HP Global Workload Manager Agent for OpenVMS Version 4.1 Installation Guide and User Guide

Table of Contents

List of Figures

List of Tables

About This Document

This document presents an overview of the techniques and tools available for using the HP Global Workload Manager (gWLM) for HP Integrity servers and HP 9000 servers. It exposes you to the essentials and allows you to get started quickly using gWLM. This document describes how to install and configure gWLM Version 4.1 on an OpenVMS operating system. It also discusses the prerequisites for installing gWLM. It also describes how to secure gWLM communications and manage workloads.

Intended Audience

This document is intended for 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 OpenVMS system administration concepts and procedures is assumed. A knowledge of HP-UX system administration concepts is also useful.

Document Organization

The document is organized as follows:

Chapter 1   "Overview" Provides an overview of gWLM and its benefits. It also describes the concepts and terms associated with gWLM and explains the gWLM management model.

Chapter 2   "Installing and Uninstalling gWLM" Describes how to install and uninstall gWLM.

Chapter 3   "Configuring gWLM to Manage Workloads" Describes how to configure gWLM to effectively manage the resources of your workload.

Chapter 4   "Monitoring Workloads and gWLM" Describes how to monitor workloads and gWLM.

Chapter 5   "Security" Highlights several security items that you should be aware of while using gWLM.

Chapter 6   "Additional Configuration and Administration Tasks" Provides information on various configuration and administration tasks.

Appendix   "Compatibility with Agents" Provides information on the compatibility of gWLM with agents.

Typographic Conventions

Table 1 lists the typographic conventions used in the document.

Table 1 Typographic Conventions

Convention	Description
. . .	A horizontal ellipsis in a figure or examples indicates the following possibilities: • Additional optional arguments in a statement have been omitted. • The preceding item or items can be repeated one or more times. • Additional parameters, values, or other information can be entered.
. . .	A vertical ellipsis indicates the omission of items from a code example or command format; the items are omitted because they are not important to the topic being described.
( )	In command format descriptions, parentheses indicate that you must enclose choices in parentheses if you specify more than one. In installation or upgrade examples, parentheses indicate the possible answers to a prompt, such as: Is this correct? (Y/N) [Y].

Table 1 Typographic Conventions *(continued)*

Convention	Description
[ ]	In command format descriptions, brackets indicate optional choices. You can choose one or more items or no items. Do not type the brackets on the command line. However, you must include the brackets in the syntax for OpenVMS directory specifications and for a substring specification in an assignment statement.
{}	In command format descriptions, braces indicate required choices; you must choose at least one of the items listed. Do not type the braces on the command line.
Example	This typeface indicates code examples, command examples, and interactive screen displays. In text, this type also identifies website addresses, OpenVMS command and pathnames, PC-based commands and folders, and certain elements of the C programming language.
italic type	Italic type indicates important information, complete titles of manuals or variables. Variables include information that varies in system output (for example, Internal error number), in command lines (/PRODUCER=name), and in command parameters in text (where dd represents the predefined code for the device type).
UPPERCASE TYPE	Uppercase indicates the name of a command, routine, file, file protection code, or the abbreviation of a system privilege.
-	A hyphen at the end of a command format description, command line, or code line indicates that the command or statement continues on the following line.
WARNING	A warning calls attention to important information that if not understood or followed will result in personal injury or nonrecoverable system problems.
CAUTION	A caution calls attention to important information that if not understood or followed will result in data loss, data corruption, or damage to hardware or software.
IMPORTANT	This alert provides essential information to explain a concept or to complete a task.
NOTE	A note contains additional information to emphasize or supplement important points of the main text.

