

HP Global Workload Manager Version 4.0 User's Guide

HP Part Number: T8672-90001
Published: May 2008



© Copyright 2004-2008 Hewlett-Packard Development Company, L.P.

Legal Notice

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

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

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group.

Table of Contents

About This Document.....	9
Intended Audience.....	9
Typographic Conventions.....	9
Related Information.....	9
Publishing History.....	10
HP Encourages Your Comments.....	10
Assumptions.....	10
Finding More gWLM Information.....	11
1 Overview.....	13
What is gWLM?.....	13
Benefits of Using gWLM.....	13
Comparing PRM, WLM, and gWLM Features.....	13
Concepts and Terms for Using gWLM.....	13
The gWLM Management Model.....	16
How gWLM Allocates CPU Resources.....	17
Available Interfaces.....	18
2 Configuring gWLM to Manage Workloads.....	19
Policy Types.....	19
Choosing a Policy Type.....	20
Choosing Between an OwnBorrow Policy and a Utilization Policy.....	20
Combining the Different Policy Types.....	21
Seeing How gWLM Will Perform Without Affecting the System.....	21
Getting Started with gWLM.....	21
Tabs and Menus.....	21
Using the Wizard.....	21
Seeing gWLM in Action.....	22
Common Uses for gWLM.....	24
Fixing the Amount of CPU Resources a Workload Gets.....	24
Resizing a Workload's Npar, Vpar, Virtual Machine, Pset, or FSS Group as Needed.....	25
Common Configuration Tasks.....	25
Setting up gWLM (Initial Setup Steps).....	25
Changing from Advisory Mode to Managed Mode.....	26
Creating a New Policy.....	26
Editing a Policy.....	27
Changing Which Policy is Associated with a Workload.....	27
Adding a New Compartment or GiCAP Group Member to an SRD.....	28
Stop Managing a Workload.....	28
Stop Managing an SRD.....	29
3 Monitoring Workloads and gWLM.....	31
Monitoring Workloads.....	31
High-Level View.....	31
Graphical Reports.....	31
Real-time Reports.....	31
Historical Reports.....	31
Viewing gWLM Reports in Monitor-Only Mode.....	31

Monitoring gWLM from the Command Line.....	32
Message Logs.....	32
Viewing HP Systems Insight Manager Events.....	34
Monitoring gWLM with GlancePlus.....	34
4 Security.....	35
General Security Topics.....	35
Securing gWLM Communications.....	35
5 Additional Configuration and Administration Tasks.....	37
Manually Adjusting CPU Resources.....	37
Manually Adjusting Memory Resources.....	38
Setting Aside Space for Historical Data.....	38
Backing up the VSE Management Software Database.....	38
Tips for Backup and Restore.....	39
Setting gWLM Properties.....	39
CMS Properties	39
Agent Properties.....	41
Communications Ports.....	42
Controlling gWLM's Startup Behavior.....	43
Automatic Restart of gWLM's Managed Nodes in SRDs (High Availability).....	44
How the Automatic Restart Works.....	44
Related Events.....	45
"Node Failed to Rejoin SRD on Start-up" Event.....	45
"SRD Communication Issue" and "SRD Reformed with Partial Set of Nodes" Events.....	45
Manually Clearing an SRD.....	45
Clearing an SRD of A.02.50.00.04 (or later) agents.....	46
Clearing an SRD of agents of any version.....	46
Nesting Partitions.....	47
Changing the gWLM Resource Allocation Interval.....	48
Changing the Interval in HP SIM.....	48
Changing the Interval on the Command Line.....	48
Using gWLM with Hyper-Threading.....	48
Using gWLM with Hosts on Multiple LANs.....	49
Creating "Golden Images"	50
A Compatibility with Agents.....	51

List of Figures

2-1	Top Controls of the Global Workload Manager Screen.....	21
5-1	Nested Partitions.....	47
5-2	Using gWLM with Hosts Separated by Firewalls.....	49

List of Tables

1	Where to Find Additional Information.....	11
1-1	Default Weights by Policy Type.....	17
2-1	Choosing a Policy Type.....	20
3-1	gWLM Log Files.....	32

About This Document

This document presents an overview of the techniques and tools available for using HP Global Workload Manager for HP Integrity servers and HP 9000 servers (gWLM). It exposes you to the essentials and allows you to get started quickly using gWLM.

Intended Audience

This document is intended to be used by Virtual Server Environment (VSE) system administrators, VSE application administrators, and other technical professionals involved with data center operations, administration, and planning. An understanding of HP-UX system administration concepts and procedures is assumed.

Typographic Conventions

This document uses the following typographic conventions.

<i>Book Title</i>	Title of a book or other document.
<u><i>Linked Title</i></u>	Title that is a hyperlink to a book or other document.
<u>http://www.hp.com</u>	A Web site address that is a hyperlink to the site.
Command	Command name or qualified command phrase.
user input	Commands and other text that you type.
computer output	Text displayed by the computer.
Enter	The name of a keyboard key. Note that Return and Enter both refer to the same key. A sequence such as Ctrl+A indicates that you must hold down the key labeled Ctrl while pressing the A key.
term	Defined use of an important word or phrase.
variable	The name of an environment variable, for example <code>PATH</code> or <code>errno</code> .
value	A value that you may replace in a command or function, or information in a display that represents several possible values.
<i>find(1)</i>	HP-UX manpage. In this example, “find” is the manpage name and “1” is the manpage section.

Related Information

The following related documents are available from <http://docs.hp.com/en/vse.html>:

- *VSE Management Software Version 4.0 Installation and Update Guide for HP-UX*
- *VSE Management Software Version 4.0 Getting Started Guide*
- *VSE Management Software Version 4.0 Release Notes*
- *HP Application Discovery Version 4.0 User's Guide*
- *HP Capacity Advisor Version 4.0 User's Guide*
- *HP Integrity Virtual Machines Manager Version 4.0 Getting Started Guide*
- *HP Integrity Virtual Machines Installation, Configuration, and Administration*
- *HP Integrity Virtual Machines Manager Release Notes*

Publishing History

Manufacturing Part Number	gWLM Release	Supported Operating Systems	Publication Date
T8672-90001	A.04.00.07	HP-UX 11i v1	May 2008
		HP-UX 11i v2	
		HP-UX 11i v3	
T2762-90057	A.03.00.00	HP-UX 11i v1	June 2007
		HP-UX 11i v2	
		HP-UX 11i v3	

HP Encourages Your Comments

Your comments and suggestions regarding product features can help HP develop future versions of the VSE Management Software.

Use the following email address to send feedback directly to the gWLM development team: gwlmfeedback@rsn.hp.com.

For comments on VSE Management Software other than gWLM, send feedback directly to the VSE Management Software development team: vse@hpuxweb.fc.hp.com.



NOTE: HP cannot provide product support through these email addresses. To obtain product support, contact your HP Support Representative, your HP Services Representative, or your authorized HP reseller. For more information about support services, visit <http://www.hp.com/go/support>. (If you need support in a country other than the U.S., visit <http://welcome.hp.com/country/us/en/othercountrieswel.html>. Click the link for your country or region, and then click **Support and Troubleshooting**.)

For a forum with other gWLM users, visit the IT Resource Center's forum for HP-UX Workload/Resource Management:

<http://forums.itrc.hp.com/cm/CategoryHome/1,,213,00.html>

For other ways to contact HP, visit http://welcome.hp.com/country/us/en/contact_us.html. (To contact HP in a country other than the U.S., visit <http://welcome.hp.com/country/us/en/othercountrieswel.html>. Click the link for your country or region, and then click **Contact HP**.)

Assumptions

It is assumed that you have already installed the following software:

- HP Systems Insight Manager (SIM)
- HP VSE Management Software CMS
- gWLM agent on your managed nodes

For information about setting up HP SIM, see the documentation available at <http://www.hp.com/go/hpsim>.

The following steps give an overview of the HP SIM installation process. When you install the VSE Management Software and gWLM, you must do the following:

1. Decide which system will be your central management server (CMS), then install the VSE Management Software CMS on that system.

This system must also have HP SIM installed and running.

2. Initialize the CMS by running the `vseinitconfig` command.
For more information, see *vseinitconfig(1M)*.
3. Decide which systems will be your managed nodes, then install the gWLM agent software on those systems. (The agent software is free, but it is functional only for a limited time. For unlimited use, purchase the agent license to use, or LTU.)
4. On each managed node, start the gWLM agent daemon `gwlmagent`.

You can perform the last two steps through HP SIM, as described in the *VSE Management Software Installation and Update Guide*.

Finding More gWLM Information

Table 1 indicates where you can find additional information.

Table 1 Where to Find Additional Information

To...	See...
View the structure (npars, vpars, ...) of your systems.	VSE Management Page in HP SIM (Tools → VSE Management)
Learn about configuring, backing up, and maintaining your CMS.	<i>vseinitconfig(1M)</i>
Use gWLM immediately, reading as little as possible.	gWLM Home Page in HP SIM (Tools → VSE Management , then click the Shared Resource Domain tab, then Tools → Global Workload Manager → Getting Started - gWLM Home) or Global Workload Manager topic in online help ¹ or <i>HP Global Workload Manager User's Guide</i> (this document) (http://docs.hp.com/en/vse.html ²)
Learn about gWLM concepts.	Global Workload Manager topic in online help ¹ or <i>HP Global Workload Manager User's Guide</i> (this document) (http://docs.hp.com/en/vse.html ²)
Learn gWLM terms.	"Concepts and Terms for Using gWLM" (page 13) or Glossary in online help ¹
Learn gWLM best practices.	"Getting the Most Out of gWLM" topic in online help ¹
Learn about other gWLM features.	<i>HP Global Workload Manager User's Guide</i> (this document) (http://docs.hp.com/en/vse.html ²)
Learn about the gWLM interface in HP SIM.	Online help ¹
Learn about the gWLM command-line interface.	<i>gwlm(1M)</i>
Learn about gWLM daemons and services	<i>gwlmcmd(1M)</i>
Learn about using secure communications with gWLM.	Securing gWLM Communications topic in online help ¹ or <i>gwlmsslconfig(1M)</i>
Learn how to update metrics when using custom policies.	<i>gwlmSEND(1M)</i>

Table 1 Where to Find Additional Information *(continued)*

To...	See...
Learn how to manually place processes in workloads based on psets or fss groups.	<i>gwlmp</i> lace(1M)
Learn about using gWLM with HP Serviceguard.	The "Technical Documentation web site for HP Virtual Server Environment (VSE)": http://docs.hp.com/en/vse.html
Learn more about npars, vpars, virtual machines, and psets.	<ul style="list-style-type: none"> • HP Virtual Server Environment web site: http://www.hp.com/go/vse • The "Technical Documentation web site for HP Virtual Server Environment (VSE)" web site: http://docs.hp.com/en/vse.html
Learn about other gWLM commands.	The SEE ALSO section of <i>gwl</i> m(5) for a list of all gWLM manpages

- 1 To access online help in HP SIM, click **Tools**→**VSE Management**, then click the **Shared Resource Domain** tab, and then click question mark (?) in the top right corner.
- 2 This documentation is in English. For documentation in other languages, see <http://docs.hp.com>.

