

Getting Started with gWLM

Version A.01.01.x



Manufacturing Part Number: T2762-90005

May 2005

© Copyright 2004-2005 Hewlett-Packard Development Company, L.P.

Publication history

Second edition

May 2005

T2762-90005

HP-UX 11i v1, HP-UX 11i v2 Update 2, and Linux

First edition

December 2004

T2762-90002

HP-UX 11i v1, HP-UX 11i v2 Update 2, and Linux

Notice

© Copyright 2004-2005 Hewlett-Packard Development Company, L.P. All Rights Reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material.

Contents

Assumptions	1
What is gWLM?	2
What are the benefits of using gWLM?	2
What do I need to know to use gWLM?	3
What is the gWLM management model?	5
What types of policies are available?	6
Choosing a policy type	8
Choosing between an OwnBorrow policy and a utilization policy	9
Combining the different policy types	9
Can I see how gWLM will perform—without actually affecting my system?	10
How do I get started with gWLM?	10
How do I stop gWLM?	13
Show me gWLM in action	14
What are some common gWLM tasks?	17
Fixing the amount of CPU a workload gets	17
Resizing a workload's npar, vpar, pset, or fss group as needed	17
How do I monitor gWLM's effect on my workloads?	18
Graphical reports	18
High-level view	18
Other administration tasks	19
Initial steps	19
How do I change from advisory mode to managed mode?	20
How do I create a new policy?	21
How do I edit a policy?	22
How do I change which policy is associated with a workload?	23
How do I add a new npar / vpar / pset / fss group to an SRD?	24
How do I stop managing a workload?	24
How do I stop managing an SRD?	26
Finding more information	27
Providing feedback	29

Contents

This document presents an overview of the techniques and tools available for using Global Workload Manager, or gWLM. It exposes you to the essentials and allows you to quickly get started using gWLM.

Assumptions

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

- HP Systems Insight Manager
- HP Global Workload Manager

For information on setting up HP Systems Insight Manager, see the *HP Systems Insight Manager Installation and User Guide*, available from <http://www.hp.com/go/hpsim>.

If you have not installed gWLM, the following steps give an overview. When you install gWLM, you must:

Step 1. Decide which system will be your central management server (CMS), then install the gWLM CMS software on that system.

This system must also have HP Systems Insight Manager installed and running.

Step 2. Initialize the CMS for gWLM by running `gwlminitconfig --initconfig`.

Step 3. Start the gWLM CMS daemon `gwlmcmsd`.

Step 4. 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, as described in the gWLM Release and Installation Notes.)

Step 5. On each managed node, start the gWLM agent daemon `gwlagent`.

For detailed instructions on installing gWLM for HP-UX and Linux, see the gWLM Release and Installation Notes.

What is gWLM?

What is gWLM?

gWLM is a tool that allows you to centrally define resource-sharing policies that you can use across multiple HP servers. These policies increase system utilization and facilitate controlled sharing of system resources. gWLM's monitoring abilities provide both real-time and historical monitoring of the resource allocation.

What are the benefits of using gWLM?

The benefits of gWLM include:

- Better use of the server capacity you already have
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 on a single server and make use of the reserve capacity.
- Confidence that your mission-critical workloads get the resources they need
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, but also to dedicate those resources to workloads when there are spikes in resource demand.
- Reduced system administration costs
With gWLM managing resource allocation, you can combine more workloads on fewer servers—reducing administration costs.

What do I need to know to use gWLM?

Here are some terms to know when using gWLM:

Workload

The collection of processes that are within a single compartment. The compartment can be an nPartition (npar), a virtual partition (vpar), 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, psets, and fss groups, see the section “What is the gWLM management model?” on page 5.)

Compartment

An npar, a vpar, a pset, or an fss group with its resource allocation being managed by gWLM.

Multiple compartments—all of the same type—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 SRD. gWLM manages each workload by adjusting the resource allocation for its compartment.