Related Information

The latest versions of manuals and white papers for HP Insight Dynamics — VSE, the VSE Management Software, and related products can be downloaded from the HP Web:

- Documentation for HP Insight Dynamics — VSE and related HP ProLiant software: http://www.hp.com/go/insightdynamics/docs
- Documentation for VSE Management Software and related HP Integrity software: http://docs.hp.com/en/vse.html

    For more information about HP Insight Dynamics — VSE, the VSE Management Software, and VSE-related products and solutions, visit the following HP websites:
    — HP Virtual Server Environment: http://www.hp.com/go/vse
    — HP Insight Dynamics — VSE: http://www.hp.com/go/insightdynamics

- *HP Instant Capacity User's Guide for Versions 8.x*
- *HP OpenVMS V8.3-1H1 and 8.3 Systems Documentation*

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 provide 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 (LTU).)

4. On each managed node, start the gWLM agent daemon `gwlmagent`.

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

Finding More gWLM Information

Table 2 indicates where you can find additional information about gWLM.

Table 2 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.	*vseinitconfig*(1M)
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 *HP Global Workload Manager User's Guide* (this document) (http://docs.hp.com/en/vse.html[2])
Learn about gWLM concepts.	Global Workload Manager topic in online help[1] or *HP Global Workload Manager User's Guide* (this document) (http://docs.hp.com/en/vse.html[2])
Learn gWLM terms.	"Concepts and Terms in gWLM" (page 11) 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.	*HP Global Workload Manager User's Guide* (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.	*gwlm*(1M)
Learn about gWLM daemons and services	*gwlmcmsd*(1M)

Table 2 Where to Find Additional Information *(continued)*

To...	See...
Learn about using secure communications with gWLM.	Securing gWLM Communications topic in online help[1] or *gwlmsslconfig*(1M)
Learn how to update metrics when using custom policies.	*gwlmsend*(1M)
Learn how to manually place processes in workloads based on psets or FSS groups.	*gwlmplace*(1M)
Learn about using gWLM with HP Serviceguard.	The "Technical Documentation website for HP Virtual Server Environment (VSE)": http://docs.hp.com/en/vse.html

Learn more about nPars, vpars, virtual machines, and psets.	• HP Virtual Server Environment website:
	  http://www.hp.com/go/vse
	
	• The "Technical Documentation website for HP Virtual Server Environment (VSE)" website:
	  http://docs.hp.com/en/vse.html

Learn about other gWLM commands.	The SEE ALSO section of gwlm(5) for a list of all gWLM manpages

1   To access the online help in HP SIM, click **Tools→VSE Management**, then click the **Shared Resource Domain** tab, and then click question mark (?) on the top right corner.

2   This documentation is in English. For documentation in other languages, see http://docs.hp.com.

HP Encourages Your Comments

HP encourages your comments and suggestions on this document. Please send comments to:

openvmsdoc@hp.com

1 Overview

HP Global Workload Manager (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 can configure gWLM and monitor your workloads from the system where the CMS software is installed. Also, you can use an agent software on the systems where you have workloads that you want gWLM to manage.

This chapter addresses the following topics:

- "Benefits"
- "Concepts and Terms in gWLM"
- "The gWLM Management Model"

Benefits

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.

Concepts and Terms in gWLM

Following is a list of concepts and terms you must be aware of when using gWLM:

Workload
The collection of processes executing within a single compartment. The compartment can be an nPartition (npar), 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, psets, and FSS groups, see "The gWLM Management Model" (page 13).) On OpenVMS, use UNIX-style pathnames to create an application workload definition For example:

`/sys$common/java$142/bin/java$java.exe`

`/sys$system/vms*.exe`

`/RED$DKA0/SYS0/SYSCOMMON/JAVA$142/BIN/JAVA$JAVA.EXE`

Compartment	An npar, a virtual machine, a pset, or an FSS group with its resource allocation being managed by gWLM.
	Multiple compartments are grouped to form a shared resource domain (SRD). All the compartments 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.
Shared Resource Domain (SRD)	A collection of compartments that can share system resources. The compartments can be nPars, 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 13) are met.
	gWLM allows you to nest compartments. gWLM then manages resources for the various levels of compartments.
Policy	A collection of settings that instruct gWLM on 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 required), as well as how much of those resources the workload can lend to other workloads.
	A single policy can be associated with, or applied to, multiple workloads.
	For more information on policies, see "Policy Types" (page 21).
Mode	Two modes are available: advisory and managed. The advisory mode allows you to see what CPU resource requests gWLM would make for a workload—without actually affecting resource allocation.
	The advisory mode is not available for SRDs containing psets, or FSS groups due to the nature of these compartments.
	Use this mode when creating and fine-tuning your policies. Once you are comfortable with your policies, use the managed mode to have gWLM automatically adjust the resource allocations for your defined workloads.
	You can only set the mode at the SRD level. All workloads within an SRD operate in the same mode, either advisory or managed.
Deploy	Enables gWLM to control an SRD.
	Deploying an SRD in managed mode enables gWLM to control the 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.)