1 Overview

This chapter provides an overview of gWLM, including benefits, key concepts and terms, and the gWLM management model.

What is gWLM?

gWLM allows you to centrally define resource-sharing policies that you can use across multiple HP servers. Using these policies can increase system utilization and facilitate controlled sharing of system resources. In addition, gWLM provides both real-time and historical monitoring of the resource allocation.

gWLM consists of a VSE Central Management Server, or CMS. You configure gWLM and monitor your workloads from the system where the CMS software is installed. Also, you use agent software on the systems where you have workloads you want gWLM to manage.

Benefits of Using gWLM

The benefits of using gWLM include:

- Better use of existing server capacity
Typically, servers are set up with a single workload and ample reserve capacity to handle the peak demand of that workload. gWLM allows you to combine multiple workloads with differing demand patterns on a single server and make use of the idle capacity—when it is not needed by your mission-critical workload.
- Confidence that mission-critical workloads get the required resources
Even with multiple workloads on a server, you can ensure your mission-critical workload gets the resources it needs. gWLM automatically adjusts resource allocation, making it easy to share resources when they are plentiful, and to dedicate resources to workloads during spikes in resource demand.
- Reduced system administration costs
With gWLM, you can combine more workloads on fewer servers, thereby reducing administration costs.

Comparing PRM, WLM, and gWLM Features

Process Resource Manager (PRM), Workload Manager (WLM), and gWLM all enable you to control system resources. Use only one of these products at a time on a system.

For a detailed comparison of WLM and gWLM features, see the white paper *Migrating from WLM to gWLM* available from <http://docs.hp.com/en/vse.html>. As for comparing PRM and gWLM features, gWLM provides most of the PRM functionality. However, gWLM does not provide the following PRM features:

- Memory controls
- Disk I/O controls
- Simultaneous management of processor sets and fss groups within a single HP-UX image
- Noncapped mode
- Integration with HP-UX Security Containment
- Integration with netgroup user lists
- Integration with HP System Management Homepage (SMH)

Concepts and Terms for Using gWLM

Here are some concepts and terms to know when using gWLM:

Workload	<p>The collection of processes executing within a single compartment. The compartment can be an nPartition (npar), a virtual partition (vpar), a virtual machine provided by HP Integrity Virtual Machines (hpvm), a processor set (pset), or a Fair Share Scheduler (fss) group. gWLM manages a workload by adjusting the system resource allocations for its compartment. (For background information on npars, vpars, virtual machines, psets, and fss groups, refer to the section “The gWLM Management Model” (page 16).)</p>
Compartment	<p>An npar, a vpar, a virtual machine, a pset, or an fss group with its resource allocation being managed by gWLM.</p> <p>Multiple compartments are grouped to form a shared resource domain, or SRD. The compartments all share the resources within the SRD. Each compartment holds a workload and can be in only one deployed SRD. gWLM manages each workload by adjusting the resource allocation for its compartment.</p>

Shared Resource Domain (SRD)	<p>A collection of compartments that can share system resources. The compartments can be npars, vpars, virtual machines, psets, or fss groups. For example, a server containing npars can be an SRD—as long as the requirements in “The gWLM Management Model” (page 16) are met. Also, a server or an npar divided into vpars can be an SRD for its vpar compartments. Similarly, a server or an npar divided into virtual machines can be an SRD for its virtual machine compartments.</p> <p>gWLM allows you to nest compartments. gWLM then manages resources for the various levels of compartments.</p>
Policy	<p>A collection of settings that instruct gWLM how to manage a workload’s resources. For example, a policy can indicate the amount of CPU resources a workload owns (and is allocated when needed), as well as how much of those resources the workload can lend to other workloads.</p> <p>A single policy can be associated with, or applied to, multiple workloads.</p> <p>For more information on policies, see “Policy Types” (page 19).</p>
Mode	<p>Two modes are available: advisory and managed. Advisory mode allows you to see what CPU resource requests gWLM would make for a workload—without actually affecting resource allocation.</p> <p>Advisory mode is not available for SRDs containing virtual machines, psets, or fss groups due to the nature of these compartments.</p> <p>Use this mode when creating and fine-tuning your policies. Once you are comfortable with your policies, use managed mode to have gWLM automatically adjust the resource allocations for your defined workloads.</p> <p>You can only set the mode on the SRD level. All workloads within an SRD operate in the same mode, either advisory or managed.</p>
Deploy	<p>Enable gWLM control of an SRD.</p> <p>Deploying an SRD in managed mode enables gWLM control of resource allocation within the SRD. For example, in an SRD based on a vpar that has psets for compartments, deploying an SRD in managed mode allows gWLM to actively migrate cores among psets. (A core is the actual data-processing engine within a processor. A single processor might have multiple cores.)</p> <p>When deploying an SRD in advisory mode, gWLM simply reports what the allocation would be—without actually affecting resource allocations on a system.</p> <p>Advisory mode is not available for SRDs containing virtual machines, psets, or fss groups due to the nature of these compartments.</p>

Undeploy

Disable gWLM's management of resources in a specified SRD.

If an SRD is in managed mode, undeploying stops the migration of system resources among workloads in the SRD. If the SRD is in advisory mode, gWLM no longer provides information on what requests would have been made.

The gWLM Management Model

gWLM enables utility computing across a data center by providing resource-sharing policies that you centrally create and monitor. gWLM moves resources among the workloads in a shared resource domain (SRD) as needed—based on the policies you specify.

gWLM allows you to manage resource allocations for several types of system divisions, as discussed below. These divisions are referred to as compartments in gWLM.

- HP-UX Hardware Partitions (npar)

A hardware partition, also known as an nPartition or npar, is a physical partition of a server, where each npar runs its own instance of the HP-UX operating system (which can include use of HP Integrity virtual machines) or is divided into virtual partitions.

Using the HP Instant Capacity product, gWLM simulates the movement of CPU resources among npars by turning off an active core in one npar then turning on a deactivated core in another npar in the same complex. Thus, the first npar has one less active core, while the second npar has one additional active core. (gWLM maintains the number of active cores, honoring the Instant Capacity usage rights. As a result, no additional costs are incurred.) Combining gWLM A.04.00.07 or later and an appropriate version of the iCAP software, gWLM's ability to manage npars using iCAP is extended across multiple complexes that are members of the same Global iCAP group.

- HP-UX Virtual Partitions (vpar)

A virtual partition is a software partition of a server or of a single nPartition, where each virtual partition runs its own instance of the HP-UX operating system. A virtual partition cannot span an nPartition boundary.

- HP Integrity Virtual Machines (hpvm)

Virtual machines are a robust soft-partitioning and virtualization technology that provides operating system isolation, with sub-core allocation granularity and shared I/O. These virtual machines can run a variety of operating systems. gWLM can manage a virtual machine regardless of the operating system running inside it.

- Processor sets (psets)

A processor set is a collection of cores (formerly known as CPUs) grouped together for the exclusive access by processes assigned to that processor set. Processor sets form partitions within a single operating system image.

- HP-UX Fair Share Scheduler groups (fss groups)

A group of processes that has its CPU resource allocation managed by the Fair Share Scheduler that is available with HP-UX. A benefit of fss groups is their granularity: You can allocate fractions of CPU resources, rather than only whole cores, to the group of processes. These groups form partitions within a single operating system image.

For more information on these partitions, visit:

- HP Virtual Server Environment web site:
<http://www.hp.com/go/vse>
 - The “Technical Documentation web site for HP Virtual Server Environment (VSE)” web site:
<http://docs.hp.com/en/vse.html>
 - The “Global Workload Manager” topic and the glossary in the online help for gWLM
- gWLM manages resources based on the following model:

1. You define an SRD by:
 - a. Deciding which of your systems you want to manage and what type of compartments you want to use. (gWLM manages existing npars, vpars, and virtual machines. It can manage your existing psets as well as create new ones. It creates fss groups for you.)
 - b. Associating each workload with a compartment. (For npars, vpars, and virtual machines, the compartment itself defines the workload. For psets and fss groups, you define the workload based on applications, users, or process IDs.)
 - c. Associating a policy with the workload indicating how gWLM should allocate resources to the workload's compartment. (gWLM comes with several policies and also lets you define your own. You can use a single policy for multiple workloads, minimizing the number of policies, if desired.)
2. Once the SRD is deployed:
 - a. gWLM monitors the CPU resource consumption of all the workloads in the SRD during the current allocation interval.
 - b. At the end of the interval, gWLM adjusts the CPU resource allocations for the compartments in accordance with the policies. It also makes the allocation data available for real-time and historical reports.
 - c. gWLM repeats the previous two substeps.

For information on what types of workloads to combine for optimal resource utilization, refer to the online help topic “Getting the Most Out of gWLM,” available in gWLM’s graphical interface in HP SIM.

How gWLM Allocates CPU Resources

gWLM addresses priority levels from highest to lowest, allocating resources to all requests at a given priority level before considering lower priority requests. If, at some priority level, all requests cannot be satisfied, the remaining resources are distributed so that the total resource allocation for each workload is as near the proportion of its weight relative to the sum of all the weights as possible. If gWLM has satisfied all resource requests at all priorities and there are resources still to be allocated, it will distribute the remaining resources by weight. Again, this is so that the total resource allocation for each workload is as near the proportion of its weight relative to the sum of all the weights as possible.

Table 1-1 lists the default weights for the various policy types. For policies with weights, you can also set the weight explicitly.

Table 1-1 Default Weights by Policy Type

Policy type	Default weight
Fixed	N/A (You cannot deploy an SRD where all the workloads with fixed policies are not satisfied.)
Utilization	1

Table 1-1 Default Weights by Policy Type (continued)

Policy type	Default weight
OwnBorrow	Equal to its owned value
Custom	1



NOTE: To ensure CPU resource allocations behave as expected for OwnBorrow policies, the sum of the CPU resources owned cannot exceed the number of cores in the SRD. (However, if the sum is less than the number of cores in the SRD, the excess is distributed to all compartments in proportion to the amounts owned. Thus, workloads will routinely get more than they are due.)

