

SmartPeak WLM with HP BladeSystem and Microsoft Virtual Server 2005 R2 – dynamic workload management in a virtual environment



Executive summary.....	3
Audience	4
Introduction.....	4
Resource virtualization in the Adaptive Enterprise	4
Adaptive Enterprise	4
Virtual infrastructure.....	5
Dynamic Systems Initiative	5
HP BladeSystem.....	5
HP ProLiant Essentials	6
HP Systems Insight Manager	6
Requirements in a Windows environment	6
Virtual Machine Management Pack.....	8
Infrastructure	8
Requirements	8
HP OpenView	9
Support	9
Microsoft Virtual Server 2005 R2	10
What's new?	10
SmartPeak WLM.....	10
Key CPU control features.....	13
Smart Scheduling with Share Factors	13
CPU reservations	13
CPU application limits.....	13
CPU soft limits	13
Thread Throttling	13
Key memory control features	14
Virtual memory optimization	14
Physical memory control.....	14

Key network control features.....	15
Network bandwidth management	15
Network quota management	15
Cost-effective licensing.....	15
Testing	15
Objectives.....	16
Performance monitoring.....	16
Test environment	16
Hardware.....	16
Host server software	17
Test scenarios.....	18
Selecting the workload for the secondary application	18
SmartPeak WLM Application Footprint	18
Test results	19
Baseline tests	19
Without SmartPeak WLM.....	20
One VM – without SmartPeak WLM	20
Two VMs – without SmartPeak WLM.....	21
Three VMs – without SmartPeak WLM.....	22
With SmartPeak WLM	23
One VM – with SmartPeak WLM.....	23
Two VMs – with SmartPeak WLM.....	24
Three VMs – with SmartPeak WLM	25
Analysis.....	26
For more information.....	29

Executive summary

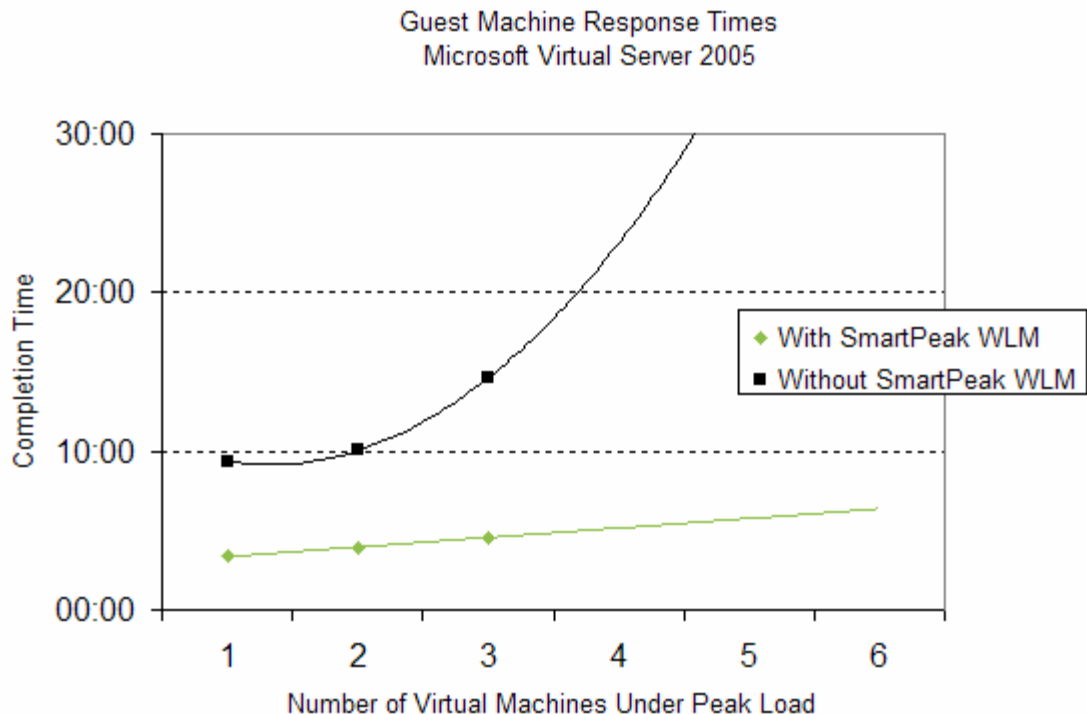
Infrastructure virtualization is changing the IT landscape and has become a cornerstone of strategies such as the HP Adaptive Enterprise and the Microsoft® Dynamic Systems Initiative. Tools such as SmartPeak WLM are key to the successful implementation of a virtual infrastructure.

SmartPeak WLM delivers a solution capable of managing diverse and demanding server systems (including web, mail, and application servers). SmartPeak WLM's workload management and optimization technologies dynamically control processor, virtual and physical memory, and network bandwidth usage in order to improve Quality of Service and increase consolidation ratios.

This white paper outlines the capabilities of SmartPeak WLM and describes the benefits of using this solution in a virtualized environment based on Microsoft Virtual Server 2005 R2 technologies. In addition, several complementary HP solutions (such as HP BladeSystem, HP ProLiant Essentials, and HP OpenView) are outlined.

In partnership with SmartPeak, HP performed a series of tests designed to demonstrate the ability of the default configuration of SmartPeak WLM to enhance the performance of an application running on Microsoft Windows Server 2003-based Virtual Machines (VMs). The tested environment included an HP BladeSystem server, HP StorageWorks storage, and Virtual Server virtualization software.

Test results clearly showed that SmartPeak WLM can add significant value in a Virtual Server guest environment. With SmartPeak WLM, three VMs were able to simultaneously complete an application task in less time than a single VM that was not using SmartPeak WLM. Overall, HP concluded that, with SmartPeak WLM, applications can run three times faster when VMs are under peak load, as demonstrated by the following test results:



Audience

The information contained in this white paper is intended for solutions architects, engineers, and project managers involved in the planning, implementation and ongoing management of virtualized infrastructures based on Virtual Server. The reader should be familiar with virtualized infrastructures, have a basic knowledge of Virtual Server, HP BladeSystem, HP StorageWorks Enterprise Virtual Array storage subsystems, HP ProLiant Essentials, and Windows Server 2003.

