

# Planning a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management



Executive summary.....	4
Audience .....	4
This white paper .....	4
Introduction.....	5
Adaptive Enterprise .....	5
Virtual infrastructure .....	5
Virtualization planning .....	5
Initial evaluation .....	6
Inventory and performance .....	6
Identifying appropriate performance metrics.....	6
Data collection.....	7
HP ProLiant Essentials Performance Management Pack (PMP).....	7
Products from AOG .....	7
Interpreting the results.....	7
Analysis.....	8
Which servers can I virtualize? .....	8
Which servers should I virtualize? .....	8
To which platforms do I consolidate? .....	9
Scale-up .....	9
Scale-out .....	9
HP BladeSystem .....	10
What are the infrastructure needs?.....	10
Storage .....	10
Networking .....	10
High availability.....	10
Backup and disaster recovery .....	11
Management .....	11
Security .....	11
Strategies and best practices.....	11
Virtualization strategies .....	12
Load balancing.....	12

Management.....	12
I/O performance .....	12
More information .....	12
Migration strategies .....	13
Storage strategies .....	13
Space needs .....	13
Service console .....	13
VMware File System (VMFS) .....	13
Additional storage.....	14
Storage Area Networks (SANs) .....	14
VMFS partitions.....	14
Number of VMs connected to a LUN .....	14
Disks.....	14
Clustering.....	15
Boot from SAN .....	15
High availability.....	15
Network strategies.....	15
ESX Server networking.....	15
Service console.....	16
VMotion.....	16
Teaming.....	17
VLAN .....	17
Virtual Guest Tagging (VGT) .....	17
External Switch Tagging (EST).....	18
Virtual Switch Tagging (VST).....	18
High Availability strategies .....	19
HP clustering solutions .....	19
Limitations .....	19
VM-based clustering .....	19
Network load balancing .....	20
Cluster-aware applications .....	21
Clustering software for VMs .....	21
Creating a cluster in a box .....	21
Multipathing in ESX Server .....	22
Disaster Recovery strategies .....	22
Backup strategies and products .....	23
HP OpenView Storage Mirroring .....	23
Backup/failover verification.....	24
Management strategies .....	24
HP Systems Insight Manager .....	24
Requirements in a Windows environment .....	24
Virtual Machine Management Pack.....	26
Infrastructure .....	26
Requirements .....	26
VMware VirtualCenter .....	27
Using the VC interface .....	27
Software components.....	27
VC server requirements .....	28
VC client requirements .....	28
VMware web service requirements .....	28
VC networking requirements .....	29
HP OpenView .....	29
Support .....	30

Security strategies .....	30
Authenticating users .....	30
Using certificates to secure remote sessions .....	31
Default permissions .....	31
TCP/IP ports for management access .....	32
Strengthening security .....	32
Trusted users only in the service console .....	32
Do not configure Promiscuous mode NICs .....	32
Consider disabling VM logging .....	33
Disable copy-and-paste in the VM .....	33
Additional security issues .....	33
Host machine .....	33
VMs .....	33
Appendix A – Using the HP ProLiant server sizer for VMware ESX Server .....	34
Using the sizer .....	34
Max Rates .....	35
Maximum CPU Utilization .....	35
Maximum Memory Utilization .....	35
Maximum Disk Utilization .....	35
Maximum Network Utilization .....	35
Additional Disk Space (GB) .....	36
Safe Calculations .....	36
Servers to Consolidate .....	36
Application .....	37
Preferences .....	37
Platform selection .....	38
Appendix B – HP ProLiant Essentials Performance Management Pack .....	39
Architecture .....	39
Hardware requirements .....	40
Analysis server .....	40
Client .....	40
Software requirements .....	40
Operating system .....	40
Additional requirements for the monitored server .....	41
Additional requirements for the client .....	42
Browser security .....	42
Appendix C – Using AOG CapacityPlanner .....	43
Appendix D – Using Microsoft Windows Performance Monitor .....	44
Perfmon counters .....	44
Evaluation best practices .....	45
For more information .....	46

## Executive summary

This white paper provides guidance on planning a VMware Virtual Infrastructure environment based on HP server, storage, and management products. The following key technology components are deployed:

- HP ProLiant servers
- HP ProLiant Essentials software
- HP StorageWorks Storage Array Network (SAN) products
- VMware ESX Server
- VMware VirtualCenter

This white paper is not designed to replace documentation supplied with individual solution components but, rather, is intended to serve as an additional resource to aid the IT professionals responsible for planning a VMware environment.

This is the first in a series of documents on the planning, deployment, and operation of an Adaptive Enterprise based on VMware Virtual Infrastructure and HP server, storage, and management technologies. The remaining documents in this series are a deployment guide, *Deploying a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management*, and a managing and operating guide, *Managing and operating a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management*. These guides can be found at: <http://h71019.www7.hp.com/ActiveAnswers/cache/71088-0-0-0-121.html>.

This white paper contains planning information that can help customers deploy a VMware ESX Server environment running on HP ProLiant servers, HP StorageWorks storage solutions, and HP ProLiant Essentials management components.

## Audience

This planning information contained in this white paper is intended for solutions architects, engineers, and project managers involved in the planning of virtualization solutions. The reader should be familiar with networking in a heterogeneous environment and with virtualized infrastructures, and have a basic knowledge of VMware ESX Server, and HP StorageWorks and HP ProLiant Essentials products.

## This white paper

This white paper includes information on the following topics:

- [Introduction](#) – outlines the HP Adaptive Enterprise strategy and virtualization
- [Virtualization planning](#) – lists questions to be answered during the initial evaluation of a virtualization project
- [Inventory and performance](#) – describes how to select appropriate metrics to monitor server performance in the current infrastructure, outlines data collection methodologies, and describes how to analyze the raw data collected
- [Analysis](#) – lists questions that can help in the analysis of the data collected in the previous section
- [Strategies and best practices](#) – outlines strategies and best practices for implementing various components of a virtual infrastructure (virtualization, migration, storage, networking, high availability, disaster recovery, management, and security)

- [Appendix A – Using the HP ProLiant server sizer for VMware ESX Server](#)
- [Appendix B – HP ProLiant Essentials Performance Management Pack](#)
- [Appendix C – Using AOG CapacityPlanner](#)
- [Appendix D – Using Microsoft Windows Performance Monitor](#)

## Introduction

A key component of an Adaptive Enterprise, as envisioned by HP, is the virtualization of resources. This section provides more information on these concepts.

### Adaptive Enterprise

The HP Adaptive Enterprise strategy combines industry-leading solutions, services, and products from HP and partners that can help organizations quickly turn challenges into opportunities. This strategy is based on four design principles – simplification, standardization, integration, and modularity – which, when applied consistently across business processes, applications, and infrastructure, will ultimately lead to an organization that can adapt to – even embrace – change. These design principles are applied to individual elements of the IT infrastructure and the entire infrastructure itself; in this way, organizations can create consistent building blocks that can be combined as needed.

Adaptive Enterprise is not a single product; it cannot be purchased “off the shelf”. It is a philosophy designed to make an organization agile and easily adaptive to changing business needs.

### Virtual infrastructure

Virtualization is one of the cornerstones to an Adaptive Enterprise. The primary benefit to virtualization may indeed be consolidation; however, a virtualized infrastructure can be beneficial in many other ways. For example, because an entire operating environment can be encapsulated in several files, that environment becomes easier to control, copy, distribute, and so on. If an organization virtualizes an operating system, its applications, configuration settings, and other desirable elements, that entire operating environment – known as a Virtual Machine (VM) – can be rolled out anywhere in the organization to maintain business continuity. To maximize availability, emerging technologies can allow VMs to automatically migrate from a potentially failing host to another virtualized platform – with little or no user intervention.

## Virtualization planning

A virtual infrastructure offers many benefits, including more efficient use of resources, reduction in server sprawl, and reduced capital expenditures for test and development environments. While ESX Server has been a boon for test and development activities, it is rapidly making its way into the mainstream production environments of many organizations.

Whatever your motive for moving to a virtualized environment, the key to a successful deployment is solid planning. This section guides you through the planning process.

## Initial evaluation

While money, knowledge, and time are always project constraints to some degree, you should always ask the following high-level questions when undertaking a consolidation or virtualization project:

- What are your currently useable resources?
  - How many servers are currently in use?
  - What is the knowledge level of virtualization?
  - What storage, networking, and software resources are available?
- How can you integrate virtualization into the current IT environment?
- How will virtualization impact current business processes?
- How will virtualization impact the current user experience?
- Which operating systems and applications can and should be virtualized?
- Which operating systems and applications should not be virtualized?
- Are you considering the use of server blades to consolidate server hardware?
- Which Physical-to-Virtual (P2V) processes should be used? What will the impact be?
- What are the implications of management, monitoring, and administration in the new infrastructure?
- Can redundancy and uptime levels be maintained with fewer servers?
- How do you make fewer servers more resistant to failure?
- At what level of virtualization does Return on Investment (ROI) become apparent?

While the above list is not a comprehensive, it should prompt the appropriate questions when starting to plan a virtualized infrastructure project.

## Inventory and performance

The first step in the evaluation process is to take a detailed inventory of the components of your computing environment. You should understand the server resources available to you and where these servers are located; it may be helpful to identify the entities that own and operate these resources.

In addition to taking an inventory, you should also understand the performance characteristics of the workloads running on the servers: not all server workloads make good virtualization candidates. There may be other barriers that prevent a particular server from being virtualized, such as the need for unsupported I/O devices.

## Identifying appropriate performance metrics

It is essential to understand your current environment when evaluating candidates for virtualization. A wide range of metrics is available to help you characterize performance: for a web server, for example, you may choose to focus on requests/second or, for a database system, you may choose transactions/second. Although readily available, these can be closely-focused metrics that describe how an application is performing but provide little information on overall server performance.

To better your computing environment, you need to understand performance at the server-level – more precisely, at the levels of major server subsystems (CPU, memory, disk, and network). When gathering or analyzing performance data, you should focus on the metrics that describe what is happening at a physical level.

## Data collection

Since performance characterization can only be as effective as the performance metrics collected, the largest and most critical part of the characterization process becomes data collection (sampling).

Data should be sampled over as long a period as possible and should be representative of your business processes and cycles. For example, if you are considering virtualizing a server or server farm responsible for month-end batch processing that would typically result in higher than average utilization rates at that time, be sure to include the month-end time period in your sampling scheme.

While creating a better representation of server operating characteristics, lengthy sampling periods may conflict with normal business operations. If sampling over an extended period is impossible, take samples during the most performance-critical business operations.

The following tools can aid inventory collection and help characterize the performance of your computing environment.

### **HP ProLiant Essentials Performance Management Pack (PMP)**

PMP can detect and analyze hardware bottlenecks on HP ProLiant servers. This information can be interactively displayed or logged to a database for later analysis or reporting.

PMP integrates with HP Systems Insight Manager (HP SIM) to provide a complete performance monitoring and inventory tool for an HP ProLiant server environment.