Available Interfaces

There are two interfaces for controlling and monitoring gWLM:

- HP Systems Insight Manager

A web-based interface accessed through the Shared Resource Domain tab reached through the **Tools**→**VSE Management** menu in HP Systems Insight Manager (Access HP Systems Insight Manager via the URL `http://hostname:280` where *hostname* represents the name of your Virtual Server Environment Management Software CMS).

To orient yourself with gWLM, refer to the gWLM Home page inside this interface by selecting **Tools**→**VSE Management** from the HP SIM menu bar and then the Shared Resource Domain tab and then from the VSE Management menu bar:

Tools→**Global Workload Manager**→**Getting Started - gWLM Home...**

(The HP SIM menu bar and VSE Management menu bar are discussed in the section “Tabs and Menus” (page 21).)

- `gwl` command

A command-line interface, described in *gwl*(1M) .

Other components of the command-line interface are: *vseinitconfig*(1M), *gwlmcmsd*(1M), *gwlagent*(1M), *gwlreport*(1M), *gwlmp*(1M), *gwlmp*(1M), *gwlmsend*(1M), *gwlmslconfig*(1M), *gwlmsstatus*(1M), and *gwlxml*(4).

2 Configuring gWLM to Manage Workloads

This chapter describes the various aspects of configuring gWLM to effectively manage the resources for your workloads.

Policy Types

You can define several types of policies to instruct gWLM how to manage the resources for your workloads. These types are:

Fixed	<p>Allocates a fixed (constant) amount of CPU resources to a workload's compartment.</p> <p>gWLM satisfies these policies before attempting to satisfy any other type of policies.</p>
Utilization	<p>Attempts to keep a workload's CPU utilization close to a target percentage by requesting more CPU resources when the workload is using too much of its current CPU resource allocation or by requesting fewer resources when the workload is using too little of its allocation. For example, assume a workload has a utilization policy with a target of 80% and an allocation of 5 cores. If the workload is consuming 4.5 cores, its utilization percentage is $4.5/5$, or 90%. gWLM would attempt to allocate additional CPU resources to the workload's compartment to meet the target. An allocation of 6 cores would result in a utilization percentage of $4.5/6$, or 75%, meeting the target.</p> <p>With a utilization policy, you specify the minimum and maximum CPU resource requests. Workloads with this type of policy are always allocated at least the minimum request. Utilization policies allow you to prioritize workloads.</p>
OwnBorrow	<p>Allows you to set the following values:</p> <ul style="list-style-type: none">• Amount of CPU resources, in cores, a workload's compartment owns.• Minimum amount of CPU resources, in cores, a workload's compartment must have (after lending resources to other workloads).• Maximum amount of CPU resources, in cores, a workload's compartment can have (after borrowing resources from other workloads). <p>The compartment of a workload with an OwnBorrow policy is allocated the owned CPU resources when needed. The minimum and maximum sizes allow you to specify how much the workload can lend (when resources are not needed) or borrow (when additional resources are needed and available). If a compartment has lent out cores and that compartment's workload becomes busy, the compartment re-acquires those lent-out cores.</p>
Custom	<p>Available for advanced users. For information on custom policies, refer to the online help or <i>gwlmxml(4)</i>.</p>
Conditional	<p>Specifies the existing policy to use when a time-based condition, a file-based condition, or a Serviceguard condition is met.</p>

You can define your own policies or use one of the numerous policies that come with gWLM. (You can use one policy for multiple workloads, minimizing the number of policies, if desired.)

Choosing a Policy Type

How do you decide which policy type to use? Table 2-1 answers this question for several common use cases. The section following the table helps you decide between using an OwnBorrow policy or a utilization policy.

Table 2-1 Choosing a Policy Type

If...	Use the following type of policy...
You want gWLM to allocate a constant amount of CPU resources to a workload.	Fixed
You have your own metric by which you want gWLM to manage a workload.	Custom
IT acts as a service provider to business units.	OwnBorrow This policy type allows you to set an owned amount of resources, while also giving you control over how workloads borrow and lend resources. gWLM provides a “topborrowers” report and a “resourceaudit” report to help you manage your data center using this model. For more information, see <i>gwlreport(1M)</i> .
You have static vpars, but you want to move to a model where cores migrate among vpars.	OwnBorrow For each vpar, set its number of owned cores to its static number of cores. The vpar gets those owned cores whenever needed.
You have npars but, you want to move to a model where CPU resources migrate among npars.	OwnBorrow Install the HP Instant Capacity product on each npar. (This software allows gWLM to simulate CPU resource movement among npars with spare capacity.) For each npar, set its number of owned cores to the number of cores you want the npar to have whenever needed.
You want to tap into a pool of resources taking or giving CPU resources as needed—with possibly no access to resources beyond a minimum request.	Utilization
You have a policy that should be in effect only for a given time period, for the duration of a file's existence, or for a certain Serviceguard condition.	Conditional Select an existing policy and a default policy and then set a time-based condition, set a file-based condition, or choose from the possible Serviceguard conditions.

Choosing Between an OwnBorrow Policy and a Utilization Policy

OwnBorrow and utilization policies both allocate resources to a workload based on the workload's use of its current allocation. Both policy types also specify minimum and maximum amounts of resources the workload should get. A workload with either type of policy can lend other workloads its unused resources—down to its minimum. (If the workload does not consume its entire minimum allocation, those unused resources are not available to other workloads.) OwnBorrow policies, however, provide greater control in lending resources because they also have an owned amount of resources. A workload always gets its owned resources back whenever needed. So, with an OwnBorrow policy, you can set a lower minimum allocation (increasing the amount of resources available for sharing among workloads), knowing the associated workloads get their owned resources whenever needed. Thus, an OwnBorrow policy provides greater flexibility in attempting to allocate a workload a certain amount of resources when needed while also lending those resources to other workloads when not needed.

Combining the Different Policy Types

Each workload in an SRD must have a policy. Starting with gWLM A.02.00.00.07, you can use any combination of the policy types within an SRD.

Seeing How gWLM Will Perform Without Affecting the System

gWLM provides an advisory mode that allows you to see how gWLM will approximately respond to a given SRD configuration—without putting gWLM in charge of your system’s resources. Using this mode, you can safely gain a better understanding of how gWLM works. In addition, you can check that your policies behave as expected—with minimal effect on the system.

Advisory mode is not available for SRDs containing virtual machines, psets, or fss groups due to the nature of these compartments.

Once you are comfortable with an SRD, change its mode to managed to let gWLM manage resource allocation for the compartments in the SRD.

For information on changing modes, refer to “Changing from Advisory Mode to Managed Mode” (page 26).

Getting Started with gWLM

gWLM is typically accessed through HP SIM. For information on the gWLM command-line interface, see *gwm(1M)*.

After performing the necessary gWLM daemon (or service) configuration as described in the *VSE Management Software Installation and Update Guide*, the quickest way to start using gWLM to manage new systems is to use the Manage Systems and Workloads wizard, as described in the following text.


Before you start the wizard though, decide:

- Which systems you want to manage with gWLM
- Whether you want to manage your workloads by migrating CPU resources among npars, vpars, virtual machines, processor sets, or fss groups. (CPU resource migration among npars with spare capacity is simulated using the HP Instant Capacity product, as explained in the section “The gWLM Management Model” (page 16).)

Tabs and Menus

The controls shown in Figure 2-1 appear at the top of the Global Workload Manager screen.

Figure 2-1 Top Controls of the Global Workload Manager Screen



- 1 The SIM menu bar
- 2 The Virtualization Manager tabs
- 3 The VSE Management menu bar

These menu bars will be referenced later.

Using the Wizard

To start the wizard:



NOTE: You must be logged in as root on the systems where you run the `mxstart`, `gwlmcmd`, and `gwlmagent` commands mentioned below. In HP SIM, you must be logged in as root or have authorizations for “All Tools” or “VSE All Tools.”

1. Configure your CMS as indicated in the *VSE Management Software Installation and Update Guide* if you have not already done so
2. On each managed node, start the gWLM agent if it is not already running:

```
# /opt/gwlm/bin/gwlmagent
```

Alternatively, you can start the agent through HP SIM, as discussed in the *VSE Management Software Installation and Update Guide*.
3. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
where *hostname* represents the hostname of the CMS.
4. Select:
Tools→VSE Management
and then click the Shared Resource Domain tab. From the VSE Management menu bar, select:
Create→Shared Resource Domain

The wizard guides you through the following steps:

1. Specify the hosts on which to run workloads that you want gWLM to manage as part of one SRD.
2. Set SRD properties.
Properties include the SRD name, mode, use of Temporary Instant Capacity (if available on the system), and resource allocation interval.
3. Specify workload and policy settings.
Settings include the workload name and policy.
4. Review and confirm the SRD.
5. Verify the SRD is configured as expected, and click **Finish** to have gWLM manage the resource allocation for the workloads in the SRD.

Seeing gWLM in Action

This section helps you see gWLM move CPU resources among vpars. You can use similar steps to see CPU resources move among npars, virtual machines, psets, or fss groups. For psets and fss groups, though, you will need to put processes in the desired pset or fss group. (Place processes by modifying the workload definition or by using the `gwlmplace` command.) In this example:

- The gWLM agent is used on two vpars, which we will call `vpar1` and `vpar2`.
- These vpars are idle and have a number of unbound cores that gWLM can move among them.
- HP SIM and the gWLM CMS software are installed, configured, and running on a vpar called `vpar3`.

To see gWLM in action:



NOTE: You must be logged in as root on the systems where you run the `mxstart`, `gwlmcmdsd`, and `gwlmagent` commands mentioned below. In HP SIM, you must be logged in as root or have authorizations for “All Tools” or “VSE All Tools.”