When deploying an SRD in advisory mode, gWLM simply reports what the allocation would be—without actually affecting resource allocations on a system.

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

Undeploy　　　　　　　　　　Disables 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 an SRD as required—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.

- OpenVMS 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 OpenVMS operating system.

  Using the HP Instant Capacity product, gWLM simulates the movement of CPU resources among nPars by turning off an active core in one nPar and 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.)

- 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.

  gWLM on OpenVMS uses process and processor capabilities to allocate resources among psets.

- 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 OpenVMS. 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.

  gWLM on OpenVMS uses Class Scheduler to allocate resources among FSS.

For more information on these partitions, visit:

- HP Virtual Server Environment website:

  http://www.hp.com/go/vse

- The "Technical Documentation website for HP Virtual Server Environment (VSE)" website:

  http://docs.hp.com/en/vse.html

- The "Global Workload Manager" topic and the glossary in the online help for gWLM, available in gWLM's graphical interface in HP SIM.

gWLM manages resources based on the following model:

1. You can 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 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 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 specified 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 the types of workloads to combine for optimal resource utilization, see the online help topic "Getting the Most Out of gWLM."

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
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, see 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, select:

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

  (The HP SIM menu bar and VSE Management menu bar are discussed in the section .)

- gwlm command

  An HP-UX command-line interface, described in *gwlm*(1M) .

  Other components of the command-line interface are: *vseinitconfig*(1M), *gwlmcmsd*(1M), *gwlmagent*(1M), *gwlmreport*(1M), *gwlmplace*(1M), *gwlmsend*(1M), *gwlmsslconfig*(1M), *gwlmstatus*(1M), and *gwlmxml*(4).

2 Installing and Uninstalling gWLM

This chapter provides information on how to install and use the gWLM agent for OpenVMS. It addresses the following topics:
- "Prerequisites"
- "Installing gWLM"
- "Uninstalling gWLM"
- "Compatibility with Other System Management Methods"

Prerequisites

Before installing gWLM, you must have the following software/applications installed on your system:

a. The HP VSE Management Software CMS to set up and monitor gWLM.

b. To complete your installation,
- HP Systems Insight Manager 6.0: http://h18013.www1.hp.com/products/servers/management/hpsim/
- HP Global Workload Manager CMS Plugin: http://h18013.www1.hp.com/cpq-products/servers/management/hpsim/plugin-apps.html

  The following steps provide 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 (LTU).)

  4. On each managed node, start the gWLM agent daemon gwlmagent.

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

- HP TCP/IP services for OpenVMS Integrity servers Version 5.6 or later.
- Coordinated Universal Time (UTC):

  Ensure that the UTC is correctly set up on the managed systems. Failure to do so might result in errors in time stamp values reported in the user interface.

c. OpenVMS installation requirements:

Table 2-1 lists the installation requirements of the gWLM agent on OpenVMS Integrity servers.

Table 2-1 Installation Requirements

System	Functionality	Dependent Kits and Patches
Cell-based	Pset, FSS, nPar, iCAP, and TiCAP	Latest WBEMCIM and iCAP patches.
Non-Cell-based	Pset and FSS	None

Before installing gWLM, perform the following steps:

1) Install the patches listed in the Table 2-1 in the same order using the `/SAVE_RECOVERY_DATA` qualifier.
2) Reboot the system after installing the patches and ensure that WBEMCIM and WBEMPROVIDERS are configured properly and running before configuring the iCAP software.