Refer to [Appendix B – HP ProLiant Essentials Performance Management Pack](#) for more information.

### **Products from AOG**

AOG (Asset Optimization Group) offers products and services to help with data and inventory collection. Their agent-less capacity planning tools help gather the data you need to quickly optimize your infrastructure and develop a server virtualization strategy.

For more information on how AOG CapacityPlanner software can help during this stage of the planning process, please visit their web site at [www.aogtech.com](http://www.aogtech.com).

Refer to [Appendix C – Using AOG CapacityPlanner](#) for information on CapacityPlanner.

## Interpreting the results

The effort required to interpret data collected by popular performance monitoring tools can vary greatly, as shown in the following examples:

- Data collected by CapacityPlanner require no further interpretation – system performance is diagrammed and plotted against profiles of industry averages
- PMP produces relatively concise reports to represent system performance; however, these reports may still require some interpretation backed by an understanding of server workloads and expected (or baseline) results.
- If Microsoft® Windows® Performance Monitor (Perfmon) is used, much evaluation is required after raw data are collected. Consider the following scenarios where interpretation is required:
  - If the Average Disk Queue Length is three or greater, the disk subsystem is over-taxed **if** there is sufficient system memory and swapping is not causing the queue to be high
  - If % Processor Time is low and idle times are high, the processor is under-used if no I/O subsystem is bottlenecked, preventing applications from utilizing the CPU

## Analysis

Once you have collected as much information as possible about your computing environment, you can begin to analyze this information prior to developing a virtualization strategy. This section discusses the types of questions that need to be answered when developing a virtualization execution plan. The following section, *Strategies and best practices*, discusses these questions in more detail and provides the tools to help answer them.

### Which servers can I virtualize?

One of the first questions most people ask when considering virtualization is, “Which servers can I virtualize?” Luckily the answer to this question is relatively simple: most servers and workloads can be virtualized. Noted exceptions include the following:

- Servers utilizing more than two CPUs or more than 3.6 GB RAM
- Workloads with high utilization (over 85%) of resources such as CPU, disk, network, and/or memory
- Servers or applications that require use of specialized hardware devices

### Which servers should I virtualize?

The next question to ask is, “Which servers **should** I virtualize?” This question is not as easy to answer and depends on your goals and expectations for the virtualized environment. However, there are some general characteristics that can be used to identify servers that make prime virtualization candidates.

These characteristics include:

- **Workloads with low utilization rates and small footprints**  
Generally, these are infrastructure- and appliance-type workloads such as file and print servers or web servers; small departmental and home-grown applications may also fall into this category. Often, these workloads follow the one-application-to-one-server model, resulting in many underutilized computing resources.
- **Servers that are chronically reconfigured**  
These are often test, development, and staging servers that go through regular cycles of reconfiguration and provisioning. By virtualizing these servers, you can dramatically reduce the time it takes to reconfigure the environment.  
Additional features of ESX Server (such as undoable disks) can also be beneficial in a development and test environment.

---

#### Note:

For more information on disk modes see the ESX Server Administration and Installation guides at <http://www.vmware.com/support/pubs/>.

---

## To which platforms do I consolidate?

Once you have identified your virtualization candidates, you now can consider which server platforms you should utilize as your virtualization hosts. The basic question becomes whether to scale-up or scale-out.

Unfortunately, virtualization is not a “one size fits all” solution. While any HP ProLiant or HP BladeSystem server makes a suitable virtualization platform, considerations like features, performance, and Total Cost of Ownership (TCO) should be carefully weighed against the proposed virtualized workload in order to identify the appropriate platform.

### **Scale-up**

Although the terms “scale-out” and “scale-up” are somewhat loosely defined, in general, scaling up implies the use of fewer, large-capacity servers such as larger 4P or 8P models.

In the past, one of the shortcomings of scaling up has been the lack of applications that can effectively scale past two or four processors. Since resources are not fully utilized, running these types of applications on a large-capacity server does not often yield the best price/performance. Virtualization makes more sense when multiple operating systems and applications are running on the host rather than a single application – an approach that scales well and leads to much better resource utilization.

Another benefit to the scale-up approach is the cost savings realized by sharing infrastructure resources such as NICs, HBAs, and their corresponding switch ports. Virtualization allows these resources to be shared by all VMs on the same physical host. By consolidating more VMs on fewer servers, you also reduce the number of infrastructure components necessary to provide network and storage connectivity to your VMs.

One potential area of concern to the scale-up approach is the impact of server downtime – both scheduled and non-scheduled. Unlike a conventional environment where server downtime typically only impacts a single application, a virtual host that has failed or is otherwise taken down for scheduled maintenance impacts every VM on the host. Because of the potential for higher-capacity servers to be hosting a large number of VMs, it is increasingly important to make sure that these servers are always available. The deployment of high-availability features such as redundant power supplies, fans, and ROMs; RAID storage and memory; and pre-failure warnings should be a top concern when evaluating virtualization platforms, particularly when using a scale-up approach.

### **Scale-out**

The scale-out approach involves consolidating servers on to more 2P and smaller 4P servers. Scaling-out with many physical servers provides greatly flexibility in distributing your virtualized workloads, which, in turn, can lead to better resource utilization. By creating large server farms, you can easily move your VMs from host to host, redistributing workloads as necessary to accommodate fluctuating resource demands.

Although having highly-available servers is always important, a scale-out approach can help mitigate some of the effects of a downed server. Because fewer VMs are typically hosted on a smaller server, the VMs from a failed server or one that has been brought down for maintenance can easily be redistributed among the remaining servers in the farm until the server can be brought back online. However, this can also be more difficult in a scale-out environment where the servers may not have enough capacity to accommodate the larger displaced VMs.

Of course, there are drawbacks to the scale-out approach: for example, more servers mean more infrastructure components. While some costs savings can be realized by VMs sharing network and HBA ports, greater savings in infrastructure components are likely when using a scale-up approach. More servers also mean greater management costs: server management and maintenance can be a significant portion of the total cost of ownership – often in the form of human resources.

## HP BladeSystem

Delivering the benefits of both scale-up and scale-out approaches, HP BladeSystem can provide an ideal platform for building a virtualized infrastructure.

Using integrated network and SAN switches, HP BladeSystem provides common network and storage connectivity to an entire enclosure of HP BladeSystem servers. As such, an HP BladeSystem enclosure can be thought of as a single large-capacity server, with the enclosure's shared network and storage connectivity providing cost savings in the same way as the scale-up model. As needed, the enclosure is then populated with individual HP BladeSystem servers, providing the flexibility and availability afforded by scaling-out.

Both 2P and 4P HP BladeSystem servers can be plugged into a single enclosure, offering flexibility in the way that capacity is expanded. Moreover, since the power, network, and storage cabling is already in place, adding servers is simply a plug-and-play operation, dramatically reducing management costs and time-to-deployment.

## What are the infrastructure needs?

While deciding on a server platform for consolidation, you should also be thinking about the additional components that make up a virtual infrastructure. Some of your choices or limitations may influence your server purchasing decision, so it is best to consider all aspects of a virtual infrastructure before moving forward.

### Storage

Storage is a critical component of a virtual infrastructure. A Storage Area Network (SAN) is often used to provide centralized storage for virtual machines, making it easier to move VMs from host to host for optimized workload distribution and enabling technologies like VMotion intelligent workload management software.

---

#### Note:

VMotion is a VMware technology to migrate running VMs from one ESX host to another without service interruption.

---

#### Note:

Proper SAN configuration is essential to providing a robust and high-performing virtual infrastructure.

---

### Networking

Networking is another important aspect to consider when planning your virtual infrastructure. When moving from a physical to a virtual environment, you may be consolidating servers from different segments of your networks with different connectivity requirements on to a single server. Providing a resilient network that meets the connectivity needs of all your VMs requires some careful, up-front planning.

### High availability

Making sure servers and applications are always running is an important part of any mission-critical computing environment.

High availability can be achieved at both the hardware and application levels. At the hardware level, high availability is usually achieved by deploying redundant components such as power supplies, I/O devices, and RAID storage.

Application availability is generally implemented through the use of clustering or load-balancing software that works in conjunction with the application. Moving these types of applications from a physical environment to a virtual environment may require special configuration that should be considered during the planning process. Note that ESX Server gives you additional options for creating a highly available environment.

### **Backup and disaster recovery**

Moving from a physical to a virtual environment may require changes to your backup and recovery practices. An Internet-based backup system can usually be left in place, unchanged; however, direct-attached tape systems generally require some reconfiguration.

You should also consider modifying your backup strategies to take advantage of the virtual architecture.

### **Management**

The planning stage is a good time to review your current management strategy and applications. Having a single system to manage your entire infrastructure – with both physical and virtual components – would certainly make life easier for your IT staff. Of course, this may mean replacing existing management systems and/or purchasing additional software, adding to your project costs.

### **Security**

Security is always a concern in an IT environment – protecting data and securing access to resources to only those that need them is of paramount importance.

VM access is enforced through user and group permissions; it may be possible to integrate this scheme into your existing directory or authentication services. Before deployment, take time to consider how you will control access to your VMs.

## Strategies and best practices

The following sections outline strategies and best practices for specific aspects of a virtual infrastructure:

- [Virtualization strategies](#) – outlines a strategy for load balancing, management, and I/O performance
- [Migration strategies](#) – provides a link to more information on HP ProLiant Essentials Server Migration Pack (SMP), which supports easy physical-to-virtual migration
- [Storage strategies](#) – describes the space needs of ESX Server and the use of SANs in a virtual infrastructure
- [Network strategies](#) – describes ESX Server networking, service console requirements, configuring VMotion technology, deploying NIC teams, and deploying VLANs
- [High Availability strategies](#) – outlines the deployment of clustering solutions in a virtual infrastructure
- [Disaster Recovery strategies](#) – outlines backup products and strategies for a virtual infrastructure
- [Management strategies](#) – outlines management products for a virtual infrastructure (including HP SIM, ProLiant Essentials Virtual Machine Management Pack (VMM), VMware VirtualCenter, and HP OpenView)
- [Security strategies](#) – describes authentication schemes for the remote console, setting permissions for access to ESX Server configuration files, configuring TCP/IP ports for management access, and strengthening overall security

# Virtualization strategies

This section provides strategies for load balancing, management, and I/O performance.

## Load balancing

When consolidating, it is important to obtain a balance across all virtual hosts. Non-resource-intensive infrastructure applications such as DNS, DHCP, and NFS can be combined into VMs (or a single VM, if appropriate) on a virtual host, while larger applications like Microsoft Exchange or Microsoft SQL Server VMs may exist on virtual hosts with fewer total VMs. This allows more headroom on the virtual host to accommodate spikes in utilization.