This white paper is not designed to replace documentation supplied with individual solution components but, rather, is intended to serve as an aid to those wishing to understand the benefits of utilizing SmartPeak WLM within a virtualized environment.

Introduction

SmartPeak WLM provides a workload management solution designed to optimize the performance of Windows servers. With SmartPeak WLM, computing resources are reliably and consistently delivered to applications or users based on business policies or other priorities.

This solution can dynamically control processor, virtual and physical memory, and network bandwidth usage in diverse, demanding application server systems. Furthermore, the ability to automatically analyze and optimize can lead to significant reductions in system memory overhead, disk paging, and application start-up times. These performance gains provide considerable increases in server capacity and Quality of Service in consolidated environments, further reducing hardware, software, and operational costs.

This white paper describes testing designed to demonstrate that SmartPeak WLM can enhance the performance of an HP BladeSystem server in a virtualized environment. Performed by HP and SmartPeak engineers using industry-standard HP ProLiant hardware and HP ProLiant Essentials software, the testing focused on processor utilization and overall throughput.

The ability to manage a virtualized environment using tools such as SmartPeak WLM is key to the successful implementation of an Adaptive Enterprise, which is described in the next section.

Resource virtualization in the Adaptive Enterprise

A critical component of the Adaptive Enterprise envisioned by HP – as well as Microsoft's Dynamic Systems Initiative (DSI) – 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 to create an operating environment where challenges can be changed quickly 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.

The 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 of 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 can virtualize an entire operating environment – operating system, applications, configuration settings, and other desirable elements – that environment, which is known as a Virtual Machine (VM), can be rolled out anywhere in the organization to maintain business continuity. To maximize availability, emerging technologies¹ can allow VMs to automatically migrate from a potentially failing host to another platform with little or no user intervention.

For more information on virtualizing IT in an Adaptive Enterprise, see <http://h71028.www7.hp.com/enterprise/cache/8886-0-0-225-121.aspx>.

Dynamic Systems Initiative

The Dynamic Systems Initiative (DSI) is a commitment from Microsoft and its partners to deliver self-managing, dynamic systems. With DSI, it becomes possible to design more manageable systems and automate operations, reducing costs and freeing up time for IT staff to focus on what is most important to the organization.

Virtual Server 2005 R2 plays an important role in this vision, helping customers simplify their IT operations, improve server hardware utilization, and more flexibly provision their data center resources.

For more information on DSI, see <http://www.microsoft.com/windowsserversystem/dsi/default.aspx>.

The following sections outline some of the systems and tools that can be used to implement and manage a virtualized infrastructure.

HP BladeSystem

Providing the ability to scale up or scale out, HP BladeSystem can provide an ideal platform on which to build a virtualized infrastructure.

Through integrated network and Storage Array Network (SAN) switches, HP BladeSystem offers common network and storage connectivity for an entire enclosure of HP BladeSystem servers. As a result, this 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 becomes a simple plug-and-play operation, dramatically reducing management costs and time-to-deployment.

¹ Such as Microsoft Virtual Server 2005 R2 host clustering

HP ProLiant Essentials

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](#) (HP SIM) and HP ProLiant Essentials [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.

Additionally, Microsoft Operations Manager (MOM) 2005, a key component of DSI, can give IT staff the knowledge to avoid the avoidable, reducing the complexity associated with managing the environment and lowering the cost of operations.

HP Systems Insight Manager

HP Systems Insight Manager (SIM) is a management application that helps IT manage all HP servers and system hardware within an IT environment. Regardless of the size or complexity of the environment, HP SIM can help make the organization more efficient and proactive in identifying, diagnosing, and fixing potential issues. 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 divided into the various components of an HP SIM environment: Central Management Server (CMS), the managed system, and the network client.

Table 1. Requirements for installing HP SIM in a Windows environment

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

Note:

HP SIM can also manage Linux servers.

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

Virtual Machine Management Pack

Integrated with HP SIM, Virtual Machine Management Pack (VMM) delivers centralized management and control capabilities for VMs and the resources of their host HP ProLiant servers. 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 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 physical-to-virtual (P2V), virtual-to-physical (V2P), or virtual-to-virtual (V2V) migration of managed VMs.

Note:

SMP has additional licensing requirements.

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 Virtual 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**

For additional information on CMS, HP SIM, or any other HP ProLiant Essentials virtualization management software, see:

<http://h18004.www1.hp.com/products/servers/proliantessentials/valuepack/vms/index.html>.

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 the 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)
- Smart Plug-in (SPI) for managing the Microsoft Virtual Server infrastructure
- IT service levels and quality maximized

Note:

For more information on the SPI for Virtual Server, see:

http://www.managementsoftware.hp.com/products/spi/spi_msvs/index.html.

The multi-platform HP OpenView solution allows you to manage a heterogeneous environment² 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 provide a quick return on investment.

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

For more information on HP OpenView, see

<http://www.managementsoftware.hp.com/index.html>.

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

Microsoft Virtual Server 2005 R2

What's new?

Virtual Server 2005 R2 (Virtual Server) includes the following new capabilities:

- **Host clustering** – Support for host-to-host connectivity allows you to cluster all VMs running on a host.
- **iSCSI support** – iSCSI clustering enables guest-to-guest connectivity across physical machines.
- **x64 support** – Virtual Server runs natively within a 64-bit Windows host operating system, providing increased performance and memory headroom.
- **Enhanced PXE booting** – PXE boot support has been added to the VM network adapter. Now, with the appropriate network infrastructure in place, you can perform network installations of guest operating systems using automated deployment tools – such as HP ProLiant Essentials Rapid Deployment Pack ([RDP](#)) – in the same way as physical servers.
- **Other** – Virtual Server now includes improved hyperthreading, support for F6 Disk (SCSI driver), and virtual disk pre-compact functionality.

For more information on Virtual Server, see

<http://www.microsoft.com/windowsserversystem/virtualserver/evaluation/vsoverview.aspx>.