Shared Resource Domain (SRD)

A collection of compartments—all of the same type—that share system resources. The compartments can be npars, vpars, psets, or fss groups. A server containing npars can be an SRD—as long as the requirements in the section “What is the gWLM management model?” on page 5 are met. A server or an npar divided into vpars can be an SRD for its vpar compartments. Similarly, a server, an npar, or a vpar containing psets can be an SRD for its pset compartments. Lastly, a server, an npar, or a vpar containing fss groups can be an SRD for its fss group compartments.

gWLM creates SRDs when you use the Manage New Systems wizard, reached through the menu

Optimize

-> **Global Workload Manager (gWLM)**

-> **Manage New Systems**

What do I need to know to use gWLM?

Policy

A policy is 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 guaranteed when needed) as well as how much of those resources the workload can lend to other workloads.

A single policy can be associated, or applied, to multiple workloads.

Mode

gWLM offers two modes: advisory and managed. Advisory mode allows you to see what CPU requests gWLM would make for a workload—without actually affecting resource allocation. You can 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.

You can only set the mode on the SRD level: All compartments within an SRD operate in the same mode, either advisory or managed.

Deploy

Enable gWLM control of an SRD.

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 CPUs between psets.

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

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 between workloads in the SRD. If the SRD is in advisory mode, gWLM no longer provides information on what requests would have been made.

What is the gWLM management model?

gWLM was designed to support utility computing across a data center by providing resource-sharing policies that you centrally create and monitor. gWLM moves resources between 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, discussed below. These divisions are referred to as compartments in gWLM.

- HP-UX Hardware Partitions (npars)

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.

Using the HP product Instant Capacity, gWLM simulates the movement of CPUs between npars by turning off an active CPU in one npar then turning on a deactivated CPU in another npar in the same complex. Thus, the first npar has one less active CPU, while the second npar has one additional active CPU. (gWLM maintains the number of active CPUs, honoring the Instant Capacity licensing limits. As a result, no costs are incurred.)

- HP-UX Virtual Partitions (vpars)

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.

- Processor sets (psets)

A processor set is a collection of processors (CPUs) grouped together for the exclusive access by applications assigned to that processor set. Each application runs only on processors in its assigned processor set.

Compartments based on psets are available on HP-UX and Linux. (On Linux, gWLM simulates psets using CPU affinity masks.)

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

A group of processes that has its CPU 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 CPUs, to the group of processes.

What types of policies are available?

For additional information on these partitions, visit:

- HP Virtual Server Environment web site:
<http://www.hp.com/go/vse>
- The “partitioning continuum and the flexible computing environment” web site:
<http://docs.hp.com/hpux/hplex/index.html>

As noted above, gWLM can manage workloads on both HP-UX and Linux.

gWLM manages resources based on the following model:

1. You decide which of your systems you want to manage and what type of compartment you want to use. (For psets and fss groups, gWLM creates the compartments for you.)
2. You assign each workload to a compartment.
3. You associate a policy to 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.)
4. gWLM monitors the CPU consumption of all the workloads in the SRD during the current allocation interval.
5. At the end of the interval, gWLM adjusts the CPU allocations for the compartments in accordance with the policies.
6. gWLM repeats the previous two steps.

For information on what types of workloads to combine for optimal resource utilization, see the online help topic “Getting the most out of gWLM,” available in gWLM’s graphical interface in HP Systems Insight Manager.

What types of policies are available?

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

Fixed policy

This type of policy guarantees that a workload’s compartment has a fixed (constant) amount of CPU resources.

What types of policies are available?

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

Utilization policy

This type of policy has a target based on utilization. With a CPU utilization policy, gWLM attempts to keep a workload's CPU utilization below the target by adding CPU resources when the workload is using too much of its current CPU allocation. For example, assume a workload has a utilization policy with a target of 80% and an allocation of 5 CPUs. If the workload is consuming 4.5 CPUs, 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 CPUs would result in a utilization percentage of $4.5/6$, or 75%, meeting the target.

With a utilization policy, you specify the minimum and maximum CPU requests. Workloads with this type of policy are only guaranteed the minimum request. Utilization policies allow you to prioritize workloads.

OwnBorrow policy

This type of policy allows you to set the:

- Amount of CPU resources a workload's compartment owns
- Minimum amount of CPU resources a workload's compartment must have (after lending resources to other workloads)
- Maximum amount of CPU resources a workload's compartment can have (after borrowing resources from other workloads)

The compartment of a workload with such a policy is guaranteed the owned amount of CPU when needed. The minimum CPU request and maximum CPU request 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 CPUs and that compartment's workload becomes busy, the compartment re-acquires those lent-out CPUs immediately.