After your performance monitoring has provided an accurate representation of the virtual infrastructure workloads, these loads should be balanced across all virtual hosts, which is likely to be an ongoing activity. The HP ProLiant sizer for VMware ESX Server, described in [Appendix A – Using the HP ProLiant server sizer for VMware ESX Server](#), is an effective tool for helping you determine the initial distribution of workloads across servers.

## Management

Once virtualization becomes accepted within an organization, the number of VMs tends to grow rapidly. To handle this rapid growth, the appropriate management tools must be in place; tools such as HP Systems Insight Manager (SIM) and HP ProLiant Essentials VMM (Virtual Machine Management Pack) can be used to organize, provision, monitor, and manage the virtual infrastructure.

HP SIM allows VMs and hosts to be categorized in a variety of ways. Logical groups of VMs can be created and controlled independently; the groups can be based on user-defined parameters such as geographical location or department.

HP SIM ([HP Systems Insight Manager](#)) and VMM ([Virtual Machine Management Pack](#)) are discussed in greater detail later in this white paper.

## I/O performance

Virtualization hosting products are inherently challenged in the areas of disk storage and network I/O. To compensate, both subsystems must be enhanced – storage by deploying larger cache disk controllers or SANs, and network by running faster topologies, such as Gigabit Ethernet or Fibre Channel.

For more details on improving I/O performance, refer to, *Managing and operating a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management*. This guide can be found at: <http://h71019.www7.hp.com/ActiveAnswers/cache/71088-0-0-121.html>

## More information

- For performance tuning information on ESX Server see: [http://www.vmware.com/pdf/esx\\_performance\\_tips\\_tricks.pdf](http://www.vmware.com/pdf/esx_performance_tips_tricks.pdf)
- For detailed information on sizing memory for a virtual infrastructure see: [http://www.vmware.com/support/esx25/doc/admin/esx25admin\\_res-mem-sizing-intro.html](http://www.vmware.com/support/esx25/doc/admin/esx25admin_res-mem-sizing-intro.html)

## Migration strategies

After an appropriate virtual infrastructure has been designed, the next step is the migration of a physical server and its applications to a VM. This process is also known as Physical to Virtual (P2V).

SMP is a wizard-driven utility that allows easy P2V migrations. Most Windows operating systems are supported on SMP including Windows NT® 4.0, Windows 2000, and Windows Server 2003. For more information see the SMP user guide, available at [http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?locale=en\\_US&taskId=101&docIndexId=179166&contentType=SupportManual&prodSeriesId=453861&prodTypeId=0](http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?locale=en_US&taskId=101&docIndexId=179166&contentType=SupportManual&prodSeriesId=453861&prodTypeId=0).

For more details on using SMP, refer to, *Managing and operating a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management*. This guide can be found at: <http://h71019.www7.hp.com/ActiveAnswers/cache/71088-0-0-121.html>

## Storage strategies

This section outlines storage strategies in a virtual infrastructure.

### Space needs

The amount of storage required by ESX Server varies widely and is mostly determined by the space requirements of the hosted VMs; however, there are a few additional requirements that must be considered when evaluating overall storage needs. This section discusses these additional requirements and provides recommendations and best practices for configuring your storage.

#### **Service console**

The service console is a privileged VM that provides a management interface for ESX Server. It is based on a popular Linux distribution. As such, service console storage is configured similarly to that of any Linux server.

The service console must be configured with a /boot, / (root), and swap partition (which should be double the size of memory allocated to the service console, or a minimum of 1GB). Additional partitions for /var, /tmp, and /home may be created and require additional space but will generally not exceed 9GB.

For more details on service console storage, refer to, *Deploying a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management*. This guide can be found at: <http://h71019.www7.hp.com/ActiveAnswers/cache/71088-0-0-121.html>

#### **VMware File System (VMFS)**

Virtual machine disks are stored as files on one or more VMFS partitions. The size of the disk files is always equal to the storage capacity of the virtual disk and does not shrink (or grow) with actual use. Thus, if you have 5 VMs, each using 2 GB of a 4 GB virtual disk, your virtual disk files take up 20 GB of space rather than 10 GB.

VMFS is also used to store VM redo and suspend files. As a result, you should allocate additional space on your VMFS partitions if you plan to use these features.

Additionally, VMFS is used to create swap files for use by the VMkernel. At least one swap file equal to the size of the server RAM should be configured for each ESX Server.

### **Additional storage**

Additional partitions may be allocated to store floppy and CD-ROM images for use by VMs and VM templates. This allocation is commonly placed on a SAN or is made network-accessible so that it can be shared by all host servers and VMs to support the quick deployment of VMs and other software, reducing the need for physical media.

Other storage media can be mapped into a VM using Raw Device Mapping (RDM).

## **Storage Area Networks (SANs)**

The use of SANs in a virtual infrastructure provides greater flexibility in your storage configuration and virtualization strategy. A SAN allows you to share storage between multiple ESX Server hosts, offering common storage for virtual disks and templates and enabling features such as clustering and VMotion. A SAN may also provide better performance and high availability options.

This section discusses considerations for configuring a SAN in a virtual infrastructure.

### **VMFS partitions**

The most common use of SAN storage in a VMware environment is for VMFS partitions. In this configuration, the service console is installed on local storage in the server but the virtual disk files reside on the SAN. ESX Server permits sharing of a SAN LUN between multiple hosts, providing support for common VMFS partitions for virtual disk files and templates.

In order to use VMotion, a VM's virtual disk file must reside on a VMFS partition on the SAN that is shared between the source and target hosts.

---

#### **Note:**

HP recommends creating a local VMFS partition to store the VMkernel swap file, rather than storing this file on a SAN.

---

### **Number of VMs connected to a LUN**

The number of VMs that can be hosted on a single LUN on the SAN varies with the disk activity of the VMs themselves. Although as many as 100 is possible, it is recommended that no more than 32 active VMs be hosted on the same LUN. However, as few as 10, or even less, may be desirable based on your particular performance criteria. Connecting too many active VMs leads to increased latency and can cause the storage LUN to become too busy for VM operations.

### **Disks**

The number, size, and speed of the disks you will need for your SAN depend on the workloads presented to the VMs. Follow the sizing guidelines provided by your storage vendor based on the anticipated collective disk I/O activity of your VMs.

---

#### **Note:**

Sequential workloads become more random when combined with other activity on the same host.

---

## Clustering

When clustering VMs between physical servers, shared disk resources must reside on a SAN or in an RDM file. This is due to the manner in which SCSI reservations are handled in this configuration.

It is recommended that each shared disk resource reside on a separate LUN with a single VMFS partition. As a result, each resource can be reserved and released independently.

## Boot from SAN

VMware ESX Server 2.5, or later, supports booting from a SAN, eliminating the need for any local storage. The following list highlights the key considerations that must be met to enable ESX Server to boot from SAN:

- Each boot LUN must be masked so that it is only seen by its own ESX Server.
- The boot LUN must be presented to the lowest numbered HBA that has any LUNs presented to it.
- For active/passive arrays (such as the HP StorageWorks Modular Smart Array (MSA) and Enterprise Virtual Array (EVA) product families), the lowest-numbered path to the boot LUN must be the active path.

RDM is not supported in conjunction with the boot-from-SAN feature.

## High availability

Configuring your SAN for redundancy through use of multiple storage controllers, fabrics, and/or HBAs is recommended for high availability. ESX Server has native support for multipath I/O and neither requires nor supports the use of additional software such as HP StorageWorks Secure Path.

## Network strategies

This section discusses important topics to consider when planning network topology in a virtualized environment.

## ESX Server networking

Each physical NIC in a host server must be allocated to the service console, the VMkernel, or shared between both.

---

### Note:

Only NICs assigned to the VMkernel (or shared) can be used by VMs.

Although possible, HP does not recommend sharing the NIC between the service console and VMkernel.

---

Virtual switches are created on the host and attached to one or more physical NICs. The virtual adapter in the VM is then attached to a virtual switch that, in turn, provides external network access via the attached physical NIC. A virtual switch that is not attached to any of the physical NICs may be used to create a private network that can only be used by the VMs hosted on that server.

For more information, refer to, *Managing and operating a VMware Virtual Infrastructure with HP ProLiant servers, storage, and management*. This guide can be found at:

<http://h71019.www7.hp.com/ActiveAnswers/cache/71088-0-0-121.html>

## Service console

The service console is a privileged VM that provides a management interface for ESX Server. The service console provides access to many resources, including:

- Web-based Management User Interface (MUI)
- Remote console (keyboard, video, and mouse) to VMs
- VMware VirtualCenter communications
- HP management agents
- VMM

It is recommended that a dedicated NIC be allocated for the use of the service console. Isolating the service console network from the VM network keeps traffic on one from affecting the performance of the other.

Furthermore, this isolation can also add an additional layer of security. For example, a VM running a web server – typically located in a DMZ – may be susceptible to many kinds of malicious attacks. If the service console and VM are on the same network and the VM becomes compromised, the service console and, potentially, all other VMs are now open to attack. By maintaining the service console on a separate and, possibly, more secure network than that used for the VMs, the risk of a compromised host is reduced. Additional security topics are covered later ([Security strategies](#)) in this white paper.

While HP recommends that the service console be assigned a static IP address, the use of DHCP is acceptable if the DNS system is capable of dynamically updating host records. However, when sharing the service console NIC, a static IP address **must** be assigned.

## VMotion

Configuring and using VMotion technology, which permits the migration of a running VM from one host to another, requires some special consideration when planning your network.

When enabling a host for VMotion, it is recommended that a 1 Gb NIC be assigned to a “VMotion” virtual switch. No VMs should be attached to this virtual switch.

During a migration, the memory contents of the running VM are transferred over the “VMotion” network to the new host. Creating a dedicated network for this task helps ensure that this operation can complete in a timely manner and result in a successful migration.

If it is not possible to dedicate a NIC to use for VMotion, VMs may use the “VMotion” network. If possible, however, VMotion traffic should be isolated from the VMs through the use of VLANs. See the VLAN section ([VLAN](#)) for more details.

---

### Note:

It is not necessary to label the VMotion network, “VMotion”. It is so named here for illustrative purposes.

When creating the VMotion virtual switch, it is best to use the same name across all hosts to ease configuration and avoid confusion.

---

Other guidelines include:

- Each host enabled for VMotion requires a statically-assigned, unique IP address on the “VMotion” network.
- Although not a requirement, HP recommends isolating the VMotion network on a separate subnet or VLAN.
- The VMotion NICs must be in the same broadcast domain.

## Teaming

NIC teaming allows you to group two or more physical NICs into a single logical network device called a bond. The main reasons for using NIC teaming are as follows:

- Performance – load balancing network I/O
- High availability – providing network fault tolerance

NIC teaming in ESX Server supports the IEEE 802.3ad static link aggregation standard.

Outgoing traffic load balancing and fault tolerance are provided transparently by ESX Server and do not require special switch functionality or drivers for the guest OS. Inbound traffic can be load-balanced using 802.3ad-capable switching equipment.