SmartPeak WLM

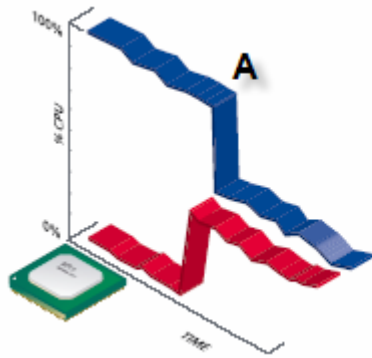
Aside from the obvious cost savings, significant benefits such as enhanced manageability, security, availability and disaster recovery can be brought to the business infrastructure through application server consolidation. However, consolidation is not without risk: conflicts between key applications such as database, messaging, web and application servers can quickly reduce the Quality of Service.

Even when each application has a dedicated VM, there can be problems. The requirement to maintain consistent, reliable performance from complex, demanding applications can be challenging: mismanaged system resources can easily cause bottlenecks; misbehaving or runaway processes can quickly saturate CPU and memory usage, causing other critical processes to fail.

However, SmartPeak WLM offers a solution, utilizing user- and application-based policies to allocate and prioritize critical system resources (such as processor, memory, and network bandwidth) according to business requirements. As a result, SmartPeak WLM can deliver extensive performance and Quality of Service enhancements in virtualized environments that are based on Microsoft technologies.

At the heart of the SmartPeak solution is patented Intelligent Process Management (IPM) technology, illustrated in Figure 1, which dynamically prioritizes resources based on default configurations or administrator-defined policies. Users continue to work as resources are dynamically re-allocated.

Figure 1. Prior to Point A, the “blue” application consumes all the CPU resource; after Point A, IPM engages to dynamically allocate the CPU fairly between “blue” and “red” applications



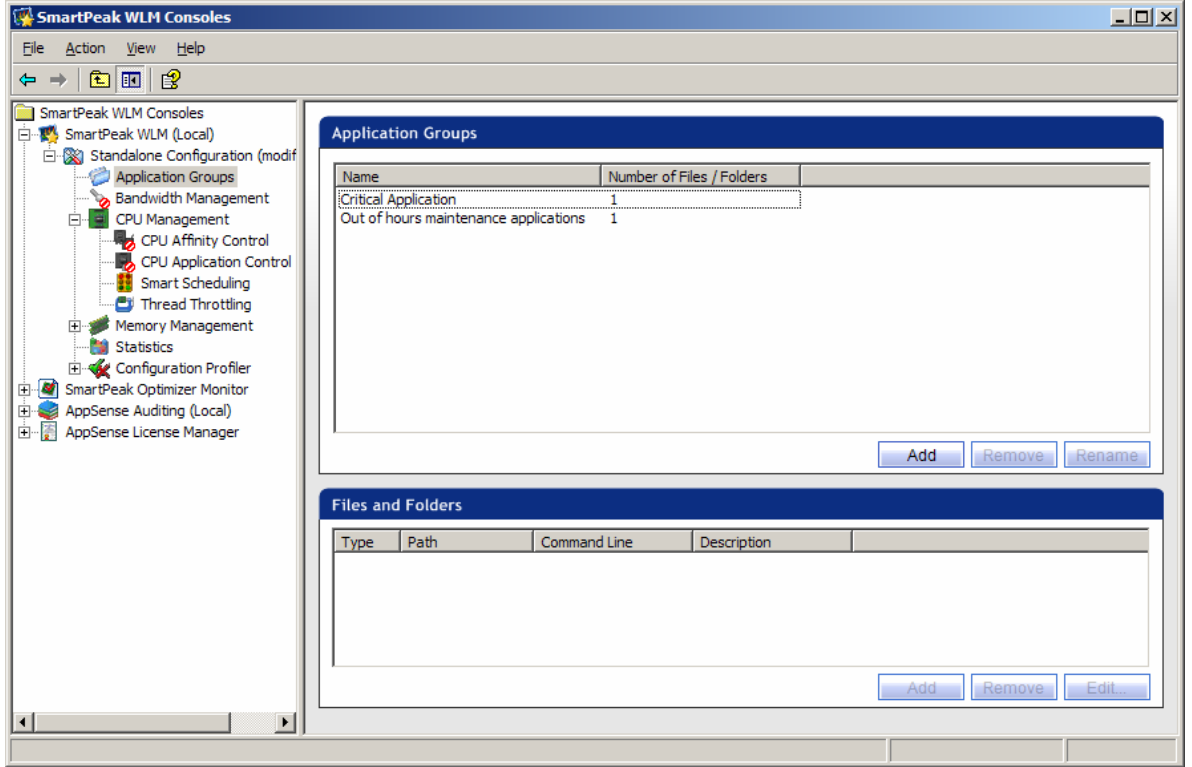
In addition to resource sharing, SmartPeak WLM features include:

- Application components are automatically analyzed and optimized, resulting in significant reductions in system memory overheads, disk paging, and application start-up times.
- The integrated auditing framework records important performance events in industry-standard formats such as system event log, e-mail, and SNMP.
- SmartPeak WLM can be centrally managed and distributed to all your physical and VM environments.
- Since SmartPeak WLM agents are based on Microsoft Windows Installer (MSI), agents and configurations can be deployed through technologies such as HP OpenView Radia.

SmartPeak WLM is a generic server solution and is not tied to specific databases, virtual environments, or other server systems, eliminating the need to deploy multiple software solutions to target specific environments. Figure 2 shows the SmartPeak WLM console, with the applications that perform scheduled maintenance tasks contained in a separate application group from critical applications.

More information on the key features of SmartPeak WLM follows.

Figure 2. Administration is simpler with SmartPeak WLM Application Groups



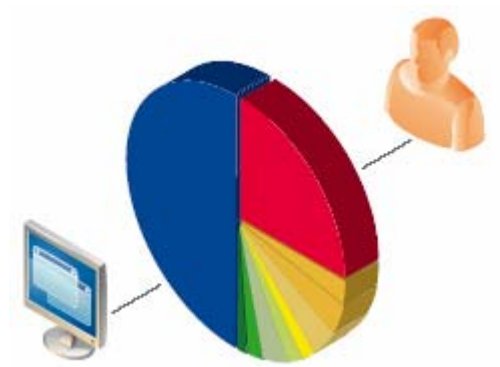
Key CPU control features

SmartPeak WLM provides a number of mechanisms for controlling the CPU resource.

