

**HP-UX Workload Manager A.03.02.xx  
Release Notes  
for HP-UX 11i v1, HP-UX 11i v2, and  
HP-UX 11i v3**



i n v e n t

**Manufacturing Part Number: B8843-90038**

**October 2007**

© Copyright 2000-2007 Hewlett-Packard Development Company, L.P.



\* B 8 8 4 3 - 9 0 0 3 8 \*

---

## Legal Notices

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

Hewlett-Packard makes no warranty of any kind with regard to this documentation, 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 direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

© Copyright 2000-2007 Hewlett-Packard Development Company, L.P.

This document contains information that is protected by copyright. All rights are reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

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

Use of this documentation and flexible discs, compact discs, or tape cartridges supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs in their present form or with alterations is expressly prohibited.

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Oracle is a registered trademark of Oracle Corporation, Redwood City, California.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

SAP is a registered trademark of SAP AG in Germany and several other countries.

The SNMP Toolkit uses a library written by CMU:

© Copyright 1998 by Carnegie Mellon University

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

CMU disclaims all warranties with regard to this software, including all implied warranties of merchantability and fitness, in no event shall CMU be liable for any special, indirect or consequential damages or any damages whatsoever resulting from loss of use, data or profits, whether in an action of contract, negligence or other tortious action, arising out of or in connection with the use or performance of this software.

BEA, Tuxedo, and WebLogic are registered trademarks and BEA WebLogic Enterprise Platform, BEA WebLogic Server, BEA WebLogic Integration, BEA WebLogic Portal, BEA WebLogic JRockit, BEA WebLogic Platform, BEA WebLogic Express, BEA WebLogic Workshop, BEA WebLogic Java Adapter for Mainframe, BEA Liquid Data for WebLogic and BEA eLink are trademarks of BEA Systems, Inc.

HP-UX Workload Manager and JFreeChart:

Starting with version A.02.01, HP-UX Workload Manager uses JFreeChart, an open source software package for displaying charts and graphs. JFreeChart is licensed under the GNU Lesser General Public License (LGPL). A copy of this license is available at </opt/wlm/lib/mongui/LGPL.txt> and at the following GNU Web site:

<http://www.gnu.org/licenses/lgpl.txt>.

The version of JFreeChart that Workload Manager uses is a modified version of JFreeChart 0.9.4. For information on the modifications, refer to </opt/wlm/lib/mongui/README>. The latest version of JFreeChart is available at the following Web sites:

<http://www.object-refinery.com/jfreechart/>

<http://www.sourceforge.net>

HP-UX Workload Manager and libxml2:

Starting with version A.03.00, HP-UX Workload Manager uses libxml2, the XML C parser and toolkit developed for the Gnome project at MIT. Though it is written in the C language, a variety of language bindings make it portable to other environments. It is licensed under the agreement available at `/opt/wlm/lib/README.libxml2`. The version of libxml2 that HP-UX Workload Manager uses is libxml2 2.6.10.

# 1 HP-UX Workload Manager Release Notes

This release notice contains the following information about HP-UX Workload Manager A.03.02.xx:

- Announcement
- New in this version
- Known problems and workarounds
- Compatibility information and installation requirements
- Patches and fixes in this version
- Software availability in native languages
- Security
- Available manuals
- WLM toolkits
- Providing feedback
- Training

---

**NOTE** For information on WLM's support policy and patch policy, visit the following Web site:

<http://www.hp.com/go/wlm>

These policies indicate the time periods for which this version of WLM is supported and patched.

---

**Announcement**

---

**Announcement**

HP-UX WLM A.03.02 runs under the following HP-UX operating systems and hardware:

<b>Operating Systems</b>	<b>Hardware</b>
HP-UX 11i v1 (B.11.11)	HP 9000 servers
HP-UX 11i v2 (B.11.23)	HP 9000 servers and HP Integrity servers
HP-UX 11i v1 (B.11.11) and HP-UX 11i v2 (B.11.23)	Servers combining HP 9000 partitions and HP Integrity partitions; in such environments, HP-UX 11i v1 (B.11.11) supports HP 9000 partitions only

HP-UX WLM A.03.02.02 runs under the following HP-UX operating systems and hardware:

<b>Operating Systems</b>	<b>Hardware</b>
HP-UX 11i v3 (B.11.31)	HP 9000 servers, HP Integrity servers, and servers combining HP 9000 partitions and HP Integrity partitions

HP-UX WLM provides goal-based workload management. This management enables automatic resource allocation and application performance management through the use of prioritized service-level objectives (SLOs). It provides this functionality by automating features of HP Process Resource Manager (PRM), HP-UX Virtual Partitions, nPartitions, and processor sets.