Custom policy

Custom policies are available for advanced users. For information on these policies, see the online help or the `gwlxml(4)` man page.

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

Choosing a policy type

How do you decide which policy type to use? The table below 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 1 **Choosing a policy type**

If...	Use the following type of policy...
You want gWLM to allocate a constant amount of CPU to a workload	Fixed
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 the gwlreport(1M) man page.
You have static vpars but you want to move to a model where CPUs migrate between vpars	OwnBorrow For each vpar, set its number of owned CPUs to its static number of CPUs. The vpar is then guaranteed those owned CPUs when needed.
You have npars but you want to move to a model where CPUs migrate between npars	OwnBorrow Install the HP product Instant Capacity on each npar. (This software allows gWLM to simulate CPU movement among the npars.) For each npar, set its number of owned CPUs to the number of CPUs you want to guarantee the npar has when needed.
You want to tap into a pool of resources taking or giving CPU resources as needed—with no guaranteed access to resources beyond a minimum request	Utilization

Choosing between an OwnBorrow policy and a utilization policy

OwnBorrow policies guarantee a certain amount of CPU resources are available to your workload. This amount is based on the owned CPU resources specified in the policy. The workload can lend unneeded resources to other workloads—and get those resources back as soon as they are needed again.

With utilization policies, gWLM guarantees CPU resources to the workload based on the minimum CPU value specified in the policy. These resources cannot be shared with other workloads.

Combining the different policy types

Each workload in an SRD must have a policy. You can assign different policy types to the workloads in an SRD as long as the types are compatible. The following table indicates the compatible combinations.

Table 2 Compatible combinations of policy types

Policy type	Fixed	Utilization	OwnBorrow
Fixed	Yes	Yes	Yes
Utilization	Yes	Yes	—
OwnBorrow	Yes	—	Yes

Combinations with custom policies are addressed in the online help and the gwlmxml(4) man page.

[Can I see how gWLM will perform—without actually affecting my system?](#)

Can I see how gWLM will perform—without actually affecting my 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.

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, see “How do I change from advisory mode to managed mode?” on page 20.

How do I get started with gWLM?

gWLM is accessed through HP Systems Insight Manager.

After performing the necessary gWLM daemon configuration as described in the gWLM installation notes, the quickest way to start using gWLM to manage new systems is to use the Manage New Systems 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 between npars, vpars, processor sets, or fss groups. (CPU migration between npars is simulated using the HP product Instant Capacity, as explained in the section “What is the gWLM management model?” on page 5.)

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 Systems Insight Manager, you must have authorizations for “All Tools” or “gWLM All Tools.”

Step 1. On your gWLM CMS:

- a.** Check that HP Systems Insight Manager is running:

```
# ps -ef | grep mxadmin
```

If there is no match for `mxadmin`, start HP Systems Insight Manager:

```
# /opt/mx/bin/mxstart
```

- b.** Start the gWLM CMS daemon if it is not already running:

```
# /opt/gwlm/bin/gwlmcmd
```

Step 2. On each managed node, start the gWLM agent if it is not already running:

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

Step 3. Connect to HP Systems Insight Manager by pointing your web browser to:

```
http://hostname:280
```

where *hostname* represents the hostname of the CMS.

Step 4. Select the menus:

Optimize

-> **Global Workload Manager (gWLM)**

-> **Manage New Systems...**

How do I get started with gWLM?

The wizard guides you through the following steps:

Step 1. Entering system names

You must enter the names of the systems (managed nodes) where the workloads you want gWLM to manage are running.

Step 2. Forming the SRDs

Based on the systems you entered, gWLM will generate default SRDs and compartments. You can influence how SRDs are formed by indicating a preference for npar compartments, vpar compartments, pset compartments, or fss group compartments, in case multiple compartments are present on a system.

NOTE

gWLM displays only the compartment types available for the system. For information on making other compartment types available, see the gWLM Release and Installation Notes.

Step 3. Indicating mode of operation for each SRD

gWLM allows you to run gWLM in either advisory mode or managed mode. In advisory mode, gWLM only reports how it would migrate CPU resources between compartments. With managed mode, gWLM actually migrates the CPU resources between compartments.