Smart Scheduling with Share Factors

Grant critical applications or users within a VM preferential access to CPU resources through weighted share factors, as shown in Figure 3. For example, if desired, an application with a greater weighting can consume a larger share of the CPU.

Figure 3. Share factors allow applications that have been assigned a greater weighting to be granted higher priority to CPU resources



CPU reservations

Ensure that critical business processes continue to run uninterrupted by applying reservations to applications, guaranteeing that a percentage of the total CPU resource granted to the guest is always available, if needed. The CPU can be used by other applications until the critical process requires additional CPU time – up to the value of the reservation – which is immediately granted.

CPU application limits

Set hard CPU limits for applications running processes that are known to be problematic, such as legacy 16-bit applications. Hard CPU limits mean that an application can never use more than a specified amount of CPU time.

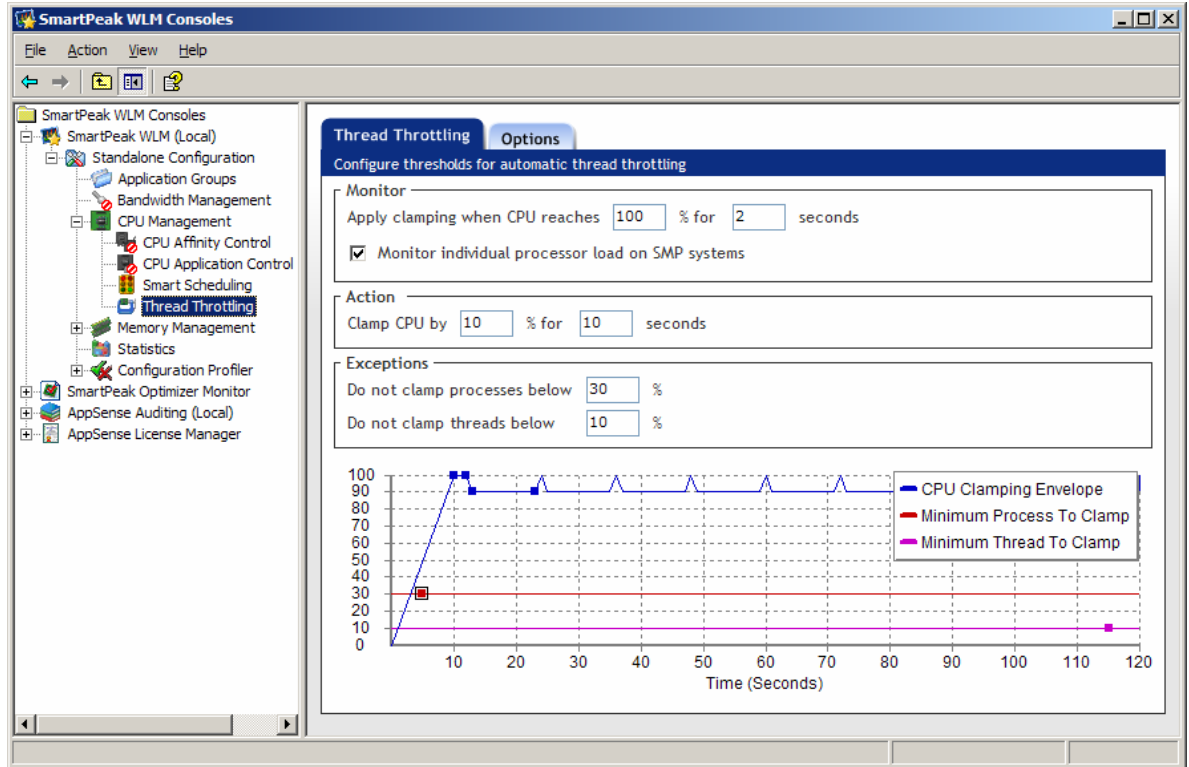
CPU soft limits

Assign a low priority to an application that is consuming more than the specified amount of CPU time. This means that the application continues to receive access to system resources until these resources are requested by another application.

Thread Throttling

Set system-wide policies that automatically trigger when the system is heavily loaded so that runaway threads within a problematic process can be gradually throttled. Figure 4 shows an example of Thread Throttling.

Figure 4. Dynamic management of runaway threads helps prevent server lock-ups



Note:

The tests described in this white paper utilized Smart Scheduling with Share Factors and Thread Throttling feature sets. For more information, see the [Testing](#) section.

Key memory control features

SmartPeak WLM provides a number of mechanisms for controlling the memory resource.

Virtual memory optimization

Significantly reduce virtual memory overhead and system paging by automatically analyzing and optimizing the Dynamic Link Libraries (DLLs) loaded by applications. Optimized DLLs are stored in a separate cache and loaded dynamically, leaving the original applications intact.

Physical memory control

Automatically trim working sets.

Key network control features

SmartPeak WLM provides a number of mechanisms for controlling the network resource.

Network bandwidth management

Continuously apply network throughput limits for applications based on network protocol, port number, or direction.

Network quota management

If pre-defined bandwidth quotas are exceeded, either block further network communication for the specific protocol, port, and direction, or limit throughput.

Cost-effective licensing

SmartPeak WLM is licensed based on the number of physical processors available on each physical server. This license allows you to install an instance of the product agent on each VM hosted on a server, providing economical workload management and delivering a cost-effective solution for any server or application consolidation.

Testing

These tests were designed to demonstrate the functional concepts of SmartPeak WLM regardless of the server applications in use. The workload was generated by a generic CPU resource loading tool in conjunction with a market-leading Anti Virus (AV) application.

The AV application was selected for its ability to provide a realistic workload consisting of disk scanning, hashing algorithms, database lookups and other tasks. The time taken to complete a manual AV scan while the server was under load from the generic tool became a key performance metric.

A goal of the testing was to demonstrate that SmartPeak WLM could enhance the performance of the AV application.

IMPORTANT:

In a production environment, it is more likely that SmartPeak WLM would be used to limit the access of management applications (such as AV or backup) to server resources that are required by core applications.