HP-UX WLM offers the following features:

- Automatic allocation of CPU resources (cores) to your most critical applications (allocation of time slices on cores as well as whole cores when managing processor sets, virtual partitions, or nPartitions)

A core is the actual data-processing engine within a processor, where a single processor might have multiple cores, and a core might support multiple execution threads; for more information, see the *WLM HP-UX Workload Manager User's Guide*.

- User-configured allocation of real memory and disk bandwidth resources
- Shares-based and goal-based SLOs

- CPU shares-per-metric allocations, which allow you to give a workload, for example, 2% of the CPU resources for each process in the workload
- Multiple SLOs at different priorities, allowing you to have a “must meet” goal and optional stretch goals for each workload
- Passive mode for analyzing a configuration’s behavior
- Audit and billing data
- Ability to set minimum and maximum amounts of cores available to a workload
- Time-based and metric-based activation of SLOs, including the ability to activate SLOs associated with a Serviceguard package on failover
- GUI for creating, modifying, and deploying WLM configurations locally and on remote systems
- Integration with HP Systems Insight Manager, allowing you to remotely control WLM running on numerous nodes from a single point of administration

---

## **New in this version**

This section lists the new or changed functionality for WLM A.03.02 and WLM A.03.02.02. WLM A.03.02 supports HP-UX 11i v1 (B.11.11) and HP-UX 11i v2 (B.11.23). WLM A.03.02.02 supports HP-UX 11i v3 (B.11.31).

This document refers to both versions of WLM collectively as A.03.02.xx.

- WLM A.03.02.02 supports HP-UX 11i v3 (B.11.31).
- WLM A.03.02.02 supports the logical CPU (Hyper-Threading) feature, which is available starting with HP-UX 11i v3 (B.11.31) for processors designed to support the feature and that have the appropriate firmware installed. A logical CPU is an execution thread contained within a core. Each core with Hyper-Threading enabled can contain multiple logical CPUs. WLM supports the Hyper-Threading feature for PSET-based groups. WLM automatically sets the Hyper-Threading state for the default PSET to optimize performance. (The default PSET, also known as PSET 0, is where all FSS groups reside.) When new PSETs are created, they inherit the Hyper-Threading state that the system had before WLM was activated (inheritance is based on the system state prior to WLM activation because WLM may change the Hyper-Threading setting for the default PSET to optimize performance). Cores can be moved from one partition to another and will take on the Hyper-Threading state of their destination PSET. You can override the default

**New in this version**

Hyper-Threading state of cores assigned to a specific PSET group; you can also modify the Hyper-Threading state of the system. (Modifications to the Hyper-Threading state should not be made while WLM is running.) For more information, see `wlmconf(4)` or the *HP-UX Workload Manager User's Guide*.

- When referring to hardware, CPUs are now referred to as cores in the documentation and in WLM displays or data reports. A core is the actual data-processing engine within a processor. A processor might have multiple cores, and a core might support multiple execution threads through Hyper-Threading. The term “CPU” is still used when referring to concepts such as “CPU resources” or “CPU utilization”; whenever the number of physical processing devices is reported by WLM or discussed by the documentation, the number will be stated explicitly in terms of cores. Note also that PSET sizes and Instant Capacity (iCAP) partitions are now expressed in terms of cores.

---

**NOTE** With Hyper-Threading disabled, each core is seen as a CPU. With Hyper-Threading enabled, each core can be seen as multiple, logical CPUs.

---

- The `wlminfo par` and `wlminfo host` commands now explicitly display core statistics, such as in the following display for the `wlminfo par` command:

```
Hostname      Intended Cores  Cores  Cores Used  Interval
```

- The `wlminfo group` command now displays memory utilization of all groups in the current deployed configuration. In addition, this command now supports the `-v` option to display each group's `gmincpu`, `gmaxcpu`, `gminmem`, and `gmaxmem` values, if they are available in the current deployed configuration. This new option is ignored if live data is not being displayed (for example, when the `-o` option is being used). If memory management is not being used, a dash (-) instead of a zero is displayed in the ‘Mem Shares’ column. If a group's `gmincpu`, `gmaxcpu`, `gminmem`, or `gmaxmem` value is not assigned in the current configuration, a dash (-) is displayed in the corresponding column.

For more information, see `wlminfo(1M)` or the *HP-UX Workload Manager User's Guide*.