Because a single NIC or switch port failure could potentially affect many VMs, fault tolerance should be a key concern when planning a virtual infrastructure. As a result, HP recommends the use of NIC teaming to provide highly available network services to your VMs.

## VLAN

A Virtual LAN (VLAN) is a logical group of network nodes that behave as if they are connected to the same wire even though they may be located on different physical LAN segments. Use of VLANs can lead to an increase in performance, improve manageability, and provide additional layers of security. VMware supports the IEEE 802.1Q VLAN standard and provides three modes for configuring VLANs when virtualizing with ESX Server. These modes are outlined below.

### Virtual Guest Tagging (VGT)

VGT mode requires that an 802.1Q trunking driver be installed in the guest OS to tag and untag frames. Frames are passed, unmodified, through the virtual switches between the VMs and the external switch, preserving the tagging information.

Because of the lack of hardware acceleration in the VM to perform the tagging, use of VGT is not recommended. However, there may be reasons to consider this mode:

- If a VM is required to be on five or more VLANs, then it may be necessary to use VGT.
- When using a P2V tool, such as the Server Migration Pack, on a server that is already running a 802.1Q trunking driver, VGT does not require any additional networking configuration on either the new VM or the host machine.

---

#### Note:

VGT is disabled by default. Additional configuration on the ESX host is required to enable VGT support. Please see the *ESX Server Administration Guide*.

---

## External Switch Tagging (EST)

EST mode relies on the external switch devices to perform frame tagging; no additional configuration on the ESX host is necessary. Use of this mode is not recommended due to its inflexibility – particularly when using port-based tagging on the switch.

Because there is only one-to-one mapping between the external switch port and virtual switch on the host server, using port-based VLANs will result in all VMs connected to the virtual switch being members of the same VLAN. Additionally, when using VMotion or moving a VM from one host to another, care must be taken to ensure that the virtual switch on the destination host is on the same VLAN(s) as the host virtual switch.

While MAC address-based tagging eliminates the problems just discussed, it also creates the potential for other problems. By default, the MAC address of a virtual NIC is auto-generated each time a VM is powered on and is not guaranteed to be the same – this obviously creates a problem if your VLANs are assigned by MAC address. You could configure the virtual NICs with static MAC addresses; however, this approach would require the assignments to be tracked to avoid duplication across the network.

## Virtual Switch Tagging (VST)

In VST mode, one port group is created for each VLAN on the virtual switch; virtual NICs are then attached to the port group rather than directly to the virtual switch. The virtual switch is responsible for tagging outbound frames and untagging inbound frames.

VST is the preferred mode for VMs because it is more flexible and easier to configure than EST, and provides better performance than VGT through the hardware acceleration available from the physical NIC for tagging frames. However, there are some limitations and considerations when using VST:

---

### Note:

Both VST and VGT cannot be used on the same host server. If VGT is enabled for a host, you will not be able to assign port groups to any virtual switches on that host. VST and EST are complementary; configuring your network switches to support VST varies with your switch vendor.

---

- Both VST and VGT cannot be used on the same host. If VGT is enabled, you cannot assign port groups to any virtual switches on that host.
- VST and EST are complementary.
- The configuration of network switches to support VST varies by switch vendor.
- A VM can be configured with up to four virtual NICs; each virtual NIC can be attached to only one port group. Thus, with VST, a VM can only be connected to a maximum of four VLANs. If a VM requires more than four VLANs, VST cannot be used.
- When using VMotion, make sure that the appropriate port groups have been configured on the destination host. Also ensure that external switches have also been configured properly. See the discussion on EST above for more information.

For more information on VLANs and ESX Server, please see the *ESX Server Administration Guide* and the [VMware ESX Server 802.1Q – VLAN Solutions](#) white paper.

## High Availability strategies

High Availability (HA) is a critical requirement in any consolidation exercise, particularly if the project involves migrating from a large number of smaller servers to a smaller number of large servers. In this case, the impact of server downtime increases with more users and applications being placed on each server.

HP offers highly available, fault-resilient solutions that directly address concerns about increased data and application vulnerability that may result from IT consolidation. These solutions are designed to maximize performance and storage capacity, while providing a 24 x 7 platform for business-critical applications.

This section outlines highly available clustering options for your virtual infrastructure.

### HP clustering solutions

HP offers a wide range of industry-standard clustering solutions, allowing you to match the processing requirements of your applications to the levels of availability and performance needed.

HP provides a broad range of configurations for Microsoft Cluster Services (MSCS) running on Microsoft Windows NT 4, Windows 2000 or Windows Server 2003; many HP ProLiant servers are certified with a choice of HP Smart Array Cluster Storage, or HP StorageWorks MSA, EVA or XP storage systems, in both single- and dual-path configurations.

#### Limitations

Currently, ESX Server only supports a two-node cluster configuration, while the current implementation of Windows Server 2003, Enterprise Edition supports up to eight nodes in a cluster.

It is possible to configure multiple two-node clusters in a virtual infrastructure, where each set of two VMs has access to the same disk file. The desirability of such a configuration would have to be weighed against business requirements.

### VM-based clustering

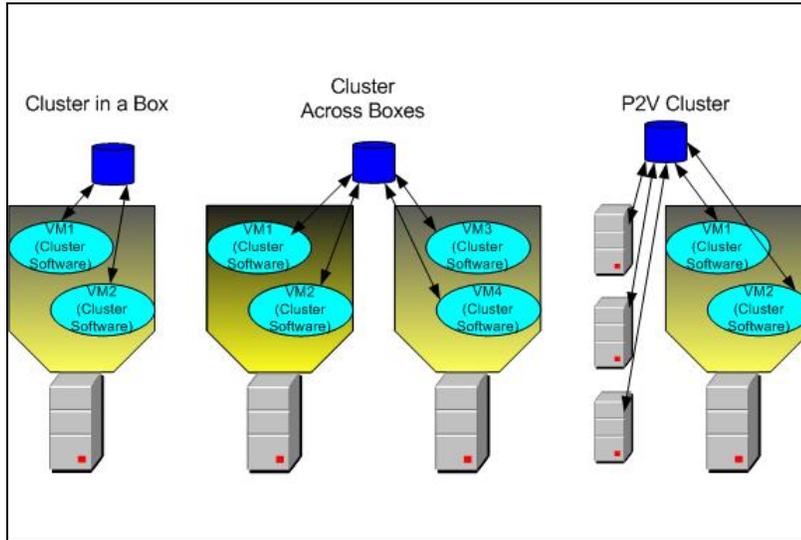
To implement high availability, you can deploy software to bind two redundant servers into a primary/standby pair. With a cluster-aware application, such an arrangement results in little or no downtime when the primary server experiences a hardware or software error. The redundancy offered by a cluster supports the elimination of single points of failure.

With virtualization, you are able to create a cluster between a physical machine running a mission-critical workload and a similarly-configured VM. In this case, the VM does not consume computing resources in stand-by mode and can be consolidated to one or a few physical platforms at a very high consolidation ratio. As a result, you can improve availability without investing in twice the amount of hardware or managing and patching a sprawl of servers. Redundancy is reduced from 2N to N+1.

P2V clustering supports the same clustering software as physical-to-physical machine clustering, greatly reducing training requirements for your IT staff. At the same time, reduced cost allows you to implement HA for more workloads.

Figure 1 outlines clustering options

**Figure 1:** Clustering options



In a typical VM cluster, each VM represents one node; disks are typically shared between nodes. Shared disks are required for applications using dynamic data, such as mail or database servers.

**Note:**

When a virtual disk is created, ESX Server pre-allocates the disk space.

There should be extra network connections between nodes for monitoring heartbeat status and a method for redirecting incoming requests.

**IMPORTANT:**

Always rigorously test and review the cluster before deploying it in a production environment.

**Network load balancing**

Network load balancing can be achieved using a cluster; for example, in a web server environment, a gateway (load balancer) distributes requests to all cluster nodes according to load and can also redirect requests to the remaining nodes if one crashes. This configuration increases availability and performance over a single-machine approach.

Services such as load balancing are easy to implement in a virtual environment. Since VMs can be replicated with relative ease, adding or removing web server VMs (as in the above example) to a load balancer is as simple as turning them on or off. The load balancer typically has the intelligence to add active VMs or remove powered-off VMs.

The benefits of network load balancing include:

- Enhance the availability of Internet server applications, such as those used on these types of servers:
  - Web
  - Proxy
  - Domain Name Service (DNS)
  - FTP
  - VPN
  - Streaming media servers
  - Terminal Services
- Scale the server's performance.
- Create the cluster with VMs on the same physical server or with VMs on multiple physical servers
- Configure up to 32 network nodes in the cluster

### **Cluster-aware applications**

To take full advantage of clustering services, applications should be cluster-aware. Such applications are typically stateless<sup>1</sup>, such as web servers or VPN servers. Cluster-aware applications often include built-in recovery features, like those in database servers, mail servers, file servers, or print servers.

### **Clustering software for VMs**

The following software solutions were designed for physical platforms but should run equally well in a virtual infrastructure; however, some configuration parameters (such as SCSI reservation and file locking) may differ.

Clustering software solutions include:

- Microsoft Clustering Service (MSCS), which provides fail-over support for two- to four-node clusters for applications such as databases, file servers, and mail servers
- Microsoft Network Load Balancing (NLB), which balances the load of incoming IP traffic across a cluster of up to 32 nodes for applications such as web servers and Terminal Services.
- VERITAS Cluster Server (VCS)
- Novell Cluster Service

---

#### **Note:**

These clustering services are tested and supported by VMware only with Windows host operating systems.

While later, more advanced versions of Windows support larger, multi-node clusters, ESX Server currently only supports two-node clusters.

---

### **Creating a cluster in a box**

With ESX Server, you can create a simple cluster within a single physical server to help mitigate the effects of software crashes or administrative problems.

The characteristics of this type of cluster include:

- Supports shared disks without any shared SCSI hardware
- Supports a heartbeat network without an extra network adapter

There may be multiple two-node VM clusters on a single physical server.

---

<sup>1</sup> As opposed to job-based applications

## Multipathing in ESX Server

ESX Server 2.5.x includes multi-pathing support to help maintain a constant connection between the host server and the storage device in case of the failure of a Host Bus Adapter (HBA), switch, storage controller, or Fibre Channel cable. For this reason, ESX Server does not need any additional multi-pathing software.

For more information on supported SAN configurations, see:

[http://www.vmware.com/partners/hw/hp.html/esx\\_SAN\\_guide.pdf](http://www.vmware.com/partners/hw/hp.html/esx_SAN_guide.pdf).

## Disaster Recovery strategies

Disaster Recovery (DR) typically involves a series of plans and processes aimed at completely restoring failed or impeded operations, or preventing failures from occurring.

Some organizations take the limited view that a disaster is an outage due to flood, earthquake, power failure, or something similar. However, in today's high-speed, always on IT world, a disaster is typically defined as the interruption or limitation of any process that supports the business operations of an enterprise. In short, if it adversely affects your business, it is a disaster.