1. Start the gWLM agent daemons on vpar1 and vpar2:

```
# vpar1> /opt/gwlm/bin/gwlmagent  
# vpar2> /opt/gwlm/bin/gwlmagent
```

Alternatively, you can start the agents through HP SIM, as discussed in the *VSE Management Software Installation and Update Guide*.
2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
where *hostname* represents the hostname of the CMS, in this case, vpar3.
3. Create a gWLM SRD containing the two vpars by following the steps given in “Getting Started with gWLM” (page 21).
 - a. For hosts, enter the names of the two vpars separated by a space in the field beneath the table.
 - b. Ensure when setting the SRD properties that the mode is Managed.
 - c. Use an OwnBorrow policy for the vpar1 workload and for the vpar2 workload.
An OwnBorrow policy has a name of the form:
`Owns_4-Max_8`
Ideally, the sum of the cores owned by the workloads will equal the total number of cores in the SRD when using this type of policy. (You may need to edit policies to achieve this equality.)
 - d. Confirm and finish the SRD creation. (You are then placed on the Shared Resource Domain View, showing the newly created SRD and its workloads.)
4. Select the vpar1 workload.
5. View gWLM’s real-time reports to show CPU resource allocation for vpar1 by selecting in the VSE Management menu bar:
Report→gWLM Real-time Reports...
You can also view the reports by clicking the workload’s bar graph.
6. Click the **Policy graph** radio button to see a graph of vpar1’s CPU resource allocation.
7. Start a CPU resource-intensive workload in vpar1.
If you already have such a workload configured in vpar1, start it now.
If you need such a workload, the following command prints its PID (24379 in this case) and then just consumes CPU resources:

```
# perl -e 'print "$$\n";while (1) {};' &  
[1] 24379
```

This command consumes most of a single core. Start multiple copies of the command to consume additional cores.
8. Wait a few minutes, then look at the “Policy Request and Workload Allocation” graph from Step 6 to see how the policy for vpar1’s workload is requesting CPU resource allocations and how gWLM is granting them.
9. Kill the workloads you started in vpar1.
10. Set the **Workload** dropdown near the top of the screen to the vpar2 workload and repeat Step 7 through Step 9 to see gWLM move cores to vpar2.

Common Uses for gWLM

gWLM is a powerful tool that allows you to manage your systems in numerous ways. The following sections explain some of the more common tasks that gWLM can do for you.

Fixing the Amount of CPU Resources a Workload Gets

gWLM allows you to give a workload a fixed amount of CPU resources. This fixed amount is in the form of a set amount of CPU resources given to an npar, a vpar, a virtual machine, a pset, or an fss group.

To fix the amount of CPU resources a workload gets, use a fixed policy provided by gWLM or create your own. Associate a fixed policy with a workload:

- When creating an SRD, as described in “Getting Started with gWLM” (page 21)
- When adding a workload to an SRD, as described in “Adding a New Compartment or GiCAP Group Member to an SRD” (page 28)
- By changing the policy associated with an existing workload, as described in “Changing Which Policy is Associated with a Workload” (page 27)

Resizing a Workload’s Npar, Vpar, Virtual Machine, Pset, or FSS Group as Needed

To ensure a workload gets the CPU resources it needs—while also allowing resource sharing when possible—gWLM provides OwnBorrow policies.

With such a policy, you indicate the amount of CPU resources a workload should own. The workload is then allocated this owned amount of CPU resources—when it needs it. However, you can configure the workload to:

- Lend CPU resources to other workloads when it is idle
- Borrow CPU resources from workloads that are idle

Associate an OwnBorrow policy with a workload:

- When creating an SRD, as described in “Getting Started with gWLM” (page 21)
- When adding a workload to an SRD, as described in “Adding a New Compartment or GiCAP Group Member to an SRD” (page 28)
- By changing the policy associated with an existing workload, as described in “Changing Which Policy is Associated with a Workload” (page 27)

gWLM’s utilization policies also allow resizing.

Common Configuration Tasks

This section discusses various configuration tasks:

- “Changing from Advisory Mode to Managed Mode” (page 26)
- “Creating a New Policy” (page 26)
- “Editing a Policy” (page 27)
- “Changing Which Policy is Associated with a Workload” (page 27)
- “Adding a New Compartment or GiCAP Group Member to an SRD” (page 28)
- “Stop Managing a Workload” (page 28)
- “Stop Managing an SRD” (page 29)

Setting up gWLM (Initial Setup Steps)

Several of the configuration tasks require the same initial set-up steps. (Each task requiring these steps indicates that the steps are needed.) These steps are given below.



NOTE: You must be logged in as root on the systems where you run the `gwlagent` command mentioned below.

1. Configure your CMS as indicated in the *VSE Management Software Installation and Update Guide*, if you have not already done so.
2. On each managed node, start the gWLM agent (if it is not already running):

```
# /opt/gwlm/bin/gwlagent
```

Alternatively, you can start the agents through HP SIM, as discussed in the *VSE Management Software Installation and Update Guide*.

Changing from Advisory Mode to Managed Mode

Advisory mode allows you to see what CPU resource requests gWLM would make for a workload—without actually affecting resource allocation. (Advisory mode is not available for SRDs containing virtual machines, psets, or fss groups due to the nature of these compartments.) Managed mode, however, allows gWLM to automatically adjust the resource allocations for your defined workloads.

To change from one mode to the other:



NOTE: In HP SIM, you must be logged in as root or have authorizations for “All Tools” or “VSE All Tools.”

NOTE: If you are changing from managed mode to advisory mode and you do not plan to change back soon, be aware that gWLM leaves the npar and vpar and pset compartments with the number of cores they had in the last allocation interval. Set the compartments to your desired sizes before changing to advisory mode by associating fixed policies with all the compartments and waiting for an allocation interval to pass.

1. Ensure HP SIM, the gWLM CMS daemon or service (`gw1mcmsd`) and all the gWLM agents (`gw1magent`) are still running, as explained in the section “Setting up gWLM (Initial Setup Steps)” (page 25).
2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
where *hostname* represents the hostname of the CMS.
3. From the HP SIM menu bar, select:
Tools→**VSE Management**
and then click the **Shared Resource Domain** tab.
4. Select the SRD for which to change the mode.
5. From the VSE Management menu bar, select:
Modify→**Shared Resource Domain**
6. Change to the desired mode.
7. Click **OK**.

Quick Link Options

In the previous procedure, instead of selecting an SRD and using the VSE Management menu bar, you can find the Details table for the SRD and then choose one of the following options:

- Click the **Change SRD to advisory mode** link
- Click the **Modify SRD** link

Creating a New Policy

A policy instructs gWLM how to manage a workload’s resources. You can create a policy when managing a workload or create a policy separately. To create a policy separately:



NOTE: In HP SIM, you must be logged in as root or have authorizations for “All Tools” or “VSE All Tools.”

1. Ensure HP SIM, the gWLM CMS daemon or service (`gw1mcmsd`) and all the gWLM agents (`gw1magent`) are still running, as explained in the section “Setting up gWLM (Initial Setup Steps)” (page 25).
2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`

where *hostname* represents the hostname of the CMS.

3. From the HP SIM menu bar, select:
Tools→VSE Management
and then click the Shared Resource Domain tab.
4. From the VSE Management menu bar, select:
Policy→Create gWLM Policy...
5. Edit the settings, selecting a policy type and specifying the required values and optional values as desired.
6. Click **OK**.

Editing a Policy

A policy instructs gWLM how to manage a workload's resources.



NOTE: You can edit the policies provided with gWLM; however, there is currently no way to restore these policies to their original definitions.

To edit a policy:



NOTE: In HP SIM, you must be logged in as root or have authorizations for "All Tools" or "VSE All Tools."

1. Ensure HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`), and all the gWLM agents (`gwlagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 25).
2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
where *hostname* represents the hostname of the CMS.
3. From the HP SIM menu bar, select:
Tools→VSE Management
4. From the VSE Management menu bar, select:
Policy→Edit gWLM Policies...
5. Select the policy to edit.
6. Click **Edit**.
7. Edit the settings.
8. Click **OK**.



NOTE: All workloads associated with this policy will automatically use the updated policy.

Changing Which Policy is Associated with a Workload

To change the policy affecting how gWLM allocates resources to a workload:



NOTE: In HP SIM, you must be logged in as root or have authorizations for "All Tools" or "VSE All Tools."

1. Ensure HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`) and all the gWLM agents (`gwlagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 25).

2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
 where *hostname* represents the hostname of the CMS.
3. From the HP SIM menu bar, select:
Tools→VSE Management
 and then click the **Shared Resource Domain** tab.
4. Select the shared resource domain containing the workload for which you want to change the policy.
5. Select the workload for which you want to change the policy.
6. From the VSE Management menu bar, select:
Policy→Change Associated gWLM Policy...
7. From the Policy dropdown in the table row for the workload, select the new policy to associate, or apply, to the workload.
8. Click **OK**.

Adding a New Compartment or GiCAP Group Member to an SRD

If you:

- Have added an *npar*, a *vpar*, or a virtual machine to your system and want to add it to an SRD,
- Have added a Global Instant Capacity (GiCAP) group member and want to add it to an SRD, or
- Want to create *psets* or *fss* groups in a host already in an SRD

You can use the gWLM wizard to accomplish those tasks. To start the wizard, select from the HP SIM menu bar:

Tools→VSE Management

and then click the **Shared Resource Domain** tab. From the VSE Management menu bar, select:

Create→Shared Resource Domain

Step 1 in the wizard allows you to add *npars*, *vpars*, and GiCAP group members. Step 3 allows you to create *psets* or *fss* groups, as well as manage existing virtual machines.

Stop Managing a Workload

When you stop managing a workload:

- gWLM stops managing resources for the workload
- The workload's definition is removed from the SRD, although it remains available for placing in another SRD



NOTE: When gWLM stops managing *npar*-based or *vpar*-based workloads, it leaves the *npar*s or *vpars* with the number of cores they had in the last allocation interval. For this reason, in Step 3 below, you associate fixed policies with the workloads based on these types of compartments. You must stop a virtual machine before you stop managing it with gWLM. When gWLM stops managing a virtual machine, it sets the entitlement of the running virtual machine to its minimum. For *psets* and *fss* groups, gWLM removes the *pset* or *fss* group and moves the processes from that compartment to the default compartment.

To stop managing workloads in an SRD:

1. Ensure HP SIM, the gWLM CMS daemon or service (`gw1mcmsd`), and all the gWLM agents (`gw1magent`) are still running, as explained in the section “Setting up gWLM (Initial Setup Steps)” (page 25).
2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
where *hostname* represents the hostname of the CMS.
3. Associate fixed policies with all workloads that you want to unmanage that are based on npars or vpars.
For information on setting the associated policy, refer to “Changing Which Policy is Associated with a Workload” (page 27).
4. Wait an allocation interval for gWLM to set CPU resource allocations based on the fixed policies.
5. Click the **Shared Resource Domain** tab.
6. Select the workload you want to stop managing in the workload table.
7. From the VSE Management menu bar, select:
Policy→Remove Associated gWLM Policy...
8. Associate policies.
Evaluate and change, if needed, the remaining workloads and their associated policies to ensure they are appropriate, given that a workload has been removed.
9. Click **OK**.