For more information, see the *HP Instant Capacity User's Guide for Versions 8.x*. These documents provide information regarding the iCAP requirements on OpenVMS.

Installing gWLM

This section provides information on how to install gWLM on your operating system. You can install the gWLM agent for OpenVMS on the systems whose workloads you want to manage.

NOTE: Before installing gWLM, you must install a `PAK` file that is appropriate to your hardware platform.

To install gWLM, follow these steps:

1. Install gWLM:

   ```
   $ product install gwlm/source=<kit-directory>
   ```

2. Start the gWLM agent:

   ```
   $ @sys$startup:gwlm$startup
   ```

   To stop the gWLM agent:

   ```
   $ @sys$startup:gwlm$shutdown
   ```

The gWLM agent gets installed on OpenVMS in the `sys$sysroot:[gwlm]` directory. Table 2-2 displays the mapping of UNIX pathnames referenced in the SIM interface with the OpenVMS pathnames:

Table 2-2 File Locations

UNIX Pathname	OpenVMS Pathname
/etc/opt/gwlm	sys$sysroot:[gwlm]
/var/opt/gwlm	sys$sysroot:[gwlm.log]
/etc/opt/gwlm/conf	sys$sysroot:[gwlm.conf]
/opt/gwlm/bin/	sys$common:[gwlm.bin]

You can perform most of the operations using either the HP SIM interface or the gWLM command–line interface on the CMS. The following agent commands are available on a system on which the OpenVMS agent is installed:

- `gwlmsslconfig`
- `gwlmimportkey`
- `gwlmexportkey`
- `gwlmplace`
- `gwlmsend`
- `gwlmstatus`

These commands are implemented as foreign commands. To use these commands, you must first execute the following command:

```
$ @sys$sysroot:[gwlm.bin]gwlm$env get_commands
```

To access help for these commands, enter the following command:

```
$ help gwlm
```

To use agent commands in `sys$sysroot:[gwlm]` directory, users must have read and write access to the files.

Managing Node Configuration

If there is a variation in the host name configuration between the CMS and the managed hosts, an error is encountered during the discovery phase. See the following example:

```
# gwlm discover test.zko.dec.com
```

```
The IP address 16.32.16.151 resolves to two different fully qualified
domain names.
```

```
The fully qualified domain name on the managed node is test; the CMS
has the managed node's fully qualified domain name as test.zko.dec.com.
```

```
Please correct and retry. Please re-enter the systems.
```

The following example deals with the problem on the managed node by providing the managed node with a fully qualified domain name definition:

```
TCPIP> show host test
LOCAL database
Host address Host name
16.32.16.151 test, TEST

TCPIP> show host test
LOCAL database
Host address Host name
16.32.16.151 test, TEST

TCPIP> set nohost "test"
TCPIP> set host "test.zko.dec.com"/address=16.32.16.151/alias=("test",test)
TCPIP> show host test
LOCAL database
Host address Host name
16.32.16.151 test, TEST

TCPIP> set nohost "test"
TCPIP> set host "test.zko.dec.com"/address=16.32.16.151/alias=("test",test)
TCPIP> show host test
LOCAL database
Host address Host name
16.32.16.151 test.zko.dec.com, test, TEST
```

Defining Workloads

The following are the guidelines for composing application workload definitions:

- You can compose an application workload definition using UNIX-style application pathnames. For example:

  `/sys$common/java$142/bin/java$java.exe`

  `/sys$system/vms*.exe`

  `/RED$DKA0/SYS0/SYSCOMMON/JAVA$142/BIN/JAVA$JAVA.EXE`

- You can use the wildcard characters, ''%'' and ''*''. ''Application Alternate Names'' are not necessary on OpenVMS; do not specify them.
- You can use OpenVMS-style pathnames for individual, non-wildcarded executable. For example:

  `RED$DKA0:[SYS0.SYSCOMMON.JAVA$142.BIN]JAVA$JAVA.EXE`

  `SYS$SYSTEM:VMSHELP.EXE`

  `SYS$COMMON:[SYSEXE]VMSHELP.EXE`

  In addition, the pathname must start with a ''/'' character; for example:

  `/RED$DKA0:[SYS0.SYSCOMMON.JAVA$142.BIN]JAVA$JAVA.EXE`

```
/SYS$SYSTEM:VMSHELP.EXE

/SYS$COMMON:[SYSEXE]VMSHELP.EXE
```

- You can use OpenVMS user account names. For example:
  - system
  - guest

Uninstalling gWLM

You can uninstall gWLM using the following command:

```
$ @sys$startup:gwlm$shutdown

$ product remove gwlm
```

Compatibility with Other System Management Methods

The gWLM agent uses of the following OpenVMS technologies:

- Class scheduler
- Process and processor capabilities

> **NOTE:** If another software is attempting resource management with any of these technologies, gWLM may not be able to place all the system processes in the workloads that you define by using the gWLM user interface.

Limitations

The following are the limitations of gWLM on OpenVMS:

- On OpenVMS managed nodes, **Restart Agents** option does not function. But, you can manually stop and start gWLM agent on the managed nodes by using the commands:

  ```
  @sys$startup:gwlm$shutdown.com

  @sys$startup:gwlm$startup.com
  ```

- gWLM agent does not support GiCAP on OpenVMS.
- On OpenVMS, gWLM agent does not support Conditional Policy, Process Map and GiCAP.

3 Configuring gWLM to Manage Workloads

This chapter provides information on how to configure gWLM to effectively manage the resources for your workloads. It addresses the following topics:

- "Policy Types"
- "Choosing a Policy Type"
- "Seeing How gWLM will Perform Without Affecting the System"
- "Getting Started with gWLM"
- "Seeing gWLM in Action"
- "Common Uses for gWLM"
- "Common Configuration Tasks"

Policy Types

You can define several types of policies to instruct gWLM on how to manage the resources for your workloads. The policy types include the following:

Fixed
: Allocates a fixed (constant) amount of CPU resources to a workload's compartment.

  gWLM satisfies these policies before attempting to satisfy any other type of policies.

Utilization
: 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 that 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%, thereby meeting the target.

  With a utilization policy, you can 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.

OwnBorrow
: Allows you to set the following values:
  - 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).

  The compartment of a workload with an OwnBorrow policy is allocated the owned CPU resources when required. The minimum and maximum sizes allow you to specify how much the workload can lend (when resources are not required) or borrow (when additional resources are required and available). If a compartment has lent out cores and that compartment's workload becomes busy, the compartment re-acquires those lent-out cores.

Custom
: Available for advanced users. For information on custom policies, see the online help or *gwlmxml*(4).

Conditional   Specifies the existing policy to use when a time-based condition, a file-based condition, or a Serviceguard condition is met.

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

Choosing a Policy Type

How do you decide which policy type to use? Table 3-1 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 3-1 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 *gwlmreport*(1M).
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 that the associated workloads get their owned resources whenever needed. Thus, an OwnBorrow policy provides

greater flexibility in attempting to allocate a certain amount of resources to a workload 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. 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 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.

The advisory mode is not available for SRDs containing psets, or FSS groups due to the nature of these compartments.

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

For information on changing modes, see "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 *gwlm*(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, decide:

- Which systems you want to manage with gWLM.
- Whether you want to manage your workloads by migrating CPU resources among nPars, 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 13).)

Tabs and Menus

The controls shown in Figure 3-1 appear at the top of the Global Workload Manager screen.

Figure 3-1 Top Controls of the Global Workload Manager Screen
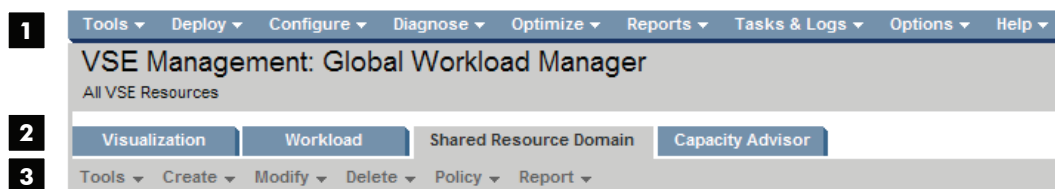
1   The SIM menu bar
2   The Virtualization Manager tabs
3   The VSE Management menu bar

These menu bars are referenced later in this guide.