You may want to change the SRD names (from what gWLM generates) to be meaningful, thereby simplifying working with the SRDs later.

Step 4. Associating, or applying, policies to workloads

gWLM comes with a number of policies for:

- Fixing a workload's CPU allocation to a certain amount of CPU resources
- Adjusting a workload's CPU allocation based on its CPU utilization
- Optimizing the borrowing/lending of CPU resources between workloads based on their CPU utilization

You must associate a policy to each workload.

You may want to change the workload names (from what gWLM generates) to be meaningful, thereby simplifying working with the workloads later.

Step 5. Review the SRD

Verify the SRD is as desired and select the [Finish] button to have gWLM manage the resource allocation for the workloads in the SRD.

How do I stop gWLM?

To turn off gWLM:

NOTE You must be logged in as root on the systems where you run the `gwlmcmsd` and `gwlmagent` commands mentioned below. In HP Systems Insight Manager, you must have authorizations for “All Tools” or “gWLM All Tools.”

Step 1. Stop all SRD management, as explained in the section “How do I stop managing an SRD?” on page 26

Step 2. On the gWLM CMS:

a. Stop the gWLM daemon:

```
# /opt/gwlm/bin/gwlmcmsd --stop
```

b. Ensure the file `/etc/rc.config.d/gwlmCtl` on HP-UX (`/etc/sysconfig/gwlmCtl` on Linux) has `GWLM_CMS_START` set to 0:

```
GWLM_CMS_START=0
```

This setting prevents `gwlmcmsd` from starting automatically when the system boots.

Show me gWLM in action

Step 3. On each managed node:

- a. Stop the gWLM daemon:

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

- b. Ensure the file `/etc/rc.config.d/gwlmCtl` on HP-UX (`/etc/sysconfig/gwlmCtl` on Linux) has `GWLM_AGENT_START` set to 0:

```
GWLM_AGENT_START=0
```

This setting prevents `gwlmagent` from starting automatically when the system boots.

Show me gWLM in action

This section helps you see gWLM move CPUs between vpars. You can use similar steps to see CPUs move between npars, psets, or fss groups. For psets and fss groups, though, you will need to put processes in the desired pset or fss group. (You place processes by using either the Applications tab on the Edit Workloads window or 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 CPUs that gWLM can move between them
- HP Systems Insight Manager and the gWLM CMS software are installed 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`, `gwlmcmd`, and `gwlmagent` commands mentioned below. In HP Systems Insight Manager, you must have authorizations for “All Tools” or “gWLM All Tools.”

Step 1. Check that HP Systems Insight Manager is running on the gWLM CMS (vpar3):

```
# ps -ef | grep mxadmin
```

If there is no match for `mxadmin`, start HP Systems Insight Manager:

```
# /opt/mx/bin/mxstart
```

Step 2. Start the gWLM CMS daemon on vpar3:

```
# vpar3> /opt/gwlm/bin/gwlmcmd
```

Step 3. Start the gWLM agent daemon on vpar1:

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

Step 4. Start the gWLM agent daemon on vpar2:

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

Step 5. Connect to HP Systems Insight Manager by pointing your web browser to:

```
http://hostname:280
```

where *hostname* represents the hostname of the CMS, in this case, vpar3.

Step 6. Create a gWLM SRD containing the two vpars by following the steps given in “How do I get started with gWLM?” on page 10

- a. For system names, enter the names of the two vpars separated by a space
- b. For the compartment type, select “Virtual Partition (vpar)”
- c. (Optional) Choose a meaningful name to use in place of the generated name for the SRD
- d. Ensure the mode is Managed

Show me gWLM in action

- e. Use an OwnBorrow policy for the vpar1 workload and for the vpar2 workload

An OwnBorrow policy has a name of the form:

Owns_4_CPUs-Max_8_CPUs

Ideally, the number of CPUs owned by the workloads will equal the total number of CPUs in the SRD when using this type of policy. (You may need to edit policies to achieve this equality.)