---

### Note:

The implementation of a DR solution should strike a balance between the risk of disaster and the vulnerability of your business if a disaster were to occur.

---

A virtual infrastructure offers several key advantages in the area of DR:

- Encapsulating an operating environment into a few files supports easy duplication, backup and restore, and management.
- VMs can be recovered independent of the hardware; recovery or redundancy can even be relegated to smaller and fewer physical servers.

While a comprehensive explanation of DR is beyond the scope of this white paper, the following sections provide an overview of DR and its relationship with a virtual infrastructure.

---

### Note:

The differences between HA and DR are sometimes blurred, with some considering HA to be a subset of DR. For the purpose of this white paper, however, these topics are considered separate. HA is primarily concerned with maintaining a constant service level to users, while DR is primarily focused on the recovery process once failure has occurred.

---

## Backup strategies and products

This section outlines backup strategies and products that can facilitate backups in a virtual infrastructure.

Currently, the following HP OpenView products are suggested for backup in a virtualized environment using ESX Server: HP OpenView Storage Data Protector and HP OpenView Storage Mirroring.

Other HP StorageWorks products are being tested with ESX Server. For more information, visit <http://h18006.www1.hp.com/storage/index.html>.

### **HP OpenView Storage Data Protector**

HP OpenView Storage Data Protector manages backup and recovery from both disks and tapes. This software centralizes backup and recovery operations integrating a variety of techniques to eliminate backup windows; these techniques include on-line backup, open file backup, and instant recovery or zero-downtime backups.

Data Protector contains instant recovery features and several other integrated DR alternatives.

Data Protector simplifies the use of complex backup and recovery procedures with a fast installation process, automated routine tasks, and easy-to-use features. The user interface simplifies the backup of VM disk files and their associated configuration files.

### **HP OpenView Storage Mirroring**

HP OpenView Storage Mirroring (OVSM) offers a host-based application that performs remote copy over an IP LAN/WAN. This application operates on a WinTel server with Microsoft Windows NT/2000/2003 operating systems.

Key OVSM features include:

- Asynchronous replication that can be scheduled to a fine granular level – LUN-, file-, or byte-level
- Multiple replication configuration options including peer-to-peer and many-to-one

OVSM offers a very cost-effective DR alternative in a number of scenarios: for example, from one host to another within a LAN or storage center or direct attached storage, or between metropolitan offices and regional centers. OVSM capabilities also include replication from small office environments.

OVSM is an ideal entry-level, host-based solution for IP networks since it does not require high-bandwidth Fibre Channel networks, high-capacity replication, and zero-down-time service levels. OVSM provides near real-time full application or file recovery with up to the last byte replication, meeting business recovery goals within minutes or hours. With its low initial investment costs compared to alternative storage-based and fabric-based replication products, OVSM is an excellent choice for low bandwidth, low storage volume changes.

OVSM and the virtual infrastructure

OVSM does not provide any specific functionality for the virtual infrastructure: it responds to VMs in the same way it responds to physical machines.

Since OVSM does not distinguish between the partition in a VM and the partition in a physical machine, failover can be accomplished between two VMs on two different host servers; due to the synchronization capabilities of OVSM, this failover can be performed from room to room on a local site, or across the country.

### **Backup/failover verification**

Many popular backup software solutions (such as Legato NetWorker and VERITAS NetBackup) have rigid failover verification processes that require a test failover server, an installed operating system, the backup agent, and adjustments to the Windows registry, as well as other system configuration changes. Since a virtual infrastructure makes failover to dissimilar servers much easier, it is often beneficial to use a VM rather than a physical server to act as a virtualized failover server or to verify the integrity of the backup process.

Furthermore, the process of creating a virtualized failover server, installing an operating system, backup agent, and making adjustments to the configuration need only be performed once. The resulting VM can then typically be copied or cloned and modified for other failover verification scenarios.

## **Management strategies**

This section outlines management products for a virtual infrastructure (including HP SIM, VMM, VirtualCenter, and HP OpenView).

### **HP Systems Insight Manager**

HP Systems Insight Manager (HP SIM) is a management application that assists IT staff in managing all HP servers and system hardware within your IT environment. Regardless of the size or complexity of your organization, HP SIM can help make you more efficient and proactive in identifying, diagnosing, and fixing potential issues on all HP hardware. Furthermore, HP SIM can increase productivity by providing inventory management, event management, and remote management, as well as role-based security.

#### **Requirements in a Windows environment**

Table 1 lists the hardware, software, and networking requirements for deploying HP SIM in a Windows environment. These requirements are separated into the three components of an HP SIM environment: Central Management Server (CMS), managed systems, and network clients.

**Table 1:** Requirements for installing HP SIM in a Windows environment

Server (CMS)			
Operating system	Hardware	Software	Networking
<ul style="list-style-type: none"> <li>Windows 2000 Server with Service Pack 4</li> <li>Windows 2000 Advanced Server with Service Pack 4</li> <li>Windows XP Professional with Service Pack 1 or later</li> <li>Windows Server 2003 Standard Edition</li> <li>Windows Server 2003 Enterprise Edition</li> </ul> <p><b>Note:</b> International – French, German, Spanish, and Japanese with the latest Service Pack are also supported</p>	<ul style="list-style-type: none"> <li>768 MB RAM (1 GB recommended)</li> <li>Minimum single CPU 1.5 GHz (2.4 GHz or greater recommended)</li> <li>500 MB free disk space</li> </ul>	<ul style="list-style-type: none"> <li>Microsoft Data Engine 2000 (MSDE) with Service Pack 3 or later or</li> <li>Microsoft SQL Server 2000 with Service Pack 3 or later</li> </ul>	<ul style="list-style-type: none"> <li>TCP/IP installed</li> <li>SNMP services installed and active</li> <li>Domain Name Services (DNS) server available in environment</li> </ul>
Managed system			
Operating system	Hardware	Software	Networking
<ul style="list-style-type: none"> <li>Microsoft BackOffice Small Business Server</li> <li>Small Business Server 2000</li> <li>Windows 2000 Server</li> <li>Windows 2000 Advanced Server</li> <li>Windows 2000 Professional</li> <li>Windows NT Server 4.0</li> <li>Windows Server 2003</li> <li>Windows Server 2003 Enterprise Edition</li> <li>Windows XP Professional</li> </ul>	<ul style="list-style-type: none"> <li>Any HP IA-32 server or</li> <li>Any HP IA-64 server</li> </ul>	<ul style="list-style-type: none"> <li>HP ProLiant Support Pack 6.30 or later</li> <li>OpenSSH 3.7.1 (optional)</li> </ul>	<ul style="list-style-type: none"> <li>TCP/IP installed</li> <li>SNMP services installed and active</li> </ul>
Network client			
Operating system	Web browser		
<ul style="list-style-type: none"> <li>Windows</li> </ul>	<ul style="list-style-type: none"> <li>Microsoft Internet Explorer 6.0 or later with Java™ JRE browser plug-in 1.4.2.03 or later</li> </ul>		

**Note:**

HP SIM can also be used to manage Linux servers.

For more information on HP SIM, see: <http://www.hp.com/servers/manage>.

# Virtual Machine Management Pack

Integrated with HP SIM, HP ProLiant Essentials Virtual Machine Management Pack (VMM) delivers centralized management and control capabilities for VMs and supporting HP ProLiant host server resources. This integrated approach provides the ability to manage both physical and virtual resources from a single management console.

VMM offers the following benefits:

- Integration with the HP SIM console to manage the VM environment
- Simplified deployment and operation of VMs
- Reduced cost and complexity for server consolidation projects
- Faster response times to your changing business demands

VMM provides tracking, monitoring, and control functions to help your IT staff organize an effective virtualized environment.

## Infrastructure

The following components are set up during VMM installation:

- **Virtual Machine Management Service**  
This service resides on the HP SIM server and controls the internal functions of VMM.
- **Virtual Machine Management Pack Console**  
This VMM user interface provides access to VM monitoring and control functionality and can be accessed locally or remotely, using an industry-standard Web browser.
- **Virtual Machine Management Agent**  
This agent can be distributed to managed hosts through the user interface.
- **HP ProLiant Essentials Server Migration Pack (SMP)**  
Through a console that is integrated with the VMM console, SMP can perform P2V or virtual-to-virtual (V2V) migration of managed VMs.

---

### Note:

SMP has additional licensing requirements. For additional information, refer to the SMP user guide.

---

## Requirements

A VMM environment consists of the following components:

- **HP SIM CMS**  
Independent of requirements for HP SIM, the CMS must meet the following requirements to successfully use VMM:
  - HP SIM must be installed on a Windows physical server; VMM does not support HP SIM operating on ESX Server or on a server running Linux or HP-UX
  - HP SIM 4.2 Service Pack 1 or later with OpenSSH and Windows Management Instrumentation (WMI) Mapper installed
  - At least 155 MB of available disk space.
  - Network connections configured appropriately for management by VMM; verify the configuration by pinging the host server from the CMS and vice versa using the primary IP address
- **VM hosts**
- **VM guests**

# VMware VirtualCenter

VirtualCenter (VC) is a management suite from VMware that offers services such as management, monitoring, and resource utilization tuning for both VMs and host servers. VC provides templates and VMotion technology to support the rapid provisioning of VMs.

For more information on VC, visit [http://www.vmware.com/products/vmanage/vc\\_features.html](http://www.vmware.com/products/vmanage/vc_features.html).

## Using the VC interface

The VC interface provides an overview of all virtual resources in your data center. From this interface, IT staff can access the following functionality:

- **Continuous workload consolidation**

By adjusting the resources dedicated to each software service, IT staff can monitor and optimize the utilization of data center resources to minimize unused capacity while maintaining application service levels.

- **Instant provisioning**

The time taken to provision servers can be reduced to tens of seconds, allowing IT staff to respond immediately to requests for services. Using server templates, IT staff can ensure that new servers are fully consistent with the current build and security policies.

- **Zero-downtime maintenance**

You can safeguard your business continuity, with no service interruptions for hardware maintenance, deployment, or migration. IT staff use VMotion to move running operating systems and their applications – without service interruption – from a managed host that needs maintenance to another managed host, then move them back when maintenance is complete.

## Software components

VC includes the following software components:

- **VC server**

At the heart of VC is the VC server, a Windows service that acts as a central administrator for host servers connected on a network, directing actions on both hosts and VMs. The VC server collects and stores per-system and environmental information and can automatically execute user-specified scheduled tasks, such as starting up or moving a shut-down VM.

The VC server runs full-time; it must have network access to the service consoles of all managed hosts and must be available for network access from any client machine.

See below for VC server machine requirements.

- **VC database**

The VC database provides persistent storage area for information on VMs, host, and users managed in the VC environment. The VC database can be local to or remote from the VC server machine.

- **VC client**

The VC client is a user interface that runs locally on a Windows machine that has network access to the VC server. The VC client can run on the same machine as the VC server or on another machine with appropriate network access.