Using the Wizard

To start the wizard:

> **NOTE:** You must be logged in as `root` on the systems where you run the `mxstart`, `gwlmcmsd`, 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:

    On HP-UX,

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

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

    On OpenVMS,

    ```
    $ @sys$startup:gwlm$startup.com
    ```

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 the SRD properties.

    Properties include the SRD name, mode, use of Temporary Instant Capacity (TiCAP) (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 nPars. 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 must place processes in the desired pset or FSS group. (Place processes by modifying the workload definition or by using the `gwlmplace` command.)

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 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 23)
- When adding a workload to an SRD, as described in "Adding a New Compartment 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 28)

Resizing a Workload's nPar, pset, or FSS Group

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 can 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 23)
- When adding a workload to an SRD, as described in "Adding a New Compartment 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 28)

gWLM's utilization policies also allow resizing.

Common Configuration Tasks

This section discusses various configuration tasks:

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 `gwlmagent` 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/gwlmagent**

   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

The 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 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 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 that HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 26).

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 on 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 that HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 26).
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 by selecting a policy type and specifying the required values and optional values as desired.
6. Click **OK**.

Editing a Policy

A policy instructs gWLM on 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 that HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`), and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 26).

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 that HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 26).
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** drop-down 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 to an SRD

You can add a new compartment to an SRD if you

- Have added an nPar to your system and want to add it to an SRD
- 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 nPar members. Step 2 allows you to create psets or FSS groups.

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 that HP SIM, the gWLM CMS daemon or service (`gwlmcmsd`), and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Setting up gWLM (Initial Setup Steps)" (page 26).
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 the nPars present in the SRD.

   When gWLM stops managing an SRD, it leaves compartments based on nPars as they were in the last allocation interval. Associating fixed policies allows you to set the sizes exactly to what you want. (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, see "Changing Which Policy is Associated with a Workload" (page 28).
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.

4 Monitoring Workloads and gWLM

This chapter describes how to monitor workloads and gWLM. It addresses the following topics:
- "Monitoring Workloads"
- "Viewing gWLM Reports in Monitor-only Mode"
- "Monitoring gWLM from the Command-line"
- "Message Logs"
- "Viewing HP Systems Insight Manager Events"

Monitoring Workloads

There are several methods for monitoring workloads, as described in the following sections.

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, see 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 permission to change gWLM configurations. To set up a user with monitor-only privileges, see 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 during 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 listed in the following section.

- `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 4-1 gWLM Log Files

Log for	Location
`gwlmcmsd` daemon or service	HP-UX: `/var/opt/gwlm/gwlmcmsd.log.0` Windows: `C:\Program Files\HP\Virtual Server Environment\logs\gwlmcmsd.log.0` OpenVMS: Not Applicable
`gwlmagent` daemon	HP-UX: `/var/opt/gwlm/gwlmagent.log.0` Windows: Not applicable OpenVMS: `sys$specific:[gwlm.log]gwlmagent_log.0`.
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 during 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 MB 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 the `/etc/opt/gwlm/conf/gwlmagent.properties` file. For all the other log files, change the values in `/etc/opt/gwlm/conf/gwlmcms.properties` file on HP-UX and in the `C:\Program Files\HP\Virtual Server Environment\conf\gwlmcms.properties` file on Windows. (The given Windows path is the default; however, a different path may have been selected during installation.)

Viewing HP Systems Insight Manager Events

gWLM allows you to configure a number of events that 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.

5 Security

This chapter highlights several security items you should be aware of while using gWLM. It addresses the following topics:

- "General Security Topics"
- "Securing gWLM Communications"

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 "Securing gWLM Communications" (page 35).
- HP SIM allows you to create user roles with different levels of privileges. For more information, see the HP SIM documentation.