- f. Confirm and finish the SRD creation

Step 7. View gWLM's real-time reports to show CPU allocation for vpar1 by selecting the following menu in HP Systems Insight Manager:

Optimize

-> **Global Workload Manager (gWLM)**

-> **View Real-time Reports...**

Step 8. Select the vpar1 workload

Step 9. Select the [View Policy Detail] button to see a graph of vpar1's CPU allocation

Step 10. Start a CPU-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:

```
# /opt/perl/bin/perl -e 'print "$$\n";while (1) {};' &
```

```
[1] 24379
```

This command will consume most of a single CPU. Start multiple copies of the command to consume additional CPUs.

Step 11. Wait a few minutes then look at the "Policy Request and Workload Allocation" graph from Step 9 to see how the policy for vpar1's workload is requesting CPU allocations and how gWLM is granting them

Step 12. Kill the workloads you started in vpar1

Step 13. Repeat Step 7 through Step 12 for the vpar2 workload if you would like to see gWLM move CPUs to vpar2

What are some common gWLM tasks?

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

Fixing the amount of CPU a workload gets

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

To fix the amount of CPU 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 “How do I get started with gWLM?” on page 10
- When adding a workload to an SRD, as described in “How do I add a new npar / vpar / pset / fss group to an SRD?” on page 24
- By changing the policy associated with an existing workload, as described in “How do I change which policy is associated with a workload?” on page 23

Resizing a workload’s npar, vpar, pset, or fss group as needed

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

With such a policy, you indicate how much CPU a workload should own. The workload is then guaranteed this owned amount of CPU—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 “How do I get started with gWLM?” on page 10
- When adding a workload to an SRD, as described in “How do I add a new npar / vpar / pset / fss group to an SRD?” on page 24
- By changing the policy associated with an existing workload, as described in “How do I change which policy is associated with a workload?” on page 23

gWLM’s utilization policies also allow resizing.

How do I monitor gWLM's effect on my workloads?

How do I monitor gWLM's effect on my workloads?

For monitoring gWLM, there are several methods, as described below.

Graphical reports

gWLM offers graphs showing either real-time or historical data through HP Systems Insight Manager. For information on interpreting these reports, see the online help.

Navigate to real-time reports through the menus:

- Optimize**
- > **Global Workload Manager (gWLM)**
- > **View Real-time Reports...**

Navigate to historical reports through the menus:

- Optimize**
- > **Global Workload Manager (gWLM)**
- > **View Historical Reports...**

High-level view

For a high-level view of the performance of your SRDs and workloads, navigate through the menus:

- Optimize**
- > **Global Workload Manager (gWLM)**
- > **View Summaries...**

Other administration tasks

This section discusses various administration tasks:

- How do I change from advisory mode to managed mode?
- How do I create a new policy?
- How do I edit a policy?
- How do I change which policy is associated with a workload?
- How do I add a new npar / vpar / pset / fss group to an SRD?
- How do I stop managing a workload?
- How do I stop managing an SRD?

Initial steps

Several of the tasks below require the same initial steps. These steps are given below.

NOTE You must be logged in as root on the systems where you run the `mxstart`, `gwlmcmsd`, and `gwlmagent` commands mentioned below.

Step 1. On your gWLM CMS:

- a.** Check that HP Systems Insight Manager is still running:

```
# ps -ef | grep mxadmin
```

If there is no match for `mxadmin`, start HP Systems Insight Manager:

```
# /opt/mx/bin/mxstart
```

- b.** Start the gWLM CMS daemon if it is not already running:

```
# /opt/gwlm/bin/gwlmcmsd
```

Step 2. On each managed node, start the gWLM agent if it is not already running:

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

How do I change from advisory mode to managed mode?

Advisory mode allows you to see what CPU requests gWLM would make for a workload—without actually affecting resource allocation. 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 Systems Insight Manager, you must have authorizations for “All Tools” or “gWLM 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, vpar, and pset compartments with the CPU counts 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 (15 seconds by default) to pass.

Step 1. Ensure HP Systems Insight Manager, the gWLM CMS daemon (`gwlmcmsd`) and all the gWLM agents (`gwlagent`) are still running, as explained in the section “Initial steps” on page 19.

Step 2. Connect to HP Systems Insight Manager by pointing your web browser to:

`http://hostname:280`

where *hostname* represents the hostname of the CMS.