The VC client requires a computer monitor to provide access to the user interface.

See below for VC client machine requirements.

- **VC web service**

The optional VC web service can be installed with the VC server.

Note that this web service is a requirement for third-party applications that use the VMware SDK Application Programming Interface (API).

See below for VC web service machine requirements.

- **VC agent**

Installed on each managed host, the VC agent collects, communicates, and executes the actions received from the VC server. It is installed automatically the first time a host is added to the VC inventory.

To install the VM agent server, install and run the following:

- A version of Windows supported for VC server (see below)
- VirtualCenter 1.2

- **VMotion**

Centrally coordinated by the VC server, VMotion technology allows running VMs from one host to be migrated to another host without service interruption.

VMotion requires licensing on both source and target hosts.

### **VC server requirements**

The VC server machine must meet the following requirements:

- Administrator privileges – required for the installation of VC server
- Windows 2000 Server, Windows 2000 Advanced Server, Windows XP Professional, or Windows Server 2003 (Web, Standard, or Enterprise Edition).
- For 50 or fewer managed hosts: a minimum of 2 GB RAM  
For more than 50 managed hosts: a minimum of 3 GB RAM  
For 100 managed hosts running 2000 virtual machines: a minimum of 4 GB RAM
- As a minimum, an Intel® Pentium® 4 2.0 GHz processor; however, dual processors are recommended for deployments with more than 25 managed hosts
- As a minimum, one 10/100 Mbps NIC; one 1 Gbps NIC recommended
- Windows Script 5.6, or later; if this particular version is not present on the server, the VC installer automatically updates the older version to 5.6.
- Disk space sufficient on the machine to support the VC database and the template upload directory

The VC server may run on the same machine as the VC client or, separately, on another Windows system. Although this is not recommended, the VC server can even be installed on a VM.

### **VC client requirements**

VC client machines must meet the following requirements:

- Microsoft .NET Framework 1.1; if an older version is present, it is automatically updated to .NET Framework 1.1.4322.573
- Windows 2000 (all versions), Windows XP Professional, Windows XP Home Edition, Windows Server 2003 (all editions), or Windows NT 4 (SP6a required).
- A minimum of 256 MB RAM (512 MB recommended).

The VC client can be installed on multiple Windows systems; each system accesses the VC server over the network. These client systems can be desktops, laptops, or VMs.

### **VMware web service requirements**

The machine must meet the minimum hardware requirements for the VC server machine (see above).

## VC networking requirements

Networking requirements for VMs and managed hosts are as follows:

- **VM**

- Up to four virtual Ethernet NICs; each may be a high-performance VMware virtual NIC or an AMD™ PCnet-PCI II-compatible virtual NIC
- Support for any protocol that the guest operating system supports over the Ethernet; multiple high-performance, Ethernet-compatible virtual networks are possible

- **Managed host**

- The minimum number of NICs is two; how the NICs are assigned depends upon the version of the VMware virtualization platform being used
- The preferred number of NICs is three: one dedicated to the managed host, one (or more) dedicated to the VMs, and one dedicated to VMotion activity

## HP OpenView

HP OpenView applications allow you to increase the performance of your IT infrastructure, anticipate and correct problems before they become critical, and automate and manage change in real time. Following the principles of simplification, standardization, and modularity, HP OpenView applications offer you a unique vision and proven results that directly impact the bottom line. HP OpenView applications enable an Adaptive Enterprise.

Benefits include:

- Focus of IT organization moved from being reactive to proactive, and towards being a valued partner of the business
- Availability and performance of critical business services managed across the enterprise
- Business processes linked to IT services
- Windows infrastructure and Microsoft applications brought under control
- Comprehensive management across all IT resources (networks, systems, applications, middleware, databases, and storage)
- IT service levels and quality maximized

This multi-platform solution allows you to manage a heterogeneous environment<sup>2</sup> and optimize service quality by monitoring and measuring the availability and performance of each element in your infrastructure. You can convert the information you have collected into actionable insight, so that the most urgent management problem can be solved first.

The depth of HP OpenView management solutions; the end-to-end, modular approach; ease of deployment and administration; and optimal customer experience combine to ensure a quick return on investment.

---

<sup>2</sup> Heterogeneous systems and applications, including networks, storage, Windows, UNIX®, Linux, Novell NetWare, Oracle, SAP, and more

## Support

HP offers a range of support options for HP OpenView, including:

- Access to a wide team of support engineers
- Over 35 response centers worldwide provide local language support
- Proactive, reactive, mission-critical, and online support services, plus a broad family of integrated solutions partners

## Security strategies

A security implementation in a virtual infrastructure should address the following considerations:

- ESX Server authenticates all remote users attempting to connect to a server using the VMware Management Interface or the VMware Remote Console.
- Security for network traffic to and from a host server depends on the security settings specified in the server configuration.
- Three or more TCP/IP ports may be used for access. Depending on remote access requirements, the firewall should be configured to allow access to one or more of these ports.

This section outlines strategies for addressing the above considerations.

### Authenticating users

ESX Server uses Pluggable Authentication Modules (PAMs) for user authentication in the remote console and the VMware Management Interface. The PAM configuration is in `/etc/pam.d/vmware-authd`.

The default installation of ESX Server uses `/etc/passwd` authentication, just as Linux does, but can easily be configured to use LDAP, NIS, Kerberos, or some other distributed authentication mechanism.

Every time a connection is made to the host server, the `inetd` process runs an instance of the VMware authentication daemon (`vmware-authd`), which requests a user name and password, then hands them off to PAM to perform the authentication.

Once a user has been authenticated, `vmware-authd` accepts a path name to a VM configuration file. Access to this configuration file is restricted as shown in Table 2.

**Table 2:** Access requirements

Read	Access		Action
	Write	Execute	
Yes			Seeing and controlling the VM in the VMware Management Interface
Yes			Viewing VM details pages
Yes			Using the local service console or connecting to the VM with the VMware Perl API
Yes		Yes	Connecting to and controlling (start, stop, reset, or suspend) a VM in a remote console using the VMware Perl API or VMware Management Interface
Yes	Yes		Changing the configuration using the Configure VM page in the VMware Management Interface

---

**Note:**

Users with list access – but not read access – may encounter errors in the VMware Management Interface.

---

If a VMware process is not running for the configuration file in use, `vmware-authd` examines `/etc/vmware/vm-list`, the file where the VMs are registered. If the configuration file is listed in `vm-list`, `vmware-authd` (not necessarily the user that is currently authenticated) starts ESX Server as owner of this configuration file.

Registered VMs (those listed in `/etc/vmware/vm-list`) also appear in the VMware Management Interface. The VMs shown on the Status Monitor page of the MUI must be listed in `vm list`; IT staff must have read access to their configuration files.

The `vmware-authd` process exits as soon as a connection to a VMware process is established. Each VMware process shuts down automatically after the last user disconnects.

**Using certificates to secure remote sessions**

When using the VMware Remote Console or the VMware Management Interface, the username, password, and network packets sent to a host server over a network connection are, by default, encrypted in ESX Server<sup>3</sup>.

With SSL enabled, security certificates are created by ESX Server and stored on the server. However, since the certificates used to secure the management interface are not signed by a trusted certificate authority, they do not provide authentication. If encrypted remote connections are used externally, consider purchasing a certificate from a trusted certificate authority. Another option is to use of an in-house developed security certificate for SSL connections.

The management interface certificate must be placed in `/etc/vmware-mui/ssl`. This certificate consists of two files: the certificate itself (`mui.crt`) and the private key file (`mui.key`), which should be readable only by the root user.

When upgrading the management interface, the certificate remains in place and, in case the management interface has been removed, the directory is not removed from the service console.

**Default permissions**

When creating a VM with ESX Server, its configuration file is registered with the following default permissions, based on the user accessing it:

- Read, execute and write — for the user who created the configuration file (the owner)
- Read and execute — for the owner's group
- Read — for users other than the owner or a member of the owner's group

---

<sup>3</sup> If Medium or High security settings were specified for the server

### **TCP/IP ports for management access**

The TCP/IP ports available for management access to the host machine vary, depending on the security settings chosen for that server.

If you need to manage host machines from outside a firewall, the firewall needs to be reconfigured to allow access on the appropriate ports. The following list below shows which ports are available based on the standard security setting selected:

- **High Security**

- 443 - HTTPS, used by the VMware Management Interface
- 902 - vmware-authd, used when you connect with the remote console
- 22 - SSH, used for a secure shell connection to the service console

- **Medium Security**

- 443 - HTTPS, used by the VMware Management Interface
- 902 - vmware-authd, used when you connect with the remote console
- 22 - SSH, used for a secure shell connection to the service console
- 23 - Telnet, used for an insecure shell connection to the service console
- 21 - FTP, used for transferring files to and from other machines
- 111 - portmap, used by the NFS client when mounting a drive on a remote machine

- **Low Security**

- 80 - HTTP, used by the VMware Management Interface
- 902 - vmware-authd, used when you connect with the remote console
- 22 - SSH, used for a secure shell connection to the service console
- 23 - Telnet, used for an insecure shell connection to the service console
- 21 - FTP, used for transferring files to and from other machines
- 111 - portmap, used by the NFS client when mounting a drive on a remote machine

The key ports used by the VMware Management Interface and the VMware Remote Console are the HTTP or HTTPS port and the port used by `vmware-authd`. The use of other ports is optional.

## **Strengthening security**

Security is integral to the ESX Server architecture, with its special isolation features and network security functionality. However, even with these capabilities, a machine is only as secure as its configuration. Here are some best practices that can maximize security in a virtualized infrastructure.

### **Trusted users only in the service console**

Since the service console has privileged access to certain areas of ESX Server, you should only give trusted users login access to this console.

In addition “root” access should be limited. Specifically, VMware recommends that SSH capability to login directly as root be restricted. ESX Server system administrators should be required to login as a regular user and then switch user (`su`) to root.

### **Do not configure Promiscuous mode NICs**

It is possible to configure virtual NICs to run in Promiscuous mode, which can enable a VM to “sniff” packets destined for other VMs just as on a physical hub. For maximum security, Promiscuous mode NICs should not be enabled. Note that they are disabled by default.

### **Consider disabling VM logging**

VMs can log troubleshooting information into a log file stored in the service console. Normal – non-root or non-administrator – users and VM processes can abuse the logging capability and cause large amounts of data to be logged. Over time, a log file can consume the file system space designated for the service console and cause a denial of service attack.

### **Disable copy-and-paste in the VM**

When VMware Tools is running on a VM, it is possible to copy and paste from that VM.

A privileged user is able to log into a VM using the remote console; however, it is possible that this user is logging in from a non-privileged account on the client machine. The remote console user's clipboard would be accessible to this non-privileged account, which could access to a privileged account as soon as the console window becomes the top most window displayed.

## **Additional security issues**

You should also consider the following security issues when securing host machines and VMs.