Thus, to avoid confusion, this white paper refers to the AV application as the **primary application**, while the arbitrary workload created by the generic tool is referred to as the **secondary application**.

While it is easy to configure SmartPeak WLM, these tests were intended to demonstrate the value of SmartPeak WLM's default configuration. In practice, however, the configuration of SmartPeak WLM would typically be optimized to deliver the most appropriate solution for a particular environment.

By focusing on the CPU scheduling features of SmartPeak WLM, this white paper focuses on just one aspect of the product in order to show its particular benefits in a reproducible scenario.

Objectives

Tests were undertaken to achieve the following objectives:

- Demonstrate that, with SmartPeak WLM, the primary application running in a guest VM can complete a known task in less time.
- With multiple VMs, demonstrate that SmartPeak WLM can increase consolidation ratios by improving performance and scalability.

In addition, it was important to ensure that the performance of other guest machines was not adversely affected. This was achieved by monitoring the Processor Queue Length and Context Switch counters as well as the CPU usage of each process.

Each test was repeated several times to obtain repeatable results. However, the nature of the primary application is such that its initial run is, in effect, a caching run; performance improves in subsequent runs. As such, an initial run – or the first runs after a VM restart – was regarded as invalid for test purposes.

Elapsed times and performance metrics quoted in this white paper reflect average values recorded over valid test iterations. As such, these metrics may differ from those indicated in screenshots, which reflect a single test iteration and are included for demonstration purposes.

Performance monitoring

Due to possible inconsistencies in performance metrics and elapsed times recorded within a virtual machine, additional measurements were recorded from outside the VM to ensure accuracy and consistency of all data provided. To this end, Windows Performance Monitor was used to record performance metrics at both the physical and virtual layers. In addition, elapsed times were verified, while the system resources consumed (and recorded) by virtual machines relate to the resources allocated to that virtual machine by the host operating system. Supporting data not included in this white paper can be found on the SmartPeak website.

This monitoring was carried out to demonstrate that SmartPeak WLM was able to enhance VM performance without adding performance overhead at the host level.

Within the guest machines, the following performance counters were monitored:

- \\%computername%\(_Total)\% Processor Time
- \\%computername%\Process\ (Primary Application)% Processor Time
- \\%computername%\Process\ (Secondary Application Process 0)% Processor Time
- \\%computername%\Process\ (Secondary Application Process 1)% Processor Time
- \\%computername%\System\Processor Queue Length
- \\%computername%\System\Context Switches/sec

Test environment

All testing was completed at the HP Solution Alliances Engineering Labs in Houston, Texas.

Hardware

The server hardware used for these tests consisted of the following physical environment:

- HP ProLiant BL20p G3 server blade
- Two Intel® Xeon® 3.6 GHz/1 MB processors with Intel Hyper-Threading (HT) enabled
- 4 GB RAM (with 512 MB configured for each VM)
- Fibre-attached HP StorageWorks 5000 Enterprise Virtual Array (EVA5000)

Host server software

- 32-bit Windows Server 2003 Enterprise Edition SP1
- Virtual Server 2005 R2

VM software (each VM)

- 32-bit Windows Server 2003 Enterprise Edition (SP1)
- Primary application (market-leading AV application)
- Secondary application (generic CPU loading tool)
- SmartPeak WLM
 - Smart Scheduling with Share Factors
 - Thread Throttling

Test scenarios

The tests were intended to emulate common business application scenarios. Each VM was subjected to a high, arbitrary workload by the secondary application running as two processes that consumed approximately 45% and 35% respectively of the available CPU resources within the guest. While the secondary application was running, the primary application was started.

Three additional VMs were also running on the server, processing moderate, arbitrary workloads that consumed a total of approximately 30% of the host CPU resource.

As stated above, the default configuration of SmartPeak WLM was utilized. With this configuration, all requests for system resources were serviced equally with no application afforded priority over any other.

Selecting the workload for the secondary application

As stated above, the secondary application was configured to require the VM to allocate approximately 80% of the VM's CPU resources before the primary application was initialized.

Without SmartPeak WLM, and taking into account the extra smaller workloads on other VMs, 80% was the highest possible secondary workload that would allow the primary application task to complete within a manageable timeframe. Higher secondary workloads would have caused the primary application to complete in hours or even days.³

With SmartPeak WLM, however, and a workload of 100%, the primary application was still able to complete in less than four minutes.

Although less dramatic, consistent improvements continued to be achieved with workloads less than 80%.

SmartPeak WLM Application Footprint

On a per-VM basis, SmartPeak WLM – with all functionality activated – requires approximately 2% CPU time, 33 MB of disk space, and 12 Megabytes of RAM. However, since the purpose of SmartPeak WLM is to manage consumption of system resources, this footprint becomes less apparent as the load placed on the system increases.

³ Due to excessive system overheads on both host and guest server layers

Test results

This section outlines the results of performance testing performed in the following areas:

- Baseline tests
- Testing without SmartPeak WLM, with one, two, and three VMs
- Testing with SmartPeak WLM, with one, two, and three VMs

Baseline tests

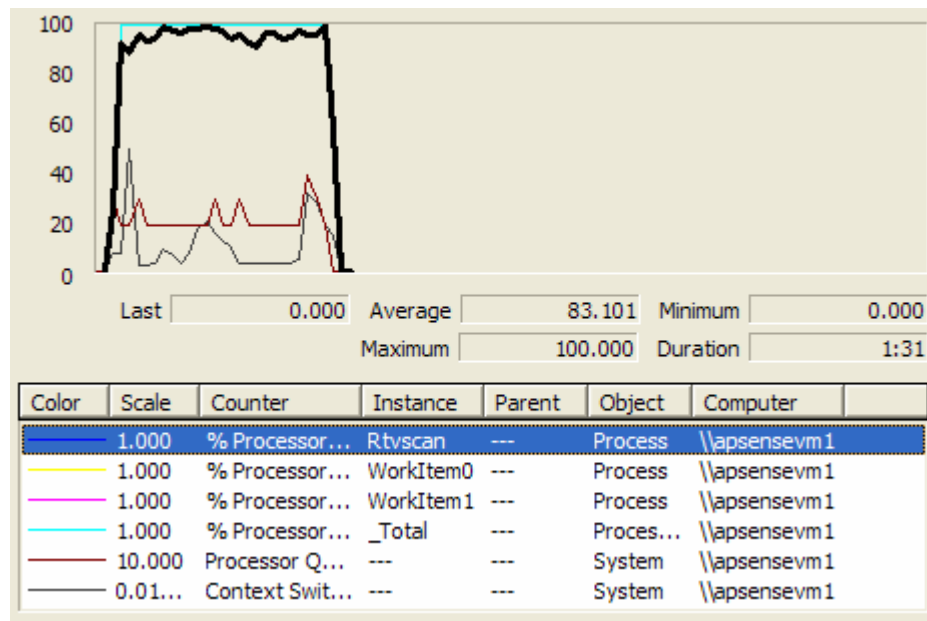
To establish an effective baseline, HP ran the primary application without SmartPeak WLM or any other application competing for system resources. In this scenario, the task was completed, on average, in one minute and 18 seconds. Table 2 indicates that the server was not required to perform at peak levels to accomplish this task.