- WLM supports the use of Extended Regular Expressions (EREs) for defining alternate names for application records. This support requires that PRM C.03.02 or later is running on the same system. For more information, see `wlmconf(4)` or the *HP-UX Workload Manager User's Guide*.
- WLM supports placement of processes and assignment of user access based on Unix groups. You can assign Unix groups to workload groups by defining the `uxgrp` record in the `prm` structure. Processes whose effective group ID (GID) matches a Unix group record

will run in the associated workload group. WLM supports Unix group records only if PRM C.03.02 or later is running on the system. For more information, see `wlmconf(4)` or the *HP-UX Workload Manager User's Guide*.

- Temporary Instant Capacity (TiCAP) activates capacity in a temporary “calling-card fashion,” such as in 30-day increments (where a day equals 24 hours for one core). With Temporary Instant Capacity on the system, any number of Instant Capacity cores can be activated as long as your prepaid temporary capacity time has not expired. By default, if fewer than 15 processing days are available, WLM stops activating Temporary Instant Capacity resources. Beginning with this release of WLM, you can change this default by setting the WLM global arbiter `utility_reserve_threshold` keyword.
- The Pay per use Toolkit (PPUTK) and the `utilitydc` command are no longer supported and have been removed from the product. Please use the simpler and more robust Temporary Instant Capacity/Pay per use solution available with `wlmpard`.
- The WLM installation script no longer detects whether the correct version of the Java™ Runtime Environment (JRE) is running or whether the correct version of PRM is running. To run the WLM GUI (`wlmgui`) and the wizard (`wlmcw`) requires JRE version 1.4.2 or later and, for PRM-based configurations, PRM C.03.00 or later (to take advantage of the latest updates to WLM, use the latest version of PRM (C.03.02 or later).

---

## Known problems and workarounds

This section discusses problems and workarounds.

### System panic when PRM is enabled; install failure in absence of PRM when certain kernel patches are present

Issue	On HP-UX 11i v1, a system panic problem occurs when Process Resource Manager (PRM) is enabled. In addition, WLM installation fails when certain core kernel patches are present and PRM is absent.
Workaround	Install the version A.01.00.00.07 or later PROCSETS product bundle, which includes critical core kernel and processor set patches. It also installs patches PHKL_30032 through PHKL_30036 or their superseding patches. For patch descriptions, refer to Table 1-4 on page 31.

**Known problems and workarounds****Capping issue**

**Issue** WLM maintains CPU allocations for workloads by capping their CPU access. Unfortunately, an algorithm in the CPU scheduler that WLM uses does not always preserve capping.

You may see symptoms of this issue in `wlminfo` output as in the portion of sample output shown below, with the `g_nice` group having its “CPU Util” value significantly higher than the “CPU shares” value:

```
# wlminfo group

Workload Group  PRMID  CPU Shares  CPU Util  Mem Shares  State
OTHERS          1      450.00     4.49      0.00      ON
g_nice          2      108.00    125.57     0.00      ON
g_nightly       3         0.00     0.00     0.00     OFF
g_team          4         6.00     0.00     0.00     ON
```

For a given group, if its “CPU Util” value is ever significantly greater than its “CPU Shares” value, your system is affected by this issue. (“CPU Util” values slightly above “CPU Shares” are normal.)

**Workaround** On HP-UX 11i v1 (B.11.11), install patches PHKL\_30034, PHKL\_30035, PHKL\_31993, PHKL\_31995, and PHKL\_32061. Install all these patches at the same time.

On HP-UX 11i v2 (B.11.23), install the BUNDLE11i patch bundle. Any version of this bundle is acceptable.

**WLM uses only the assigned CPU resources even with utilitypri set**

**Issue** In an Instant Capacity (iCAP) environment, with `utilitypri` set in your WLM global arbiter configuration, WLM ensures all your owned cores are active. However, if Instant Capacity is not configured in the environment (no designated Instant Capacity cores), WLM uses only the cores that were assigned to virtual partitions when the WLM global arbiter (`wlmpard`) was started.

**Workaround** Be sure to assign all the owned cores using `vparmodify` before you start `wlmpard`. If `wlmpard` is already running, stop it (with the `-k` option) and assign all the owned cores using `vparmodify`.

## Temporary Instant Capacity (TiCAP) expires while WLM is managing nPartitions

**Issue** WLM manages nPartitions using its wlm<sub>p</sub>ard daemon. Assume wlm<sub>p</sub>ard is started on a system that has Temporary Instant Capacity in use. If that temporary capacity expires, wlm<sub>p</sub>ard will still be able to deactivate cores without any problems. However, wlm<sub>p</sub>ard may attempt to activate cores based on the expired capacity. These attempts will fail because the temporary capacity no longer exists. wlm<sub>p</sub>ard will not abort, but it may continue to attempt to activate unavailable cores, generating a message of the following form in /var/opt/wlm/msglog:

```
Error increasing core count on partition par_name (has x needs y) .
```

You will also see the message:

```
Unable to set the local partition to z cores. Check the partition status.
```

where *x*, *y*, and *z* represent integer values.