### **Host machine**

A key issue is the security of the disk images and settings files for the VMs. This information should be stored on a NTFS disk so that NTFS permissions can be used to restrict access, preventing users from deleting or otherwise harming these files.

Furthermore, deploying an NTFS disk can also enhance performance.

### **VMs**

An important point to remember is that there is nothing extraordinary about a VM that can excuse IT staff from securing it. If accessing the network through the host machine, a VM can do anything that a physical machine could do on that network: this means that a VM **must** follow all normal security practices and guidelines, securing the VM to the same standards used to secure a physical machine.

With a host machine on home network, this means installing anti-virus software on the VM, consider firewalling the VM and remember to apply operating system and application updates.

With a host machine on a corporate network, ensure that VMs comply with your network and security policies. Manage rights carefully; giving a particular user administrator rights on a VM is risky if you would not give this user the same rights on a physical machine. As with the machine on a home network, install anti-virus software and apply operating system and application updates.

# Appendix A – Using the HP ProLiant server sizer for VMware ESX Server

The HP ProLiant server sizer for VMware ESX Server is an automated tool that helps you estimate the size and scope of a server environment supporting ESX Server. The sizer calculates the best way to consolidate your current physical machines on to new target machines and then generates a bill of materials for the new hardware. The printable view option gives a detailed report of hardware specifications as well as a chart that reports server resource utilizations.

The sizer can be downloaded from: [HP ProLiant server sizer for VMware ESX Server](#).

## Using the sizer

The sizer requires you to provide input in three areas:

- **Max Rates**

The Max Rates section allows you to specify the maximum utilization rates for the target ESX Servers. You can specify additional capacity to accommodate growth and/or occasional spikes in resource usage.

- **Servers to Consolidate**

In the Servers to Consolidate section, you enter detailed information about the current servers you wish to consolidate. The required data can be collected using HP ProLiant Essentials Performance Management Pack ([Appendix B – HP ProLiant Essentials Performance Management Pack](#)), AOG CapacityPlanner ([Appendix C – Using AOG CapacityPlanner](#)), Windows Performance Monitor (Perfmon) ([Appendix D – Using Microsoft Windows Performance Monitor](#)), or other such tools. To facilitate data entry, you can cut and paste from a Microsoft Excel spreadsheet.

- **Platform Selection**

In the Platform Selection section, you can specify up to four platforms as targets for your consolidation. The tool performs a separate sizing for each platform and presents each as a solution for comparison.

Detailed descriptions of all required user inputs are given below.

## Max Rates

**Figure 2:** Setting the maximum desired utilization rates

### Maximum Utilization Rates

Max Rates Servers to Consolidate Platform Selection

Specify the maximum utilization rates desired for the target servers that will run the virtualized environments.

**Maximum CPU Utilization**  
This specifies the maximum CPU utilization rate for the target server. The average CPU utilization for allocated VMs plus the calculated ESX Server overhead will not exceed the percentage assigned here.

75 %

**Maximum Memory Utilization**  
This specifies the maximum memory utilization for the target server. The RAM assigned to all VMs plus the calculated ESX Server overhead will not exceed the percentage of configured RAM assigned here.

75 %

#### Maximum CPU Utilization

You specify the maximum desired CPU utilization rate for the target server. The calculated ESX Server overhead plus the average CPU utilization of the VMs will not exceed the percentage of CPU utilization assigned here.

The default is 75%.

#### Maximum Memory Utilization

You specify the maximum desired memory utilization for the target server. The calculated ESX Server overhead plus the total RAM assigned to the VMs will not exceed the percentage of configured RAM assigned here.

The default is 75%.

#### Maximum Disk Utilization

You specify the maximum desired disk operations for the target server. The calculated ESX Server overhead plus the disk operations performed by the VMs will not exceed the percentage of disk operations assigned here.

The default is 75%.

#### Maximum Network Utilization

You specify the maximum desired network throughput for the target server. The calculated ESX Server overhead plus the network throughput for all VMs will not exceed the percentage of network throughput assigned here.

The default is 75%

### Additional Disk Space (GB)

You specify the amount of additional disk space to configure on the target server. This will be added to the disk space required for all VMs plus the ESX Server service console.

The default is 75%

### Safe Calculations

Checking the Safe Calculations box causes the sizer to become more conservative in its analysis, making it less likely that the target servers will exceed the specified maximums.

The default is 75%

## Servers to Consolidate

**Figure 3:** Adding current physical server usage rates

Physical Server					Application								Preferences	
Current Server Name	Current Server Model	CPU	# of CPUs	CPU Speed (MHz)	OS version	App Disk Space (GB)	% CPU Utilization	Max RAM Usage (MB)	Avg Disk Throughput (OPS)	Max Disk Throughput (IOPS)	Avg Network Throughput (MBps)	Max Network Throughput (MBps)	Min RAID Level	SMP

### Physical Server

**Current Server Name** – This is a label used to identify the server to be consolidated. The label can be any alphanumeric string but is generally the name of the physical server. Although useful, the label is not required to be unique.

**Current Server Model** – This is another label to help identify the server. The label can be any alphanumeric string but is generally the model of the physical server.

**CPU** – This is the type of CPU deployed in the server. The available options are:

**Figure 4:** Adding CPU information

Pentium	Pentium III	Xeon MP
Pentium Pro	Xeon™	AMD Opteron™
Pentium II	Xeon DP	[AMD Model]

**# of CPUs** – This is the number of CPUs in the physical server.

**CPU Speed (MHz)** – This is the speed of the CPUs in the server, given in MHz. For example, for a 1.4 GHz CPU, you enter “1400.”

## Application

**OS Version** – This is the name of the operating system running on the server. Possible entries are:

**Figure 5:** Adding OS version

Windows 2003	Windows NT 4.0	SuSE Linux
Windows XP	RedHat Linux	Netware
Windows 2000		

**App Disk Space (GB)** – This is amount of disk space (in GB) to assign to the VM. This is **not** the size of the disks in the physical server, nor is it the total amount of disk space used; rather, it specifies the size of the specific DSK file to be allocated to the VM.

**% CPU Utilization** – This is the CPU utilization on the physical server. In many cases, it is sufficient to use the true average CPU utilization; however, in an environment where the server may experience prolonged increases in utilization (such as overnight batch processing), you may prefer to use the average utilization during these peak times.

**Max RAM Usage (MB)** – This is the maximum amount of RAM (in MB) used by the server. However, if you would prefer to guarantee a specific amount of RAM to a VM, enter that amount instead.

**Avg Disk Throughput (IOPS)/Max Disk Throughput (IOPS)** – These are the average and maximum disk I/O operations per second. Typically, it is sufficient to use true average throughput values; however, in an environment where the server may experience prolonged increases in utilization (such as overnight batch processing), you may prefer to use the average utilization during these peak times.

For the maximum throughput, use a value in the 90<sup>th</sup> – 95<sup>th</sup> percentile rather than the true maximum value. Maximum values are usually the result of unusual spikes and do not represent normal server activity; as such, using a true maximum value may result in an ultra-conservative sizing.

**Avg Network Throughput (MBytes/s)/Max Network Throughput (MBytes/s)** – These are the average and maximum network throughput values (in MBytes/sec). Typically, it is sufficient to use true average throughput values; however, in an environment where the server may experience prolonged increases in utilization (such as overnight batch processing), you may prefer to use the average utilization during these peak times.

For the maximum throughput, use a value in the 90<sup>th</sup> – 95<sup>th</sup> percentile rather than the true maximum value. Maximum values are usually the result of unusual spikes and do not represent normal server activity; as such, using a true maximum value may result in an ultra-conservative sizing

## Preferences

**Min RAID Level** – This is the minimum RAID level desired for the target server and is used by the sizer when calculating the disk space required to host the VM.

The RAID levels are prioritized in the following order 0, 1, and 5, with 0 being the lowest.

Disks are sized on the target server based on the VM with the highest RAID level.

**SMP** – The SMP feature of ESX Server allows you to configure a VM with two processors. Check this option if you would like to configure your VM with SMP.

## Platform selection

Select up to four platforms you wish to review as candidates for your target servers. The tool performs a separate sizing for each platform and presents each for comparison.

---

**Figure 6:** Selecting the virtualization target host

Select the server platform you would like to consolidate to. You can compare up to four platforms at a time by checking each option that applies.

- ProLiant BL20p G2 blade server
  - ProLiant BL20p G3 blade server
  - ProLiant BL26p blade server
  - ProLiant BL40p blade server
  - ProLiant DL380 G3
  - ProLiant DL380 G4
  - ProLiant DL380 G3
  - ProLiant DL380 G4
  - ProLiant DL385
  - ProLiant DL660
  - ProLiant DL660 G2
  - ProLiant DL660 G3
  - ProLiant DL685
  - ProLiant DL740
  - ProLiant DL780 G2
  - ProLiant ML670 G2
  - ProLiant ML670 G3
-

## Appendix B – HP ProLiant Essentials Performance Management Pack

Performance Management Pack (PMP) is a software solution that detects and analyzes hardware bottlenecks on HP ProLiant servers. PMP provides the tools needed to receive proactive notification of impending bottleneck conditions and debug existing performance issues.

### Architecture

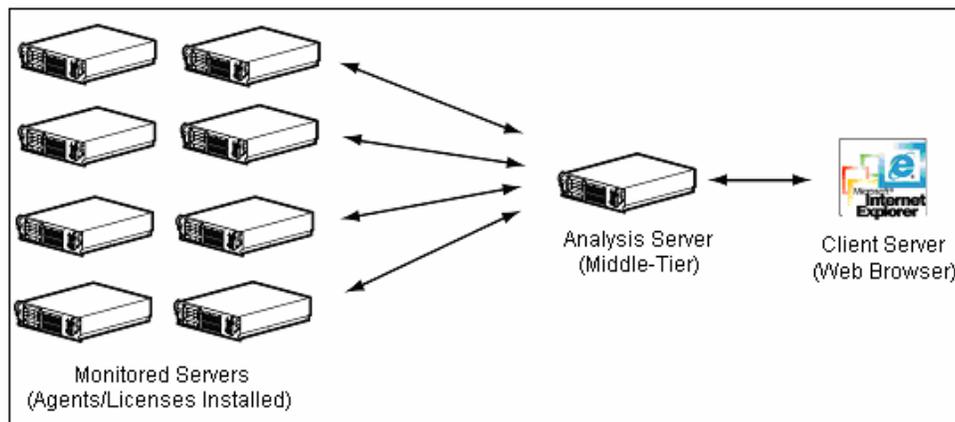
PMP is designed to run on a separate middle-tier server, known as the analysis server, which also acts as the HP SIM console server and performs the bulk of the processing for PMP, collecting SNMP performance data, for example.

The servers monitored by PMP are only required to run standard HP Insight Management (IM) Agents, installed as part of the HP ProLiant Support Pack.

In much the same way as HP SIM, PMP uses SNMP to collect information on the monitored servers from the IM Agents. There is minimal impact on the monitored servers.