Table 2. CPU consumption during the baseline test

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
01:18	92.35	96.07	2

Figure 5 shows a range of performance metrics.

Figure 5. The primary application executed alone as a baseline.



Without SmartPeak WLM

Without SmartPeak WLM activated, HP ran the primary and secondary applications on one, two, and three VMs. Meanwhile, three additional VMs processed a moderate, arbitrary workload, loading the CPU by approximately 10% within each guest.

These tests generated the following results.

One VM – without SmartPeak WLM


For this test, the primary and secondary applications were launched within a single VM without the use of SmartPeak WLM. Table 3 summarizes CPU consumption.

Table 3. CPU consumption with one VM; SmartPeak WLM is not activated

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
09:22	18.02	100	7

As shown in Figure 6, the primary application (highlighted) was starved of CPU time due to competition with the secondary processes. As a result, the time taken to complete the assigned task was over eight minutes longer than in the baseline test.

Figure 6. Resource conflicts occurring as the applications compete for resources



Two VMs – without SmartPeak WLM

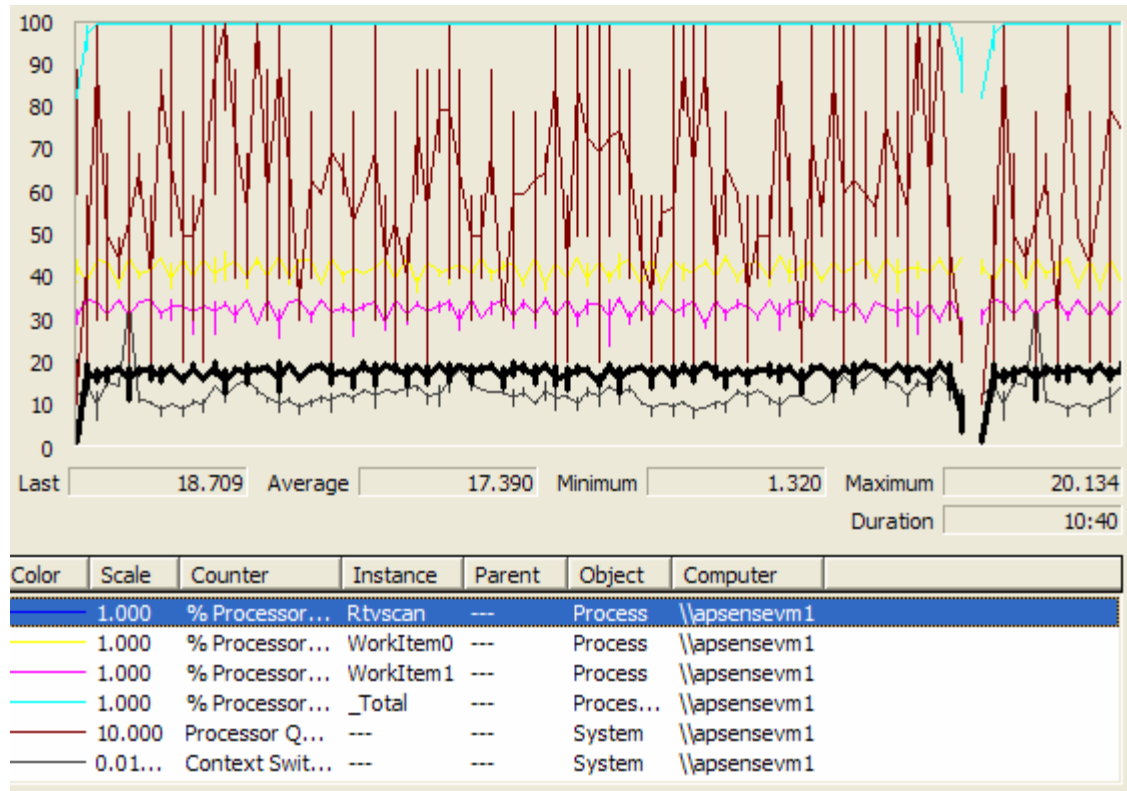
For this test, the primary and secondary applications were launched simultaneously within two different VMs without the use of SmartPeak WLM. Table 4 summarizes CPU consumption.

Table 4. CPU consumption with two VMs; SmartPeak WLM is not activated

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
10:02	17.72	99.27	6

Figure 7 shows the resource conflict.

Figure 7. Increased levels of resource conflict, resulting in additional delays completing the assigned task



Three VMs – without SmartPeak WLM


For this test, the primary and secondary applications were launched simultaneously within three different VMs without the use of SmartPeak WLM. Table 5 summarizes CPU consumption.

Table 5. CPU consumption with three VMs; SmartPeak WLM is not activated

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
14:37	17.22	98.65	6

Figure 8 shows the resource conflict.

Figure 8. The assigned task now takes almost 15 minutes to complete



With SmartPeak WLM

With SmartPeak WLM activated, HP ran the primary and secondary applications on one, two, and three VMs. Meanwhile, three additional VMs processed a moderate, arbitrary workload, loading the CPU by approximately 10% within each guest.

These tests generated the following results.