Stop Managing an SRD

To stop gWLM from managing an SRD and its workloads, returning resource allocation to HP-UX:



NOTE: In HP SIM, you must be logged in as root or have authorizations for “All Tools” or “VSE All Tools.”

1. Ensure HP SIM, the gWLM CMS daemon or service (`gw1mcmsd`), and all the gWLM agents (`gw1magent`) are still running, as explained in the section “Setting up gWLM (Initial Setup Steps)” (page 25).
2. Connect to HP SIM by pointing your web browser to:
`http://hostname:280`
where *hostname* represents the hostname of the CMS.
3. Associate fixed policies with all npars or vpars that were in the SRD.
When gWLM stops managing an SRD, it leaves compartments based on npars or vpars as they were in the last allocation interval. Associating fixed policies allows you to set the sizes exactly to what you want. (For virtual machines, gWLM sets the entitlements of the running virtual machines to their minimums. psets and fss groups are removed in this situation, with their processes going to the default pset or default fss group.)
For information on setting the associated policy, refer to “Changing Which Policy is Associated with a Workload” (page 27)
4. Click the **Shared Resource Domain** tab.
5. Select the SRD that you want to stop managing (undeploy).
6. From the VSE Management menu bar, select:
Modify→Shared Resource Domain
7. Change to the Undeployed state.
8. Click **OK**.

Quick Link Option

In the previous procedure, instead of selecting an SRD and using the VSE Management menu bar, you can find the Details table for the SRD and click the **Undeploy SRD** link.

3 Monitoring Workloads and gWLM

This chapter describes how to monitor workloads and gWLM.

Monitoring Workloads

There are several methods for monitoring workloads, as described below.

High-Level View

To see a high-level view of the performance of your SRDs and workloads:

1. From the HP SIM menu bar, select:
Tools→VSE Management
2. Click the **Shared Resource Domain** tab.

Graphical Reports

gWLM provides graphs showing either real-time or historical data through HP SIM. For information on interpreting these reports, refer to the online help.

Real-time Reports

To view real-time reports:

1. From the HP SIM menu bar, select:
Tools→VSE Management
2. Click the **Shared Resource Domain** tab.
3. Select a workload.
4. From the VSE Management menu bar, select:
Report→gWLM Real-time Reports...

Quick Link Option

In the previous procedure, instead of selecting a workload and using the VSE Management menu bar, you can click the workload's "CPU Utilization" bar graph.

Historical Reports

To view historical reports:

1. From the HP SIM menu bar, select:
Tools→VSE Management
2. Click the **Shared Resource Domain** tab.
3. From the VSE Management menu bar, select:
Report→gWLM Historical Reports...

Viewing gWLM Reports in Monitor-Only Mode

gWLM allows you to specify users who should be allowed only to monitor gWLM reports. These users do not have the ability to change gWLM configurations. To set up a user with monitor-only privileges, refer to the online help topic "Authorizations and Read-only Monitoring."

Monitoring gWLM from the Command Line

There are several command-line tools for monitoring gWLM. These commands are added to the path during installation. On HP-UX systems, the commands are in /opt/gwlm/bin/. On Microsoft Windows systems, the commands are in C:\Program Files\HP\Virtual Server Environment\bin\gwlm\ by default. However, a different path may have been selected at installation.



NOTE: You must be logged in as root on HP-UX or into an account that is a member of the Administrators group on Windows to run the commands below.

- `gwlm monitor`

The `gwlm` command, available only on a CMS, has a `monitor` subcommand that displays policy, workload, and SRD statistics.

- `gwlmreport`

This command, available only on a CMS, provides various types of reports: `topborrowers`, `resourceaudit`, `abnormalutil`, and `extract` (which provides data in comma-separated values for use in other tools). It also provides an `ovpafeed` option that extracts data for use with OpenView Performance Agent via data source integration (DSI). In addition, it can generate a `config` report showing the history of configuration changes.

- `gwlmstatus`

The `gwlmstatus` command, available only on managed nodes, displays status information for a managed node's agent and SRD. The information displayed includes:

- Whether `gwlmagent` is running
- The version of the installed `gwlmagent`
- The SRD to which the current node belongs (if a member of a deployed SRD)
- Master node of the deployed SRD
- Whether hosts in the SRD are nonresponsive
- Whether the host's SRD is the most recently deployed SRD on the CMS (This knowledge can be useful in determining failed deployments.)
- Whether any hosts are unable to rejoin the SRD

For more information, see *gwlm(1M)*, *gwlmreport(1M)*, or *gwlmstatus(1M)*.

Message Logs

For messages about gWLM's ongoing operations, gWLM logs error and informational messages, as well as statistics, to one of several log files:

Table 3-1 gWLM Log Files

Log for	Location
<code>gwlmcmds</code> daemon or service	HP-UX: /var/opt/gwlm/gwlmcmds.log.0 Windows: C:\Program Files\HP\Virtual Server Environment\logs\gwlmcmds.log.0
<code>gwlmagent</code> daemon	HP-UX: /var/opt/gwlm/gwlmagent.log.0 Windows: Not applicable

Table 3-1 gWLM Log Files *(continued)*

Log for	Location
gWLM interface in HP Systems Insight Manager	HP-UX: /var/opt/gwlm/gwlm.log.0 Windows: C:\Program Files\HP\Virtual Server Environment\logs\gwlm.log.0
gwlm command	HP-UX: /var/opt/gwlm/gwlmcommand.log.0 Windows: C:\Program Files\HP\Virtual Server Environment\logs\gwlmcommand.log.0



NOTE: On systems running Windows, log files are in C:\Program Files\HP\Virtual Server Environment\logs\ by default. However, a different path may have been selected at installation.

The name of the current log always ends in .log.0. Once this file grows to a certain size, it is moved to a filename ending in .log.1 and a new .log.0 file is started. If a .log.1 file already exists, it is renamed .log.2. If a .log.2 file already exists, it is overwritten.

By default, the log file size is limited to 20 Mbytes and the number of log files is limited to 3. You can change these defaults using the following properties:

```
com.hp.gwlm.util.Log.logFileSize = 20
```

```
com.hp.gwlm.util.Log.logNFiles = 3
```

For the gwlmagent log, change the values of these properties in /etc/opt/gwlm/conf/gwlmagent.properties. For all the other log files, change the values in /etc/opt/gwlm/conf/gwlmcms.properties on HP-UX and in C:\Program Files\HP\Virtual Server Environment\conf\gwlmcms.properties on Windows. (The given Windows path is the default; however, a different path may have been selected at installation.)

Viewing HP Systems Insight Manager Events

gWLM allows you to configure a number of events you can monitor through HP SIM. Set these events in HP SIM as follows:

1. From the HP SIM menu bar, select:
Tools→VSE Management
2. Click the **Shared Resource Domain** tab.
3. From the VSE Management menu bar, select:
Tools→Global Workload Manager→Events...

After configuring these events, you can monitor them through the various Events items in the left pane of HP SIM. Also, the Shared Resource Domain View provides links to any uncleared gWLM events, filtering out all other events.

Monitoring gWLM with GlancePlus

You can use HP's optional performance and monitoring tool GlancePlus to view data on workloads based on fss groups that are managed by gWLM. (Use GlancePlus C.03.35.00 or later. Also, you must have the bundle PRMLibraries C.03.03 or later installed. This bundle is available from <http://www.hp.com/go/softwaredepot>.)

GlancePlus has both a text interface (`glance`) and an X-Windows interface (`gpm`).

Items to be aware of when using GlancePlus:

- Data is displayed only for workloads based on fss groups. Data for the workload with FSS ID 0 is shown without a name.
- gWLM workloads are referred to as 'PRM Groups' in GlancePlus.
- The workload names displayed by GlancePlus may vary slightly from those used in gWLM due to differences in characters allowed in names.
- GlancePlus data is in terms of shares percentages of CPU resources, whereas gWLM data is in terms of physical cores.
- GlancePlus and gWLM use different methods for gathering resource utilization, so data between the products does not always match.
- Data is not displayed for more than 64 workloads.

4 Security

This chapter highlights several security items you should be aware of.

General Security Topics

The following items are a few general topics on security:

- HP provides the HP-UX Bastille product, available from <http://software.hp.com> at no charge, for enhancing system security.
- You can secure gWLM's communications as explained in the following section.
- HP SIM allows you to create user roles with different levels of privileges. For more information, see the HP SIM documentation.

For information on authorizations needed to run the VSE Management Software, see the *VSE Management Software Version 4.0 Getting Started Guide* or the online help topic "Authorizations and Read-only Monitoring."

Securing gWLM Communications

By default, gWLM's communications are not secure, meaning:

- The communications between the CMS and the managed nodes are not encrypted
- The source and destination of gWLM's communications are not authenticated

When securing communications, you must do so for every managed node in every SRD managed by a given CMS.

To secure gWLM's communications, assuming OpenSSH is installed and configured for HP SIM on each of the managed nodes, select from the HP SIM menu bar:

Configure→Configure VSE Agents→Configure gWLM Communications...

For more information, see the online help topic "Securing gWLM Communications."

Alternatively, you can secure communications manually by following the steps outlined in *gwlmsslconfig(1M)*.



NOTE: HP strongly recommends always using gWLM with its communications secured.

5 Additional Configuration and Administration Tasks

This chapter covers various configuration and administration tasks.

Manually Adjusting CPU Resources

When an SRD is created, it has a certain number of cores. gWLM manages the SRD using the same number of cores. If the SRD—or a policy used in the SRD—is configured to use Temporary Instant Capacity (TiCAP), gWLM can automatically activate that additional capacity to meet policies. If neither the SRD or its policies are configured to use TiCAP, you may be able to temporarily provide additional resources to a deployed SRD by:

- Using an available core from the vpar monitor free pool.
- Activating an iCAP core.
- Deleting a core from an unmanaged vpar and then adding it to a vpar in the SRD.
- Deactivating a core in an npar and then activating one in an npar in the SRD.



NOTE: If gWLM detects activated cores for which there is no request, it deactivates them to avoid spending money on the unneeded capacity.

NOTE: After you manually change system resources (by modifying unmanaged partitions or changing bindings, for example), you might see resize errors on one or more of the managed nodes. However, gWLM should recover (and stop issuing errors) by the next resource allocation interval—unless gWLM can no longer access the required resources.

NOTE: Deployed SRDs do not accept manual decreases in the available resources. gWLM will attempt to reclaim any removed resources.

NOTE: Although a deployed SRD might recognize added resources, policy maximum values are still in effect and can clip resource requests. Consider adjusting policy settings to use the added resources.