Step 3. Select, in HP Systems Insight Manger, the menu item

Optimize
-> **Global Workload Manager (gWLM)**
-> **Edit SRDs...**

Step 4. Select the SRD for which to change the mode

Step 5. Select the [Edit] button

Step 6. Scroll down to see the new section that has been added to the page:

Edit Shared Resource Domain

Step 7. Select the radio button corresponding to the mode you desire

Step 8. Select the [OK] button

How do I create a new policy?

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

NOTE In HP Systems Insight Manager, you must have authorizations for "All Tools" or "gWLM All Tools."

Step 1. Ensure HP Systems Insight Manager, the gWLM CMS daemon (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Initial steps" on page 19.

Step 2. Connect to HP Systems Insight Manager by pointing your web browser to:

`http://hostname:280`

where *hostname* represents the hostname of the CMS.

Step 3. Select, in HP Systems Insight Manger, the menu item

Optimize
-> **Global Workload Manager (gWLM)**
-> **Edit Policies...**

Step 4. Select the [New] button

Step 5. Scroll down to see the new section that has been added to the page:

New Policy

Step 6. Edit the settings, selecting a policy type and specifying the required values and optional values as desired

Step 7. Select the [OK] button

Other administration tasks

How do I edit 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 Systems Insight Manager, you must have authorizations for "All Tools" or "gWLM All Tools."

Step 1. Ensure HP Systems Insight Manager, the gWLM CMS daemon (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section "Initial steps" on page 19.

Step 2. Connect to HP Systems Insight Manager by pointing your web browser to:

`http://hostname:280`

where *hostname* represents the hostname of the CMS.

Step 3. Select, in HP Systems Insight Manger, the menu item

Optimize
-> **Global Workload Manager (gWLM)**
-> **Edit Policies...**

Step 4. Select the policy to edit

Step 5. Select the [Edit] button

Step 6. Scroll down to see the new section that has been added to the page:

Edit Policy

Step 7. Edit the settings

Step 8. Select the [OK] button

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

How do I change which policy is associated with a workload?

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

NOTE In HP Systems Insight Manager, you must have authorizations for “All Tools” or “gWLM All Tools.”

Step 1. Ensure HP Systems Insight Manager, the gWLM CMS daemon (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section “Initial steps” on page 19.

Step 2. Connect to HP Systems Insight Manager by pointing your web browser to:

`http://hostname:280`

where *hostname* represents the hostname of the CMS.

Step 3. Select, in HP Systems Insight Manger, the menu item

Optimize
-> **Global Workload Manager (gWLM)**
-> **Edit Associations...**

Step 4. Select the workload for which you want to change the policy

Step 5. Select the [Change Policy for Workload] button

Step 6. Scroll down to see the new section that has been added to the page:

Edit Association

Step 7. Select the new policy to associate, or apply, to the workload from the pull-down menu

Step 8. Select the [OK] button

How do I add a new npar / vpar / pset / fss group to an SRD?

If you have added an npar or a vpar to your system, set gWLM to manage the new compartment by selecting, in HP Systems Insight Manager, the menu item:

- Optimize**
- > **Global Workload Manager (gWLM)**
- > **Manage Workloads...**

You can use the same menu item to have gWLM create multiple psets or fss groups for you.

How do I 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 npars or vpars with the CPU counts they had in the last allocation interval. For this reason, in Step 4 below, you associate fixed policies with the npar-based or vpar-based workloads. 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:

Step 1. Ensure HP Systems Insight Manager, the gWLM CMS daemon (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section “Initial steps” on page 19.

Step 2. Connect to HP Systems Insight Manager by pointing your web browser to:

`http://hostname:280`

where *hostname* represents the hostname of the CMS.