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→Secure gWLM Communications…

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

Alternatively, on OpenVMS and HP-UX, you can secure communications manually by following the steps outlined in *gwlmsslconfig*(1M).

NOTE: HP strongly recommends that you always use gWLM with its communications secured.

6 Additional Configuration and Administration Tasks

This chapter covers various configuration and administration tasks. It addresses the following topics:

- "Manually Adjusting CPU Resources"
- "Setting Aside Space for Historical Data"
- "Setting gWLM Properties"
- "Automatic Restart of gWLM's Managed Nodes in SRDs (High Availability)"
- "Nesting Partitions"
- "Changing the gWLM Resource Allocation Interval"

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 nor its policies are configured to use TiCAP, you may be able to temporarily provide additional resources to a deployed SRD by:

- Activating an iCAP core.
- 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 that 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 that policies used in the SRD do not unintentionally limit their associated workloads' resource requests.

To use 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 that policies used in the SRD do not unintentionally limit their associated workloads' resource requests.

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

> **NOTE:** After manually adjusting the number of cores in an SRD, always confirm that the changes after two gWLM resource allocation intervals have passed. Changes may not be as expected due to the gWLM behavior. For example:
>
> In an SRD with nested partitions, gWLM samples the inner partitions for their sizes before it samples the outer partitions. Adjusting resources between these samplings can cause gWLM to report incorrect sizes. If you encounter this issue, try readjusting.

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 use 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 that the latest gWLM data is backed up, issue the following command before using the --backup option:

  ```
  gwlm history --flush
  ```

- If the CMS or the managed nodes have been modified (such as changes in the 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 all the managed nodes. Read the files for information on the behaviors they control.

CMS Properties

The CMS properties are in the `/etc/opt/gwlm/conf/gwlmcms.properties` file on HP-UX and in the `C:\Program Files\HP\Virtual Server Environment\conf\gwlmcms.properties` file on Windows. (The given Windows path is the default; however, a different path may have been selected during 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. See the file for information on when individual properties are read.
>
> Restarting `gwlmcmsd` temporarily disables the 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: 2008/12/02 20:17:18 $
# $Revision: 1.1 $
#
# Contents:
# This file contains the default configuration values for the
# Global Workload Manager CMS.
#
# You must restart gwlmcmsd for changes made to this file
# to take effect--unless otherwise noted.
#
#
# Set FileHandler log level for the following log files:
# /var/opt/gwlm/gwlmcmsd.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
#
```