As already mentioned, gWLM can take advantage of the additional CPU resources only temporarily. To take full, persistent advantage of the extra resources using the gWLM interface in HP SIM:

1. Modify the size of the SRD.
 - a. Select the SRD affected by the additional resources in the Shared Resource Domain View.
 - b. Select the menu item **Modify**→**Shared Resource Domain**
 - c. Click the **Workload and Policies** tab.
 - d. Adjust the size of the SRD by editing the value, beneath the table, labeled “Total Size.”
 - e. Click **OK**.
2. Edit policies used in the SRD to ensure they do not unintentionally limit their associated workloads' resource requests.

To take full, persistent advantage of the extra resources using the gWLM command-line interface:

1. Undeploy the SRD containing the systems that were adjusted.
2. Re-create and re-deploy the SRD.
3. Ensure policies used in the SRD do not unintentionally limit their associated workloads' resource requests.

gWLM cannot take advantage—even temporarily—of resources added by:

- Adjustments to entitlements for virtual machines.
- Changes to a virtual machine's number of virtual CPUs while gWLM is managing the virtual machine.

- Creation or deletion of a pset using `psrset` on a system where gWLM is managing pset compartments.
- Performing online cell operations using `parolrad`
- Enabling and disabling Hyper-Threading.

To make use of these additional resources using the gWLM command-line interface:

1. Undeploy the SRD containing the systems that you want to adjust.
2. Make your adjustments.
3. Re-create and re-deploy the SRD
4. Ensure policies used in the SRD do not unintentionally limit their associated workloads' resource requests.

To make use of these additional resources using the gWLM interface in HP SIM, follow the procedure given for that interface above.



NOTE: NOTE: After manually adjusting the number of cores in an SRD, always confirm the changes after two gWLM resource allocation intervals have passed. Changes may not be as expected due to gWLM behaviors such as the ones described below.

- In an SRD with nested partitions, gWLM samples the inner partitions for their sizes before sampling the outer partitions. Adjusting resources between these samplings can cause gWLM to report incorrect sizes. If you encounter this issue, try making your adjustment again.
 - In an SRD with nested partitions that includes vpars, assume you manually add cores from an unmanaged vpar. If you later remove those cores—without returning them to an unmanaged vpar before gWLM samples compartment sizes—those cores are deactivated.
-

Manually Adjusting Memory Resources

The `vparmodify` command enables you to move memory from one vpar to another. However, gWLM cannot move CPU resources while a `vparmodify` operation is in progress. If a memory move takes longer than gWLM's resource allocation interval, gWLM will not be able to satisfy CPU resource requests for the missed intervals. gWLM resumes allocating resources once the memory move is complete.

Setting Aside Space for Historical Data

HP recommends that you allocate 4 GB of storage for every 100 workloads you will manage with gWLM. With a 5-minute sampling interval, this is enough space to store 2 years of data, which you can then use for capacity planning and performance management.

On HP-UX, the provided HP System Management Database (HPSMDB), also known as PostgreSQL, stores its data in the `/var/opt/` file system. On HP-UX and Windows systems using Oracle, set aside the space in the file system used by the configured database.

Backing up the VSE Management Software Database



NOTE: This section applies only to the provided HP System Management Database (HPSMDB).

The VSE Management Software database contains configuration data as well as historical performance data. To create a backup of your gWLM database, use the `vseinitconfig --backup` command. To make use of a backup file, use the `vseinitconfig --restore` command.

For more information, see `vseinitconfig(1M)`.

Tips for Backup and Restore

Here are a few tips to make the best use of backup and restore:

- To ensure the latest gWLM data is backed up, issue the following command before using the `--backup` option:
`gwlmcms history --flush`
- If the CMS or the managed nodes have been modified (such as changes in number of cores or hostnames) between a backup and a restore or you are trying to do a restore on a different CMS, gWLM will probably not function as expected after the restore.

Setting gWLM Properties

gWLM provides two properties files that allow you to control various gWLM behaviors. One file is for the CMS daemon or service, and the other is for use on all the managed nodes. Read the files for information on the behaviors they control.

CMS Properties

The CMS properties are in `/etc/opt/gwlm/conf/gwlmcms.properties` on HP-UX and in `C:\Program Files\HP\Virtual Server Environment\conf\gwlmcms.properties` on Windows. (The given Windows path is the default; however, a different path may have been selected at installation.)



NOTE: Some values are read by gWLM only when a daemon or service is started; while other values are read when an SRD is deployed. Refer to the file for information on when individual properties are read.

Restarting `gwlmcmd` temporarily disables HP Virtualization Manager and HP Capacity Advisor.

The `gwlmcms.properties` file is shown below.

```
#
# (C) Copyright 2004-2008 Hewlett-Packard Development Company, L.P.
#
# $Date: 2008/02/01 19:00:32 $
# $Revision: 1.28 $
#
# Contents:
# This file contains the default configuration values for the
# Global Workload Manager CMS.
#
# You must restart gwlmcmd for changes made to this file
# to take effect--unless otherwise noted.
#
#
# Set FileHandler log level for the following log files:
#   /var/opt/gwlm/gwlmcmd.log.0
#   /var/opt/gwlm/gwlmcommand.log.0
#   /var/opt/gwlm/gwlm.log.0
#   /var/opt/vse/logs/gwlminitconfig.log.0
# Valid levels, from most severe to least, are:
#   SEVERE
#   WARNING
#   INFO
#   CONFIG
#   FINE
#   FINER
#   FINEST
# When you set the level, you will see messages only from that level and
# the levels that are more severe. So, the SEVERE level produces the fewest
```

```

# messages, while the FINEST level includes messages from all seven levels.
#
com.hp.gwlm.util.Log.logLevel = INFO

#
# Specify the size (in MB) and number of files to use
# for logging. For a single file of unlimited size, set
# logFileSize to negative one (logFileSize=-1).
# Otherwise, total log file size is
#   logFileSize * logNFiles
#
com.hp.gwlm.util.Log.logFileSize = 20
com.hp.gwlm.util.Log.logNFiles = 3

#
# Support for automatic database statistics gathering. These properties
# control how often row-level statistics are gathered from the database in
# order to optimize performance.
#
#   com.hp.gwlm.cms.db.analyze.time:
#       Frequency, in minutes, in which statistics are gathered. The
#       default is to attempt to gather database statistics every 60
#       minutes. When the analysis runs, statistics will only be gathered
#       if a certain number of transactions have been processed (which is
#       configured in the property below).
#
#   com.hp.gwlm.cms.db.analyze.limit:
#       Number of consecutive transactions that must take place before a
#       database analysis is performed.
#
com.hp.gwlm.cms.db.analyze.time = 60
com.hp.gwlm.cms.db.analyze.limit = 50

#
# Support for the database cache on the CMS.
#
#   cachesize:
#       The number of historical configurations to cache in memory.
#       A larger historical configuration cache reduces time spent
#       in database lookups. The valid range is 1-1000.
#
com.hp.gwlm.cms.cachesize = 100

#
# Support for local data caching on a remote node for report generation.
# These properties are defined on the CMS but are pushed out to the remote
# node managers during deployment of an SRD. The cached objects on the
# agent consume an amount of memory approximated by:
#
#   Memory = 5000 * workloads * cachesize * (60 / resource_domain_interval)
#
# bytes of memory. For example, if there are 4 workloads deployed with a
# 15 second interval and a cachesize of 20 minutes, the agent will need:
#
#   Memory = 5000 * 4 * 20 * (60 / 15) ~ 2.5 MB.
#
#   cachesize:
#       The number of minutes of real-time data to maintain on the remote
#       node for future CMS access. This value must be at least three
#       times the 'samples' value specified below. The default value is
#       20 minutes.
#
#   samples:
#       The number of minutes of real-time data used to aggregate into a
#       historical data point. The default is to aggregate the data into

```



```

# 5-minute averages.
#
com.hp.gwlm.node.cachesize = 20
com.hp.gwlm.node.samples   = 5

#
# Support for real-time graphing properties.
#
#   viewport:
#       The size of the displayed real-time graph (in minutes).
#
#   refresh:
#       The refresh rate of the real-time graphs and tables (in seconds).
#
com.hp.gwlm.ui.monitor.viewport = 20
com.hp.gwlm.ui.monitor.refresh  = 15

```

Agent Properties

The agent properties are in `/etc/opt/gwlm/conf/gwlmagent.properties` on HP-UX and in `C:\Program Files\HP\Virtual Server Environment\conf\gwlmagent.properties` on Windows. (The given Windows path is the default; however, a different path may have been selected at installation.)



NOTE: You must restart the `gwlmagent` daemon on each managed node where you have modified the properties file for the changes to the file to take effect.

The `gwlmagent.properties` file is shown below.

```

#
# (C) Copyright 2004-2007 Hewlett-Packard Development Company, L.P.
#
# $Date: 2007/04/12 15:46:19 $
# $Revision: 1.17 $
#
# Contents:
# This file contains the default configuration values for the
# Global Workload Manager Agent on a given managed node. The
# agent on each managed node uses the default values unless you
# edit that node's gwlmagent.properties file.
#
# You must restart gwlmagent for changes made to this file to
# take effect.
#

#
# Set FileHandler /var/opt/gwlm/gwlmagent.log.0 log level.
# Valid levels, from most severe to least, are:
# SEVERE
# WARNING
# INFO
# CONFIG
# FINE
# FINER
# FINEST
# When you set the level, you will see messages only from that level and
# the levels that are more severe. So, the SEVERE level produces the fewest
# messages, while the FINEST level includes messages from all seven levels.
#
com.hp.gwlm.util.Log.logLevel = INFO

#
# Specify the size (in MB) and number of files to use
# for logging. For a single file of unlimited size, set

```

```

# logFileSize to negative one (logFileSize=-1).
# Otherwise, total log file size is
#   logFileSize * logNFiles
#
com.hp.gwlm.util.Log.logFileSize = 20
com.hp.gwlm.util.Log.logNFiles = 3

#
# Set the number of seconds for the high availability (HA) minimum
# timeout. Communication is considered lost on a managed node
# once the timeout period has elapsed.
#
# By default, communication is not considered lost until 10 allocation
# intervals have elapsed without communication. The default value of the
# property (60) is used only when the allocation interval is less than 6
# seconds.
#
com.hp.gwlm.node.HA.minimumTimeout = 60

#
# Enable/disable use of Instant Capacity (iCAP) to simulate movement of
# cores across nPartitions. To use iCAP, you must enable this property on
# each managed nPartition where you want to take advantage of iCAP. Also,
# each nPartition must meet the requirements outlined in the online help
# topic "Getting the most out of gWLM" as well as in the section "USING
# NPARS AS COMPARTMENTS IN AN SRD" in the gwlmxml(4) man page on HP-UX or
# the gwlmxml(5) man page on Linux.
#
com.hp.gwlm.platform.icap.manageWithIcap = on

#
# Set the minimum number of Temporary Instant Capacity (TiCAP) minutes
# that must be available for TiCAP activation. gWLM will stop using
# TiCAP when the available balance goes below this threshold. The
# same value should be set on all agents managing the SRD. To use
# TiCAP, it must be available on the complex and enabled at the policy
# level or the SRD level.
#
com.hp.gwlm.node.ticap.minimumBalance = 30

```

Communications Ports

gWLM uses the following ports for communications:

Managed nodes: 9617

CMS: 9618

If you need to change these ports, add the following lines:

```
com.hp.gwlm.cms.port = portX
```

```
com.hp.gwlm.node.port = portY
```

to both properties files:

- `gwlmcms.properties`

On HP-UX, this file is in `/etc/opt/gwlm/conf/`. On Windows, it is in `C:\Program Files\HP\Virtual Server Environment\conf\`. (The given Windows path is the default; however, a different path may have been selected at installation.)