- Step 3.** Associate fixed policies with all workloads that you want to unmanage that are based on npars or vpars:
- a. Select, in HP Systems Insight Manager, the menu item
Optimize
-> **Global Workload Manager (gWLM)**
-> **Edit Associations...**
 - b. Select an npar-based or vpar-based workload to unmanage
 - c. Select the [Change Policy for Workload] button
 - d. Scroll down to see the new Edit Association section
 - e. Select a fixed policy
 - f. Select the [OK] button
 - g. Repeat Step b through Step f for any additional npar-based or vpar-based workloads to unmanage
- Step 4.** Wait an allocation interval (15 seconds by default) for gWLM to set CPU allocations based on the fixed policies.
- Step 5.** Select, in HP Systems Insight Manger, the menu item
Optimize
-> **Global Workload Manager (gWLM)**
-> **Unmanage Workloads...**
- Step 6.** Select one or more workloads
- a. Select one or more (but not all) workloads you want to unmanage (You cannot select all workloads to unmanage because a deployed SRD must have at least one workload. To stop managing an SRD, undeploy it.)
 - b. Select the [Next] button
- Step 7.** Associate policies
- a. Evaluate and change, if needed, the remaining workloads and their associated policies to ensure they are appropriate given that a workload has been removed
 - b. Select the [Next] button

Other administration tasks

Step 8. Review summary

- a. Review the resulting SRD configuration
- b. Select the [Prev] button to make edits
- c. Select the [Finish] button if the configuration is as desired and you are ready for gWLM to begin managing the reconfigured SRD

How do I stop managing an SRD?

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

NOTE In HP Systems Insight Manager, you must have authorizations for “All Tools” or “gWLM All Tools.”

Step 1. Ensure HP Systems Insight Manager, the gWLM CMS daemon (`gwlmcmsd`) and all the gWLM agents (`gwlmagent`) are still running, as explained in the section “Initial steps” on page 19.

Step 2. Connect to HP Systems Insight Manager by pointing your web browser to:

`http://hostname:280`

where *hostname* represents the hostname of the CMS.

Step 3. Associate fixed policies with all npars or vpars that were in the SRD

When gWLM stops managing an SRD, it leaves npar compartments and vpar compartments 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.)

- a. Select, in HP Systems Insight Manager, the menu item
Optimize
-> **Global Workload Manager (gWLM)...**
-> **Edit Associations...**
- b. Select an npar-based or vpar-based workload to unmanage
- c. Select the [Change Policy for Workload] button
- d. Scroll down to see the new Edit Association section

- e. Select a fixed policy
- f. Select the [OK] button
- g. Repeat Step b through Step f for any additional npar-based or vpar-based workloads to unmanage

Step 4. Select, in HP Systems Insight Manger, the menu item

Optimize
 -> **Global Workload Manager (gWLM)...**
 -> **Edit SRDs...**

Step 5. Select the SRD you want to stop managing (undeploy)

Step 6. Scroll down to see the new section that has been added to the page:

Edit Shared Resource Domain

Step 7. Select the radio button corresponding to Undeployed for the SRD State

Step 8. Select the [OK] button

Finding more information

The table below indicates the documents that explain various topics.

Table 3 **Where to find additional information**

To...	See...
Use gWLM immediately, reading as little as possible	gWLM Home Page in HP Systems Insight Manager (Optimize -> Global Workload Manager (gWLM) -> Getting Started - gWLM Home) or Concepts topic in online help or Getting started with gWLM (this document) (/opt/gwlm/doc/getting.started.with.gWLM.pdf)

[Finding more information](#)

Table 3 **Where to find additional information (Continued)**

To...	See...
Learn about gWLM concepts	Concepts topic in online help or Getting started with gWLM (this document) (/opt/gwlm/doc/getting.started.with.gWLM.pdf)
Learn gWLM terms	“What do I need to know to use gWLM?” on page 3 or gWLM glossary in online help
Learn gWLM best practices	“Getting the most out of gWLM” topic in online help
Learn about other gWLM features	gWLM: Reference and additional topics (/opt/gwlm/doc/gWLM.reference.pdf)
Learn about the gWLM interface in HP Systems Insight Manager	online help
Learn about the gWLM command-line interface	gwlm(1M) man page
Learn about gWLM daemons	gwlmcmd(1M) man page
Learn about using secure communications with gWLM	gwlmsslconfig(1M) man page
Learn more about npars, vpars, and psets	<ul style="list-style-type: none"> • HP Virtual Server Environment web site: http://www.hp.com/go/vse • The “partitioning continuum and the flexible computing environment” web site: http://docs.hp.com/hpux/hplex/index.html
Learn about using secure communications with gWLM	gwlmsslconfig(1M) man page

Providing feedback

- Email your feedback to the gWLM development team at the following address:
gwlfeedback@rsn.hp.com
- 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>

Providing feedback