One VM – with SmartPeak WLM

For this test, the primary and secondary applications were launched within a single VM running SmartPeak WLM. Table 6 summarizes CPU consumption.


Table 6. CPU consumption with one VM; SmartPeak WLM is activated

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
03:22	43.5	91.1	4

The benefits of using SmartPeak WLM are clearly demonstrated: there is a 64% reduction in the time taken to complete the assigned task (compared to the equivalent test without SmartPeak WLM). Furthermore, the processor queue is shorter.

Figure 9 shows the absence of resource conflicts.

Figure 9. Resource conflicts are resolved, allowing the assigned task to be completed in just over three minutes



Two VMs – with SmartPeak WLM

For this test, the primary and secondary applications were launched simultaneously within two VMs running SmartPeak WLM. Table 7 summarizes CPU consumption.


Table 7. CPU consumption with two VMs; SmartPeak WLM is activated

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
03:58	43.28	90.72	4

The benefits of using SmartPeak WLM are clear: there is a 60% reduction in the time taken to complete the assigned task (compared to the equivalent test without SmartPeak WLM). Furthermore, the processor queue is shorter.

Figure 10 shows the continued absence of resource conflicts.

Figure 10. SmartPeak WLM continues to allow the assigned task to complete in significantly less time



Three VMs – with SmartPeak WLM

For this test, the primary and secondary applications were launched simultaneously within three VMs running SmartPeak WLM. Table 8 summarizes CPU consumption.

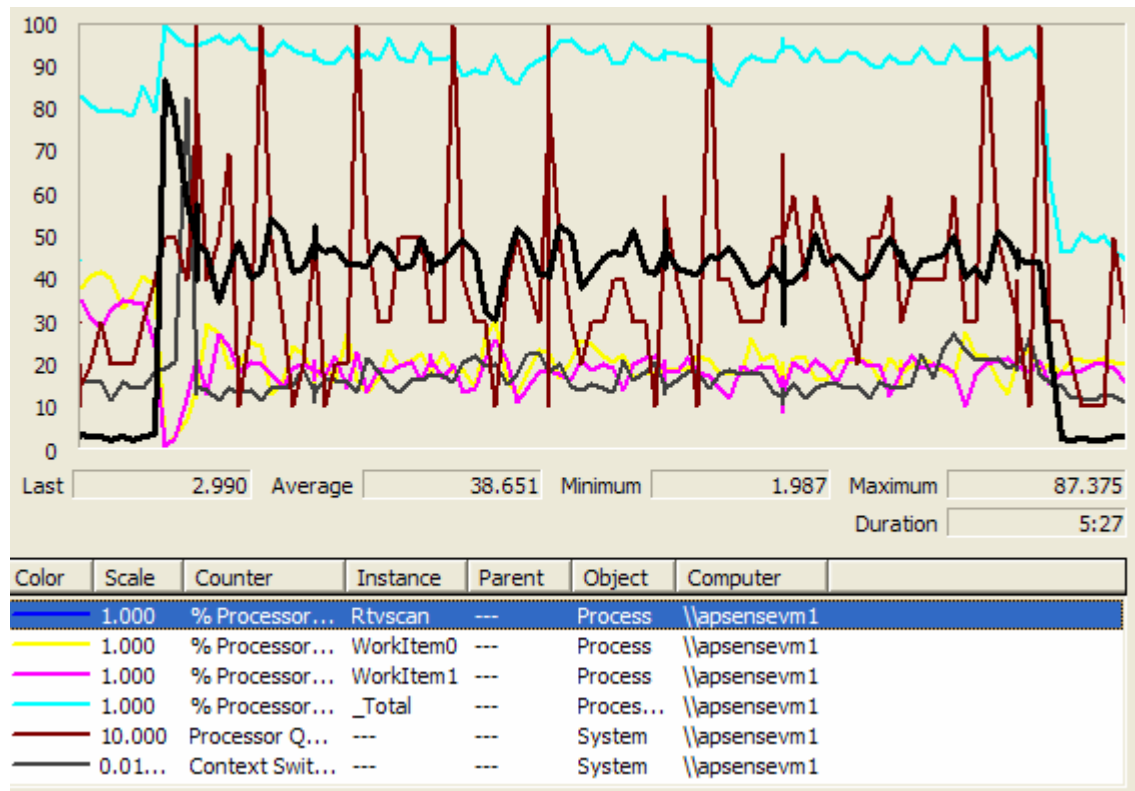
Table 8. CPU consumption with three VMs; SmartPeak WLM is activated

Elapsed time (mm:ss)	% Processor Time		Processor Queue Length
	Primary application	Total	
04:34	44.95	91.95	5

The benefits of using SmartPeak WLM are clear: there is a 69% reduction in the time taken to complete the assigned task (compared to the equivalent test without SmartPeak WLM). Furthermore, the processor queue is shorter.

Figure 11 shows the continued absence of resource conflicts.

Figure 11. With less resources provided to each of three virtual machine processing high workloads, SmartPeak WLM still allows the primary application to complete in approximately ten minutes.



Analysis

HP concluded that the stated goals of the tests were met. These were as follows:

- Demonstrate that, with SmartPeak WLM, the primary application running in a VM can complete a known task in less time.
- With multiple VMs, demonstrate that SmartPeak WLM can increase consolidation ratios by improving performance and scalability.

As shown in Figure 12, the assigned task was clearly completed in less time when SmartPeak WLM was utilized.