- `/etc/opt/gwlm/conf/gwlmagent.properties`

The `portX` and `portY` values cannot be the same value.

The `com.hp.gwlm.cms.port` property must have the same `portX` value in all the properties files—across all managed nodes and on the CMS. Similarly, the `com.hp.gwlm.node.port` property

must have the same *portY* value in all the properties files—across all managed nodes and on the CMS.

You must restart `gwlmcmsd` and `gwlagent` for the changes to take effect. If you are using gWLM through HP SIM, you must also restart SIM. (On HP-UX, restart SIM daemons using `/opt/mx/bin/mxstop` and `mxstart`. On Windows, restart the HP Systems Insight Manager service.)



NOTE: Restarting `gwlmcmsd` temporarily disables HP Virtualization Manager and HP Capacity Advisor.

Controlling gWLM's Startup Behavior

The gWLM CMS daemon or service (`gwlmcmsd`) is set to start at boot by `vseinitconfig`. Generally, you should not change this setting as it disables all VSE Management Software from starting at boot.

On HP-UX, you can control whether the gWLM CMS daemon (`gwlmcmsd`) and its agent daemons on the managed nodes (`gwlagent`) start at boot by:

- Manually setting variables in the `/etc/rc.config.d/gwlmCtl` file. (The `/sbin/init.d/gwlmcmsd` and `/sbin/init.d/gwlagent` scripts use these variables.)
- Using the `--enable_start_on_boot` and `--disable_start_on_boot` options to `gwlmcmsd` and `gwlagent`.

On Windows, to control whether the gWLM CMS daemon (`gwlmcmsd`) starts at boot: Right-click the **My Computer** icon, click **Manage**, double-click **Services & Applications**, double-click **Services**, right-click **HP Global Workload Manager Central Management Server**, select **Properties**, and change the **Startup type** to the desired setting.



NOTE: HP recommends starting these daemons at boot only when used in a secure operating environment.

Setting `GWLM_CMS_START` to 0 prevents automatic use at boot of HP Virtualization Manager and HP Capacity Advisor.

The `gwlmCtl` file is shown below:

```
#####
# (C) Copyright 2004-2008 Hewlett-Packard Development Company, L.P.
#
# gWLM Configuration File
#
# $Revision: 1.2 $
# $Date: 2008/02/05 21:30:52 $
#####

# Set GWLM_CMS_START to 1 to have the init process start the gWLM CMS
# daemon. (HP recommends setting this variable to 1 only when used in a
# secure operating environment.)
#
# NOTE: GWLM_CMS_START=0 prevents automatic use at boot of
#       HP Virtualization Manager and
#       HP Capacity Advisor.
GWLM_CMS_START=0

# Set GWLM_AGENT_START to 1 to have the init process start the gWLM agent
# daemon. (HP recommends setting this variable to 1 only when used in a
# secure operating environment.)
GWLM_AGENT_START=0

# Set GWLM_HOME to the location where gWLM is installed.
```

```
# Default is /opt/gwlm.  
GWLM_HOME=/opt/gwlm
```

Automatic Restart of gWLM's Managed Nodes in SRDs (High Availability)

Whenever a managed node boots, the node's gWLM agent attempts to automatically rejoin the node in its SRD, providing high availability. The only configuration steps you need to perform for this behavior to happen are:

1. Ensure the `/etc/rc.config.d/gwlmCtl` file on each managed node has `GWLM_AGENT_START` set to 1. You can run the following command on each system where `gwlmagent` is running to make this change for you:

```
# /opt/gwlm/bin/gwlmagent --enable_start_on_boot
```

In the same file, you also need `GWLM_CMS_START=1` on the system where `gwlmcmd` is running. However, when you ran `vseinitconfig` during installation, this change was automatically made.

2. (Optional) Edit the property

`com.hp.gwlm.node.HA.minimumTimeout`

in the file `/etc/opt/gwlm/conf/gwlmagent.properties` to set the minimum number of seconds that must pass before a managed node considers itself separated from its SRD. Set this property to ensure that minor network problems do not cause a managed node to prematurely consider itself separated.

gWLM uses this value only if it is larger than 10 multiplied by gWLM's allocation interval. For example, with an allocation interval of 15 seconds, a node can go 2.5 minutes without communicating with its SRD before the node's gWLM agent attempts to re-connect with the SRD.

This feature works best when one managed node is lost at a time or all managed nodes are lost.



NOTE: If a vpar is borrowing cores from other vpars when it loses contact with its SRD, those borrowed cores may be separated from the SRD. If the vpar may be down for an extended time, check that the SRD has reformed without that vpar and that it has enough cores to meet its commitments. If not, try using `vparmodify` to reclaim some of the cores. (With the vpar down, you will not be able to modify it locally, and only some versions of HP-UX Virtual Partitions allow you to easily modify a remote vpar.)

Similarly, if an npar has several active cores (due to Instant Capacity) when it loses contact with its SRD, you may have to manually size the npar to reclaim those cores for npars still in the SRD. For more information, see the Instant Capacity documentation.

How the Automatic Restart Works

When a managed node boots, the gWLM agent (`gwlmagent`) starts automatically if `GWLM_AGENT_START` is set to 1 in the file `/etc/rc.config.d/gwlmCtl`. The agent then checks the file `/etc/opt/gwlm/deployed.config` to determine its CMS. Next, it attempts to contact the CMS to have the CMS re-deploy its view of the SRD. If the CMS cannot be contacted, the SRD in the `deployed.config` file is deployed as long as all nodes agree.

In general, when an SRD is disrupted by a node's going down, by a CMS's going down, or by network communications issues, gWLM attempts to reform the SRD. gWLM maintains the concept of a cluster for the nodes in an SRD. In a cluster, one node is a master and the other nodes are nonmasters. If the master node loses contact with the rest of the SRD, the rest of the SRD can continue without it, as a partial cluster, by unanimously agreeing on a new master. If a nonmaster loses communication with the rest of the SRD, the resulting partial cluster continues operation without the lost node. The master simply omits the missing node until it becomes available again.

You can use the `gwlmsstatus` command to monitor availability. It can tell you whether any hosts are unable to rejoin a node's SRD as well as whether hosts in the SRD are nonresponsive. For more information, see `gwlmsstatus(1M)`.



NOTE: Attempts to reform SRDs may time out, leaving no SRD deployed and consequently no management of resource allocations. If this occurs, see the *VSE Management Software Release Notes* and follow the actions suggested in the section titled “Data Missing in Real-time Monitoring.”

Related Events

You can configure the following HP SIM events regarding this automatic restart feature:

- Node Failed to Rejoin SRD on Start-up
- SRD Reformed with Partial Set of Nodes
- SRD Communication Issue

For information on enabling and viewing these events, refer to gWLM’s “Configure Events” menu.

You can then view these events using the Event Lists item in the left pane of HP SIM.

The following sections explain how to handle some of the events.

“Node Failed to Rejoin SRD on Start-up” Event

If you see the event “Node Failed to Rejoin SRD on Start-up”:

1. Restart the `gwlmsagent` on each managed node in the affected SRD:

```
# /opt/gwlm/bin/gwlmsagent --restart
```
2. Verify the agent rejoined the SRD by monitoring the Shared Resource Domain View in HP SIM or by using the `gwlm monitor` command.
3. If the problem persists, check the files `/var/opt/gwlm/gwlmsagent.log.0` and `/var/opt/gwlm/gwlmsmcsd.log.0` for additional diagnostic messages.

“SRD Communication Issue” and “SRD Reformed with Partial Set of Nodes” Events



NOTE: Reforming with a partial set of nodes requires a minimum of three managed nodes in the SRD.

NOTE: “SRD Communication Issue” events are not enabled by default. To see these events, configure your events in HP SIM through the VSE Management menu bar using **Tools→Global Workload Manager→Events**.

If you have an SRD containing n nodes and you get $n - 1$ of the “SRD Communication Issue” events but no “SRD Reformed with Partial Set of Nodes” events within 5 minutes (assuming an allocation interval of 15 seconds) of the first “SRD Communication Issue” event, you may need to restart the `gwlmsagent` on each managed node in the affected SRD:

```
# /opt/gwlm/bin/gwlmsagent --restart
```

Manually Clearing an SRD

If gWLM is unable to reform an SRD, you can manually clear the SRD, as described in the following section.

Clearing an SRD of A.02.50.00.04 (or later) agents

The following command is an advanced command for clearing an SRD. The recommended method for typically removing a host from management is by using the `gwlms undeploy` command.

Starting with A.02.50.00.04 agents, you can manually clear an SRD with the following command:

```
# gwlms reset --host=host
```

where *host* specifies the host with the SRD to be cleared.

If this command does not work, use the procedure given in the following section.

Clearing an SRD of agents of any version

The procedure in this section clears an SRD regardless of the version of the agents in the SRD.

The `gwlms` command is added to the path during installation. On HP-UX systems, the command is in `/opt/gwlms/bin/`. On Microsoft Windows systems, the command is in `C:\Program Files\HP\Virtual Server Environment\bin\gwlms\` by default. However, a different path may have been selected at installation.



NOTE: You must be logged in as root on HP-UX or into an account that is a member of the Administrators group on Windows to run the commands below.

