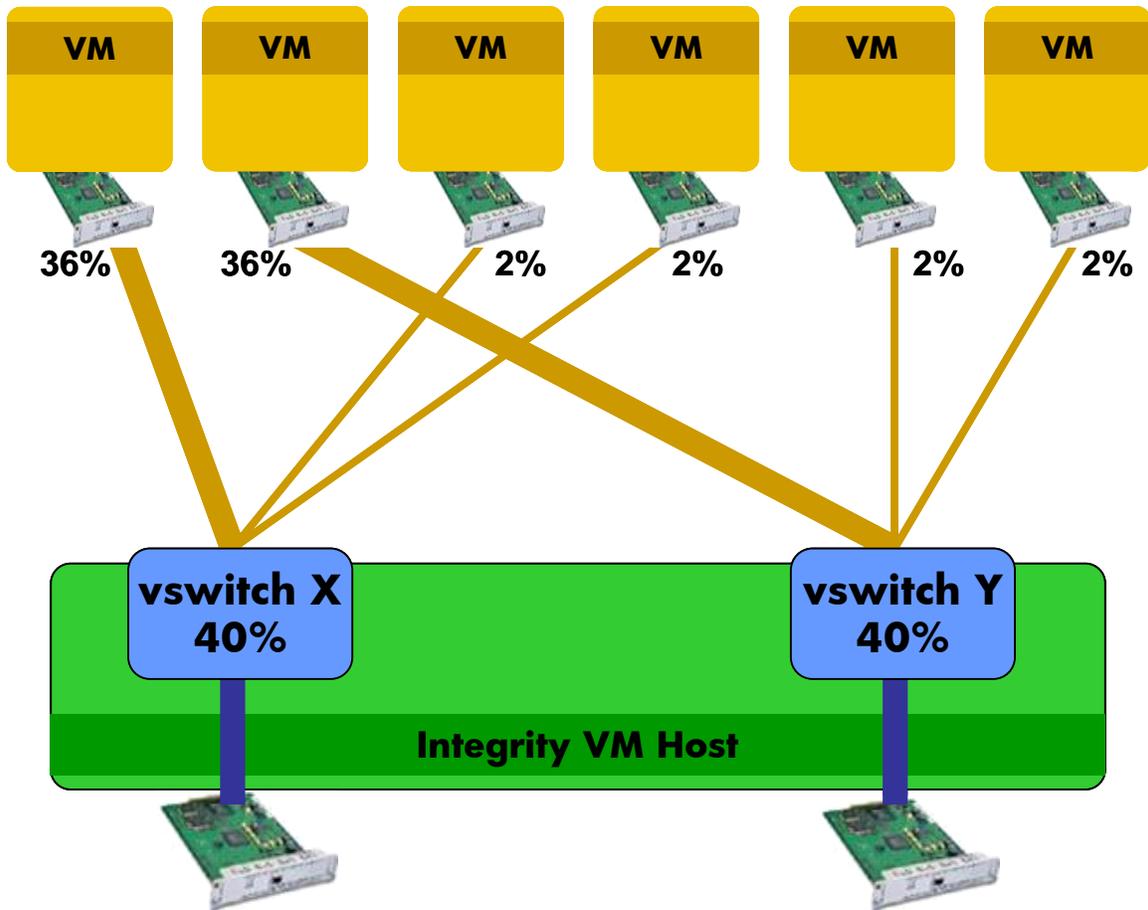
Once the utilization calculations are complete, do not forget about them when configuring virtual I/O devices. For example, if you have sized the VM Host for two physical NICs to be shared by six VMs, do not be careless when assigning those virtual NICs to virtual switches. Doing so creates a performance bottleneck as in the example shown in Figure 2. That is, be careful to distribute the virtual NICs to virtual switches according to their targeted utilization as illustrated in Figure 3.

The same concept applies to mass storage I/O as well. Be careful to configure the number of logical storage devices (e.g., LUNs) according to the consolidation plan and then connect the appropriate virtual disks to those logical storage devices when defining the virtual machines.

**Figure 2 – Avoid creating performance bottlenecks by careless assignment of virtual I/O to logical I/O devices.**



**Figure 3 – Careful assignment of virtual I/O devices is a critical step in server virtualization. Be sure logical I/O devices are loaded according to your plan**



# The Benefits of Consolidating with Integrity VM

Hardware consolidation with Integrity VM can achieve high levels of efficiency with little effort, as compared to consolidating workloads within a single OS instance or shared middleware.

The most likely area where consolidation with Integrity VM may prove troublesome is in failure to consider all of the hardware requirements of the workloads: network, storage, memory, as well as CPU. After you identify the workloads and the server capacity, plan the virtual hardware layout. This advance planning will lead to a smooth consolidation experience and increased hardware utilization.

## For More Information

More information about Integrity Virtual Machines is available online at <http://www.hp.com/go/integrityvm>

© 2010 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.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

Windows, and Windows Server are trademarks or registered trademarks of Microsoft Corporation in the U.S. and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

SPEC, SPECint, and CINT2000 are all registered trademarks of the SPEC organization.

TPC-C and TPC Benchmark are trademarks of the Transaction Processing Performance Council.

Rev 4.3, 1/2011, IC