Figure 12. Completion times clearly improved with the use of SmartPeak WLM

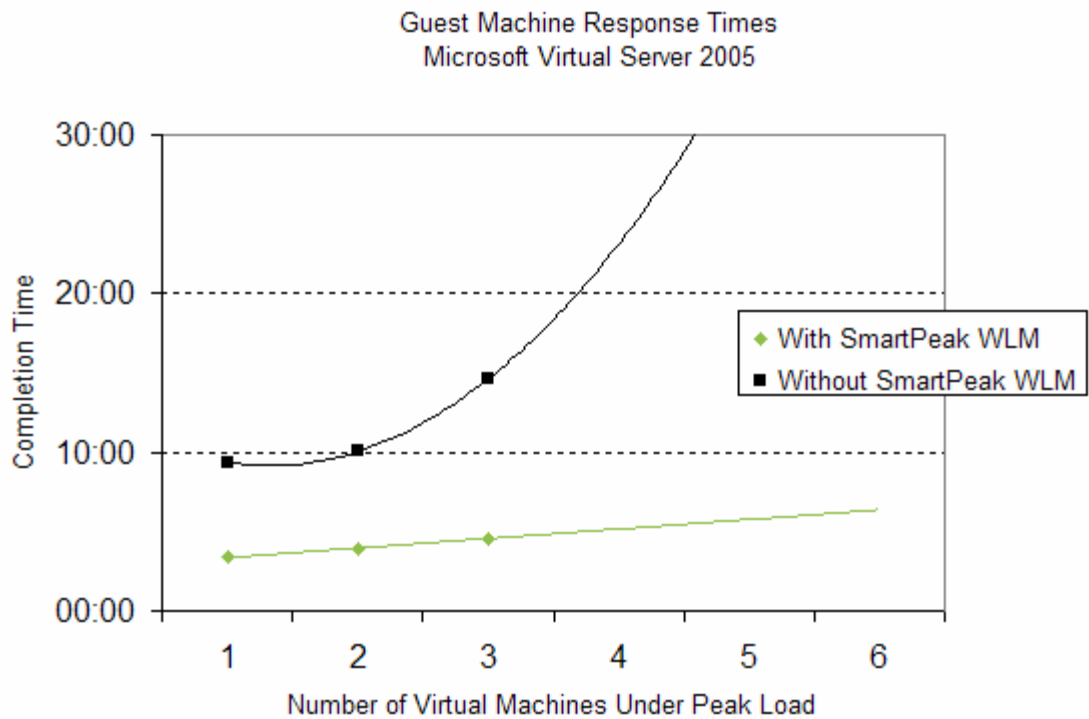


Table 9 summarizes the test results.

Table 9. Time taken for the primary application to complete

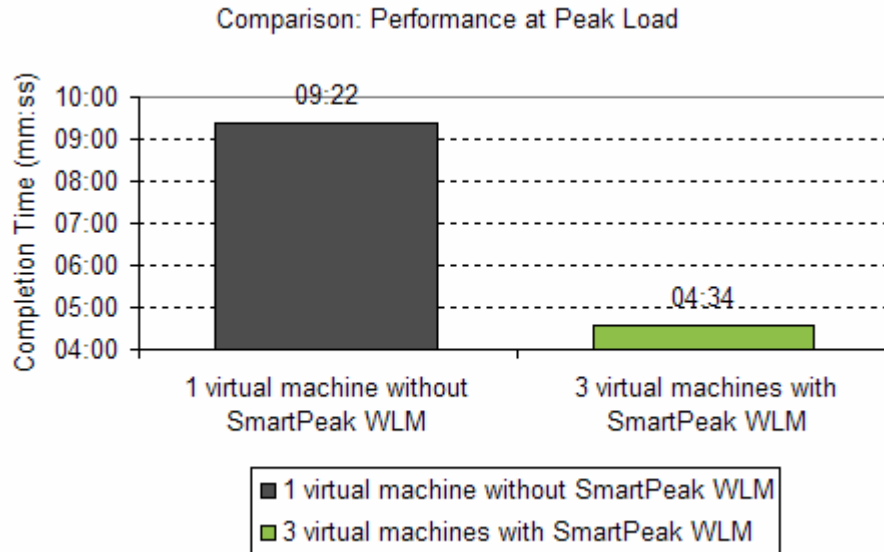
Elapsed Time	VMs under peak load		
	1	2	3
With SmartPeak WLM	03:22	03:58	04:34
Without SmartPeak WLM	09:22	10:02	14:37

Table 9 clearly demonstrates the ability of SmartPeak WLM to add significant value in a Virtual Server guest environment. When VMs were required to perform system-intensive tasks that resulted in resource conflicts, SmartPeak WLM allowed the primary application to complete in 30 – 36% of the time taken without SmartPeak WLM.

Furthermore, with SmartPeak WLM, three VMs simultaneously completed the primary application in less time than a single VM that was not using SmartPeak WLM, as shown in Figure 13. Thus,

SmartPeak WLM was able to more than triple the total workload supported. Overall, this translates into applications that run three times faster when VMs are under peak load.

Figure 13. SmartPeak WLM allowed the primary application to complete simultaneously on three VMs in less time than it took to complete on one VM not using SmartPeak WLM



The test results also demonstrate that use of SmartPeak WLM consistently reduced the Processor Queue Length when the system was under load, as shown in Table 10.

Table 10. Queue length

Processor Queue Length	VMs under peak load		
	1	2	3
With SmartPeak WLM	4	4	5
Without SmartPeak WLM	7	6	6

System metrics were also recorded at the physical machine level to demonstrate that performance improvements within virtual machines were not obtained at the expense of overall system resources. Tables 11 and 12 summarize the use of system resources at the physical machine level. Supporting data can be found on the SmartPeak website.

Table 11. Physical CPU Utilization

Physical CPU Utilization	VMs under peak load		
	1	2	3
With SmartPeak WLM	49.15	71.44	84.89
Without SmartPeak WLM	48.90	74.10	89.20

Table 12. Physical System Queue Length

Processor Queue Length	VMs under peak load		
	1	2	3
With SmartPeak WLM	1	2	3
Without SmartPeak WLM	1	2	4

For more information

HP BladeSystem	http://h71028.www7.hp.com/enterprise/cache/80316-0-0-0-121.aspx
HP ProLiant Essentials software	http://h18004.www1.hp.com/products/servers/management/
HP StorageWorks	http://h18006.www1.hp.com/storage/
SmartPeak WLM	http://www.smartpeak.com
Virtual Server	http://www.microsoft.com/virtualserver

© 2006 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 NT and Windows are U.S. registered trademarks of Microsoft Corporation. Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a U.S. registered trademark of Linus Torvalds. SmartPeak and SmartPeak WLM are trademarks of SmartPeak Ltd. Java is a US trademark of Sun Microsystems, Inc. Oracle is a registered US trademark of Oracle Corporation, Redwood City, California. UNIX is a registered trademark of The Open Group.

5/2006