1. Delete the `deployed.config` file on each managed node:

```
# rm -f /etc/opt/gwlms/deployed.config
```
 2. Force an undeploy of the SRD (named *SRD* below) to ensure the CMS and the managed nodes agree on the SRD's state. Run the following command on the CMS:

```
# gwlms undeploy --srd=SRD --force
```
 3. Restart the `gwlmsagent` daemon on each managed node:

```
# /opt/gwlms/bin/gwlmsagent --restart
```
-



NOTE: If the gWLM CMS and agent disagree about whether an SRD is deployed or undeployed, you can use the `--force` option with the `gwlms deploy` or `gwlms undeploy` commands.


Nesting Partitions

gWLM allows you to form SRDs consisting of various compartment types. This ability provides flexibility in dividing your complex. For example, you can divide your complex as shown in Figure 5-1. The complex has four npar, two of which are divided into vpars. One npar is hosting virtual machines, and the fourth npar is not divided. gWLM allows you to create an SRD containing the two virtual machine guests, the two vpars from npar 2, the two vpars from npar 3, and npar 4. The workloads in any of these compartments can then borrow resources from any of the other compartments in the SRD. If TiCAP is available on the complex, gWLM can migrate the usage rights to where they are needed.



NOTE: No more than one deployed SRD per complex should have nested partitions.

Figure 5-1 Nested Partitions



For more information on nesting partitions, see the online help or to *gwlm(1M)*.

Changing the gWLM Resource Allocation Interval

The frequency of gWLM's changes in the CPU resource allocations is an attribute of the SRDs. Once you create an SRD, you can change how often gWLM adjusts the CPU resource allocations of the workloads in that SRD using either of the methods discussed in the following sections.

Changing the Interval in HP SIM



NOTE: If you are managing virtual machines, specify a resource allocation interval that is a multiple of the `vm_essaygt` interval, which defaults to 10 seconds. If you need a resource allocation interval of fewer than 10 seconds, set the interval for the VM agents (using `vm_essaygt -i`) to the same value as your resource allocation interval.

Using HP SIM, you can set the interval in two places:

- When creating an SRD
From the HP SIM menu bar, select:
Tools→**VSE Management**
Then, click the **Shared Resource Domain** tab. From the VSE Management menu bar, select:
Create→**Shared Resource Domain**
- When editing an SRD
From the HP SIM menu bar, select:
Tools→**VSE Management**
Then, click the **Shared Resource Domain** tab. From the VSE Management menu bar, select:
Modify→**Shared Resource Domain**

Changing the Interval on the Command Line

Use the `gwlmm` command and a text editor to change the interval on the command line:

1. Use `gwlmm export` to get a copy of the SRD's XML definition from the gWLM configuration repository.
2. Edit the "interval" attribute, which is in seconds.



NOTE: If you are managing virtual machines, set the interval attribute value to a multiple of the `vm_essaygt` interval, which defaults to 10 seconds. If you need a gWLM interval attribute value of fewer than 10 seconds, set the interval for the VM agents (using `vm_essaygt -i`) to the same value.

3. Use `gwlmm import --clobber` to place the updated definition in the gWLM configuration repository.

If an SRD of the same name is already deployed, the import operation puts the new interval into effect. Otherwise, the interval is used the next time you deploy the SRD you just defined and imported.

Using gWLM with Hyper-Threading

Hyper-Threading, available starting with HP-UX 11i v3 (B.11.31), enables you to use multiple execution threads per core. (Cores were formerly known as CPUs.) Each execution thread is a logical CPU.

Utilization is based on the usage of a core—not a logical CPU. For example, assume a four-core system has Hyper-Threading enabled so that it has eight logical CPUs. If the system has four processes, each consuming an entire logical CPU, the reported utilization depends on where

those processes are. If the processes are on only two cores, the utilization is 50% (2/4). With the processes distributed across all four cores though, each process can consume an entire core, resulting in a utilization of 100%.

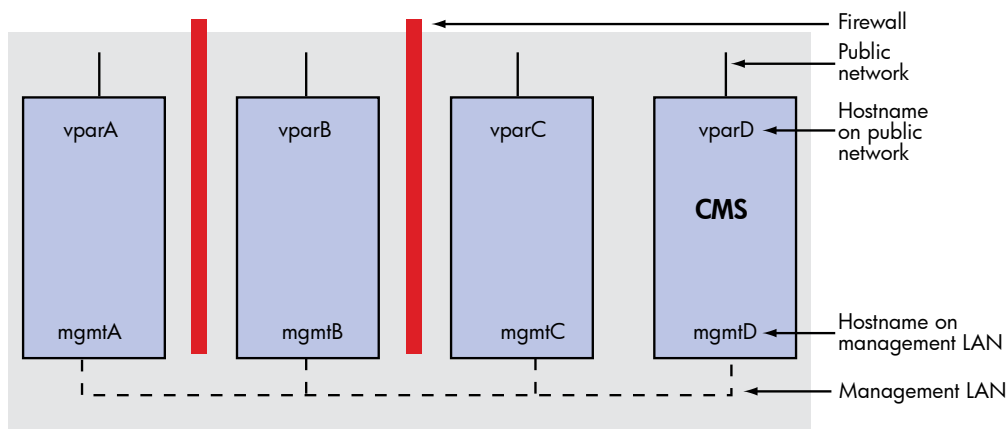
When fss groups are being used, gWLM disables Hyper-Threading for the default pset, where fss groups are created, to optimize workload performance.

When an SRD is undeployed, gwlmagent restores the lcpu_attr tunable to the value it had when the system booted. This value can be different from the value in effect before deployment if kctune was used without a reboot.

Using gWLM with Hosts on Multiple LANs

In data center operations and other cases, you may have hosts on multiple LANs or separated by firewalls. These hosts may be addressable on a public network or an internal network. However, they cannot communicate with each other because of the separation by the LANs or firewalls. Figure 5-2 shows a possible scenario.

Figure 5-2 Using gWLM with Hosts Separated by Firewalls



On the public network, the hosts can be accessed through their hostnames: vparA, vparB, vparC, and vparD. However, they cannot access each other through those names.

Although gWLM may be able to discover the hosts and you can even configure an SRD including the hosts, when you attempt to deploy the SRD, gWLM will eventually time out and display a blank screen. No error message is displayed. However, there will be events from each managed node similar to the following event:

```
gWLM Agent MySystem.MyDomain.com
Information Unable to manage the following hosts:
Associated Exception Unable to manage the following hosts: MySystem.MyDomain.com: The gWLM agent process on
the host is not running -- start the agent and retry.
```

A CMS can only manage hosts on the same LAN as the CMS itself. Thus, if you set up a separate LAN (sometimes referred to as a management LAN or a backup LAN) that includes the CMS and all the hosts to be managed, you can manage these hosts in a single SRD. Figure 5-2 shows a management LAN in which the hosts are known as mgmtA, mgmtB, mgmtC, and mgmtD. With this management LAN, gWLM can manage the hosts in a single SRD. Complete the following procedure to set up gWLM to manage such hosts in an SRD:

1. For each host in the management LAN that you want to manage in an SRD:
 - a. Edit the `/etc/opt/gwlm/conf/gwlmagent.properties` file to include the following property:


```
com.hp.gwlm.security.virtualLocalHostName=hostnameOnLAN
```

 For example, with the host `mgmtA`, its property would be:


```
com.hp.gwlm.security.virtualLocalHostName=mgmtA
```
 - b. Restart `gwlmagent` on the host:


```
# /opt/gwlm/bin/gwlmagent --restart
```
2. The CMS must also be in the management LAN. If the primary hostname for the CMS is not the name it uses in the management LAN:
 - a. Edit the `gwlmagent.properties` file on the CMS to include the property:


```
com.hp.gwlm.security.virtualLocalHostName=hostnameOnLAN
```

 On HP-UX, the `gwlmagent.properties` file is in `/etc/opt/gwlm/conf/`. On Windows, it is in `C:\Program Files\HP\Virtual Server Environment\conf\`. (The given Windows path is the default; however, a different path may have been selected at installation.)
 - b. Restart HP SIM and `gwlmcmsd`.
3. On each host in the SRD (CMS and managed nodes), ping every other host in the SRD—using the hostname you intend to have gWLM discover (its `hostnameOnLAN`)—to verify communication is possible.
4. Discover (or re-discover) the hosts using the gWLM Manage Systems and Workloads wizard.

Creating “Golden Images”

If you create golden images of a system’s applications and operating system to store in an IgniteUX server for distribution to other systems, here are tips for creating images that include gWLM.

When creating an image of a managed node:

- Ensure the gWLM agent is set to start automatically.
In the file `/etc/rc.config.d/gwlmCtl`, set the variable `GWLM_AGENT_START` equal to 1.
- Ensure no deployed SRD includes the node. (If the file `/etc/opt/gwlm/deployed.config` exists, an SRD is deployed on the node. Temporarily undeploy the SRD or unmanage the node.)

If a node based on a virtual partition is part of a deployed SRD and you make a golden image of the virtual partition, once you push that golden image out, the gWLM agent on the new system will still consider itself part of the SRD. The agent will then try to rendezvous with that SRD’s other managed nodes. You can reset the agent by deleting the `deployed.config` file (mentioned previously), then stopping and restarting `gwlmagent`.

A Compatibility with Agents

The gWLM A.04.00.07 CMS runs on HP-UX 11i v1 (B.11.11), HP-UX 11i v2 (B.11.23), HP-UX 11i v3 (B.11.31), and Microsoft Windows systems. It works with the following versions of the agents:

- gWLM A.02.00.00.07: HP-UX 11i v1, HP-UX 11i v2
- gWLM A.02.50.00.04: HP-UX 11i v1, HP-UX 11i v2
- gWLM A.03.00.00.05: HP-UX 11i v1, HP-UX 11i v2, HP-UX 11i v3
- gWLM A.03.00.01.05: HP-UX 11i v1, HP-UX 11i v2, HP-UX 11i v3
- gWLM A.04.00.07: HP-UX 11i v1, HP-UX 11i v2, HP-UX 11i v3



NOTE: While gWLM does support the earlier agent versions listed above, you are encouraged to upgrade to the latest version.

Use the `gwlmsatus` command on a managed node to determine the version of the gWLM agent software on that node.

Also, all gWLM agents within an SRD must be the same version: A gWLM CMS can manage an SRD that uses only A.04.00.07 agents, only A.03.00.01.05 agents, only A.03.00.00.05 agents, and so forth.

The new features introduced with a new agent version are typically not backward-compatible. For information on the key features supported with each supported agent version, see the section "Working with earlier gWLM versions" in `gwlxml(4)`.