IT staff access functionality and performance data from Microsoft Internet Explorer running on a client system; an applet on Internet Explorer displays this information.

**Figure 7:** PMP architecture, showing the analysis server monitoring multiple servers



Database software on the analysis server is required for storing and cataloging the functionality and performance data. The database requires Windows authentication to be set from HP SIM. For more information on setting authentication levels with HP SIM, refer to the [HP Systems Insight Manager Installation and User Guide](#).

## Hardware requirements

This section outlines hardware requirements for the analysis server, monitored server, and client.

### Analysis server

Table 3 lists the hardware required to implement an analysis server.

**Table 3:** Hardware requirements for the analysis server

Component	Requirement
Server	For a list of supported HP ProLiant servers, refer to the <a href="#">HP ProLiant Essentials Performance Management Pack Version 3.1.2 Support Matrix</a> .
System memory	<ul style="list-style-type: none"><li>• 192 MB RAM with Microsoft SQL Server or Microsoft Database Engine (MSDE) on the same server</li><li>• 256 MB RAM with Version Control Repository Manager on the same server with SQL Server or MSDE</li></ul> <p><b>Note:</b> For normal installations of PMP on HP SIM, HP recommends a minimum of 512 MB in addition to the 192 MB or 256 MB specified above.</p>
Disk space	<ul style="list-style-type: none"><li>• 45 MB on the Windows system drive</li><li>• 110 MB for HP SIM software</li><li>• 25 MB for PMP</li><li>• 2 MB for the database, the size of which increases in proportion to the amount of information logged for offline analysis</li></ul>

### Monitored server

For a complete list of supported monitored servers and their requirements, refer to the [HP ProLiant Essentials Performance Management Pack Version 3.1.2 Support Matrix](#).

### Client

PMP is best viewed with 256 colors and a monitor resolution of 1024 x 768. The browser window should be maximized.

## Software requirements

This section outlines operating system requirements and other software requirements for the analysis server, monitored server, and client.

### Operating system

For a complete list of the operating systems supported on the monitored servers, analysis server, and client system, refer to the [HP ProLiant Essentials Performance Management Pack Version 3.1.2 Support Matrix](#).

## Additional requirements for the analysis server

Table 4 lists additional software requirements for the analysis server.

**Table 4:** Software requirements for the analysis server

Component	Requirement
Operating system	<ul style="list-style-type: none"><li>• Microsoft Windows 2000, SP4</li><li>• Microsoft Windows 2000 Advance Server, SP4</li><li>• Microsoft Windows Server 2003, Standard Edition</li><li>• Microsoft Windows Server 2003, Enterprise Edition</li><li>• Microsoft Windows XP Professional, SP1 or later</li></ul>
Server software	<ul style="list-style-type: none"><li>• TCP/IP installed</li><li>• SNMP services installed and active</li><li>• Internet Explorer 6.0, or later</li></ul>
Database	<ul style="list-style-type: none"><li>• SQL Server 2000 Standard, SP3 or later</li><li>• SQL Server 2000 Enterprise, SP3 or later</li><li>• MSDE 2000, SP3 or later (local only)</li></ul>
HP SIM	<ul style="list-style-type: none"><li>• HP SIM 4.1 or 4.2</li></ul>

### Note:

MSDE 2000 is included on the HP Management CD and can be automatically installed if a database is not already available and running. A database engine must be installed for HP SIM to be successfully installed.

### Note:

For a remote database to work with a second database instance, Microsoft Data Access Components (MDAC) 2.7 SP1 or later is required.

### Note:

If SNMP services are installed after installing a Windows service pack, reinstall the service pack.

## Additional requirements for the monitored server

Servers monitored by PMP must have the standard IM Agents installed and running. For more information about IM Agents, refer to the [HP Systems Insight Manager Installation and User Guide](#).

## Additional requirements for the client

PMP requires Internet Explorer 6.0, or later, be installed on all client systems.

---

### Note:

To determine the current version of Internet Explorer, open the browser and select **Help→About Internet Explorer** from the menu bar. An information box displays the version currently installed.

---

## Browser security

Appropriate browser security and privacy options must be set on the local intranet web content zone for PMP to function as intended. For example, PMP cannot properly display graphs unless Internet Explorer ActiveX control settings are properly configured.

The required Internet Explorer security settings and privacy options on the client are:

- Enable **Download signed ActiveX controls**
- Prompt to **Download unsigned ActiveX controls**
- Enable **Run ActiveX controls and plug-ins**
- Enable **Script ActiveX controls marked safe for scripting**
- Enable **Active scripting**
- Allow stored and per-session cookies

Follow these steps to modify security and privacy settings for the local intranet web content zone:

1. Open the Internet Explorer browser.
2. Select **Tools→Internet Options** from the menu bar.
3. Select the **Security** or **Privacy** tab as necessary.
4. Perform any necessary modifications by selecting the appropriate options.

After all the requirements have been met, you are ready to install PMP and start monitoring. For information on these next steps, refer to the [HP ProLiant Essentials Performance Management Pack User Guide](#).

## Appendix C – Using AOG CapacityPlanner

AOG CapacityPlanner can take a full inventory of your current installed-base. This tool takes a three- or four-week snap-shot of the data center, providing a baseline for the consolidation process.

CapacityPlanner offers the following capabilities:

- **Inventory**
  - Physical description of every platform
  - Application
  - Location
- **Performance**
  - Main performance counters – using Perfmon counters (described in [Appendix D – Using Microsoft Windows Performance Monitor](#)), CapacityPlanner can collect performance data and identify trends affecting the main server subsystems (such as CPU, memory, disk I/O, and network); these trends can be used to help you design a suitable consolidation solution
  - Peak load – indicated by the maximum value recorded on a particular counter during the sampling period; you should monitor the frequency and duration of these peaks to determine if a bottleneck exists
  - Prime time load – the average workday load, specified over a particular timeframe (such as 9:00 am – 5:00 pm, Monday – Friday); by comparing your prime time loads with your peak loads, you can easily identify bottlenecks
  - Industry average

Just as with performance counters, peak load and prime time load sampling can be manipulated to observe performance trends in a specific time slot; for example, you can omit weekends or holidays, or periods with known productivity downtime. As a result, you can focus the sample on your most significant workloads and accurately identify when a server is really being used.

## Appendix D – Using Microsoft Windows Performance Monitor

Appendix D describes the use of Performance Monitor (Perfmon) to evaluate server performance, providing necessary input for the HP ProLiant server sizer for VMware ESX Server.

### Perfmon counters

At a minimum, HP suggests monitoring the following Perfmon counters to evaluate the performance of server subsystems.

`Logical Disk: Average Disk Queue Length`

This counter reflects the number of read and write disk requests awaiting service. This key metric characterizes hard disk performance, with increasing values indicating that performance is declining.

This counter should typically remain below two.

`Physical Disk: Average Disk Queue Length`

This counter is similar to the Logical Disk: Average Disk Queue Length counter just described. The primary difference is that, in this case, physical disks are monitored.

This counter should typically remain below two.

`Logical/Physical Disk: % Disk Time`

This counter represents how often the disk is busy and is an aggregate of data read from and written to the disk sub-system.

`Memory: Pages/sec`

This counter measures the number of pages read to or written from secondary storage (disk) when Windows cannot locate the requested information in primary storage (RAM). This is a key metric, helping determine if there is an excessive number of hard page faults<sup>4</sup>, which can result in disk thrashing.

Increasing values may indicate a RAM shortage.

`Network Interface:Bytes Total/sec`

Essentially, this counter reflects how often individual NICs are busy and should be tracked over time. Analyzing this counter allows you determine, for example, if you need to deploy an additional NIC.

---

#### Note:

Use of this counter assumes the server is running TCP/IP with SNMP service installed.

SNMP adds several TCP/IP-related counters to Perfmon.

---

`Objects:Processes`

This counter reflects the number of processes (essentially, programs) running at the moment the measurement is taken.

---

<sup>4</sup> Caused by Windows having to retrieve information from secondary storage

Processor: % Processor Time

This counter reflects how often the selected processor is busy. If this value is over 75 percent for an extended period (measured in days and weeks), the processor is over-utilized.

To input this data into the sizer, use the average percent utilization

### **Using disk I/O values with the sizer**

The disk counters described above do not provide metrics that can be entered directly into the HP ProLiant server sizer for VMware ESX Server.

If possible, you should compare the metrics to those obtained from a disk subsystem whose utilization is known. Otherwise, you can use the following methodology to infer utilization values:

- Theoretically add twice the workload to the disk system; if this new workload could be expected to consume 100% of disk resources, the original utilization is considered to be 50%.
- Theoretically add three times the workload to the disk system; if this new workload could be expected to consume 100% of disk resources, the original utilization is considered to be 25%.
- Continue extrapolating with different theoretical workloads until you obtain the inferred disk utilization value.

---

#### **Note:**

More accurate results are possible if the disk workload can be varied physically rather than theoretically.

---

### **Evaluation best practices**

- In order to evaluate performance correctly, do not allow the performance of one subsystem to characterize another. For example, if disk queue times are increasing due to an overactive swap file, it is probable that RAM is the cause rather than the disk subsystem.
- Ensure that no server resource is consistently over-burdened during the evaluation. If this should occur, correct the problem before continuing the evaluation.

## For more information

VMware documentation	<a href="http://www.vmware.com/support/pubs/">http://www.vmware.com/support/pubs/</a>
Performance tuning	<a href="http://www.vmware.com/pdf/esx_performance_tips_tricks.pdf">http://www.vmware.com/pdf/esx_performance_tips_tricks.pdf</a>
Sizing memory for a virtual infrastructure	<a href="http://www.vmware.com/support/esx25/doc/admin/esx25admin_res-mem-sizing-intro.html">http://www.vmware.com/support/esx25/doc/admin/esx25admin_res-mem-sizing-intro.html</a>
Information on the installation and administration of HP SIM	<a href="http://www.hp.com/servers/manage">http://www.hp.com/servers/manage</a>
Server, SAN, and option compatibility guides	<a href="http://www.vmware.com/hp">http://www.vmware.com/hp</a>
HP StorageWorks products	<a href="http://h18006.www1.hp.com/storage/index.html">http://h18006.www1.hp.com/storage/index.html</a>
VMware VirtualCenter	<a href="http://www.vmware.com/products/vmanage/vc_features.html">http://www.vmware.com/products/vmanage/vc_features.html</a>

To help us improve our documents, please provide feedback at [www.hp.com/solutions/feedback](http://www.hp.com/solutions/feedback).

© 2005, 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.

Microsoft, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation. AMD and AMD Opteron are trademarks of Advanced Micro Devices, Inc. Intel, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Java is a US trademark of Sun Microsystems, Inc. Linux is a U.S. registered trademark of Linus Torvalds. UNIX is a registered trademark of The Open Group.

4AA1-0357ENW, Revision 2, November 2007