```
# Support for securing Oracle communication.
#
# com.hp.gwlm.jdbc.oracle.secure:
# Whether communication with Oracle server is secure or not. Possible
# values are 'on' and 'off'. Default is off.
#
# oracle.net.encryption_types_client:
# Secure communication encryption type. Possible values are RC4_256,
# RC4_128, RC4_56, RC4_40, 3DES112, 3DES168. See Oracle documentation
# for details.
#
# oracle.net.crypto_checksum_types_client:
# Encryption type for Oracle secure communication integrity checking.
#
# Make sure that the parameter values specified here match those on
# the Oracle server.
#
com.hp.gwlm.jdbc.oracle.secure = off
oracle.net.encryption_types_client = 3DES112
oracle.net.crypto_checksum_types_client = MD5
```

Agent Properties

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

📝 **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: 2008/12/02 20:17:18 $
# $Revision: 1.1 $
#
# 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 want to change these ports, add the following lines to both the properties files:

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

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

- `gwlmcms.properties`

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

  On OpenVMS, it is in sys$specific:[gwlm.conf]:

  sys$specific:[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 `gwlmagent` 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 the HP Virtualization Manager and HP Capacity Advisor.

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 that you must perform for this behavior to happen is to edit the property`com.hp.gwlm.node.HA.minimumTimeout` in the file `sys$specific:[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 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 (LOGIN.COM). The agent then checks the file `sys$specific:[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 `gwlmstatus` command to monitor availability of the node. 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 *gwlmstatus*(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, see the gWLM's "Events" menu.

You can then view these events using the Event Lists item on 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 `gwlmagent` on each managed node in the affected SRD:

   `$ @sys$startup:gwlm$shutdown.com`

   `$ @sys$startup:gwlm$startup.com`

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 `sys$specific:[gwlm.log]gwlmagent_log.0` and `/var/opt/gwlm/gwlmcmsd.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 view 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 receive *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 `gwlmagent` on each managed node in the affected SRD:

`$ @sys$startup:gwlm$shutdown.com`

`$ @sys$startup:gwlm$startup.com`

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 `gwlm undeploy` command.

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

gwlm 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 gwlm command is added to the path during installation. On HP-UX systems, the command is in the /opt/gwlm/bin/ directory. On Microsoft Windows systems, the command is in the C:\Program Files\HP\Virtual Server Environment\bin\gwlm\ by default. On OpenVMS systems, the command is in sys$common:[gwlm.bin].

> **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 or System group on OpenVMS to run the following commands:

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

    $ delete sys$specific:[gwlm]deployed.config;*

2. Force an undeploy of the SRD (named *SRD* below) to ensure that the CMS and the managed nodes agree on the SRD's state. Run the following command on the CMS:

    # **gwlm undeploy --srd=*SRD* --force**

3. Restart the gwlmagent daemon on each managed node:

    $ @sys$startup:gwlm$shutdown.com

    $ @sys$startup:gwlm$startup.com

> **NOTE:**   If the gWLM CMS and agent disagree about whether an SRD is deployed or undeployed, you can use the --force option using the gwlm deploy or gwlm undeploy commands.
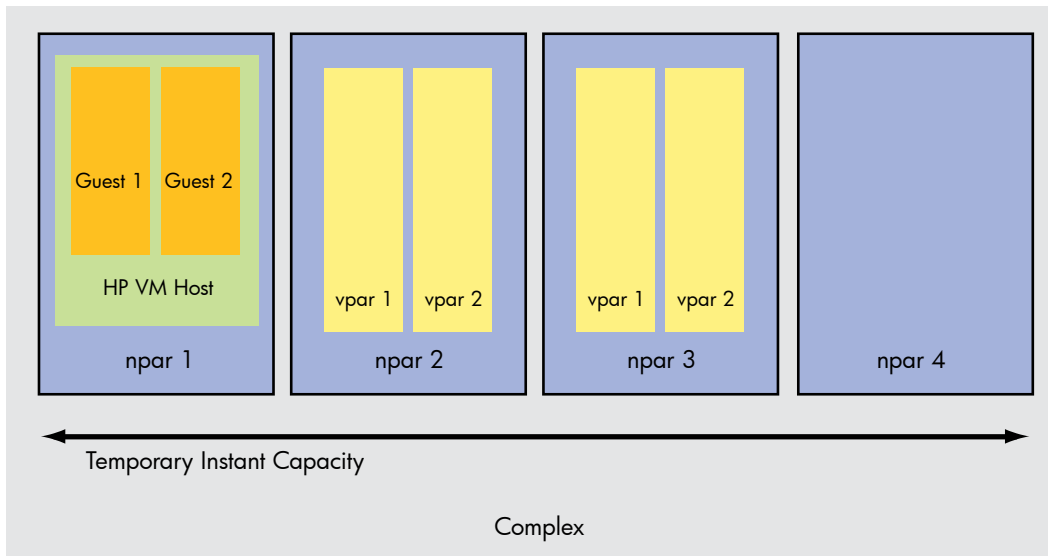
Nesting Partitions

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 6-1. The complex has four nPars, 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 6-1 Nested Partitions

For more information on nesting partitions, see the online help or *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

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 gwlm command and a text editor to change the interval on the command line:

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

3. Use the `gwlm import --clobber` command 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.

A Compatibility with Agents

The gWLM A.04.01.00.* 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.02.50.00.04: 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
- gWLM A.04.01.00.*: 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 gwlmstatus 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 belong to 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 gwlmxml(4).