**Workaround** Add a utility<sub>p</sub>ri statement to your wlm<sub>p</sub>ard configuration, say *configuration\_file*, and then load the new file:

```
# /opt/wlm/bin/wlmpard -a configuration_file
```

The utility<sub>p</sub>ri keyword allows WLM—when Temporary Instant Capacity is available—to adjust the total cores to meet demand.

Specifying this priority ensures WLM maintains compliance with your Temporary Instant Capacity usage rights. When your prepaid amount of temporary capacity expires, WLM no longer attempts to use the temporary resources.

---

**NOTE** Beginning with WLM A.03.02, you can set a threshold that determines when WLM will stop allocating temporary capacity resources. Prior to WLM A.03.02, the threshold was fixed at 15 processing days (where WLM stops allocating temporary capacity if 15 or fewer processing days of temporary capacity remain available). For more information, see “WLM Temporary Instant Capacity 15-day threshold too limiting” on page 42 and wlm<sub>p</sub>ardconf(4).

---

## Known problems and workarounds

### **Automatic activation of Instant Capacity core without authorization**

- Issue** An Instant Capacity (iCAP) core was automatically activated without customer authorization.
- Workaround** Please contact your HP representative.
- If you have Instant Capacity or Pay per use (PPU) software installed, either:
- Do not use WLM virtual partition management, or
  - Use vPars version A.03.01 or later

### **Application hangs in FSS group**

- Issue** On HP-UX 11i v2 (B.11.23), an application inside a workload group based on an FSS group may hang when running in a single-processor virtual partition, nPartition, or system.
- Workaround** Install patch PHKL\_33052.

### **Shutdown slow; “Waiting for shutdown confirmation” and “Shutdown initiated; however, ... unable to acquire confirmation” messages displayed**

- Issue** In some situations, WLM might take longer than expected to shut down, especially when WLM is in the process of modifying the distribution of CPU resources among partitions. In such cases, the shutdown request will not be honored until the modifications are completed. Beginning with WLM A.03.02, you might see the following message after 30 seconds:
- ```
Waiting for shutdown confirmation...
```
- Then, if no shutdown confirmation is received within the next 90 seconds, WLM will display the following message:
- ```
Shutdown initiated; however, we were unable to acquire confirmation. Check the messages in /var/opt/wlm/msglog for more details.
```
- In versions of WLM prior to A.03.02, under similar circumstances you would get one of the following messages after 30 seconds:
- ```
wlmd -k failed: Resource temporarily unavailable  
wlmpard -k failed: Resource temporarily unavailable
```

```
wlmcmd -k failed: Resource temporarily unavailable
```

These messages are misleading in that they imply that the shutdown request had failed when it might not have. These messages have been replaced by the more accurate messages reported above.

**Workaround** If a shutdown request has been delayed for 30 seconds, WLM issues the “Waiting for shutdown confirmation” message. WLM is likely delaying the shutdown request while waiting for partition modifications to complete. If after a total of 120 seconds the shutdown has still not completed, WLM issues the “Shutdown initiated” message. This most likely means that the partition modifications have not yet completed. They can take longer than 120 seconds. When the modifications have been made, WLM will honor the shutdown request. You can verify the shutdown has succeeded by using the `ps` command (if necessary, issue your shutdown command again). In addition, check the messages in `/var/opt/wlm/msglog`.

## Unable to get CPU allocation due to number of processes

**Issue** WLM provides a workload group its allocation on a system by granting the group its allocation on each core. If the group does not have at least one process for each core, WLM increases the allocations for the processes to compensate. For example, for a workload group with a single-threaded process, 10% of four cores is allocated as 40% of one core.

Assume this same group were allocated 50% of the four cores. WLM would allocate 100% of two cores to the workload group. However, because the group has only the one thread, it can use only one core, resulting in an allocation of 25%.

**Workaround** There is no workaround. However, be aware of how your applications run so that you do not give them resource allocations they cannot use.

## Collectors abort when updated while running

**Issue** If you update (overwrite) a data collector executable (be it a binary or a script) while it is providing data to WLM, the collector may abort.

**Workaround** There are two workarounds to this issue:

- Update the data collector in place
  1. Stop WLM (`wlmd -k`)
  2. Update the data collector

## Known problems and workarounds

3. Re-start WLM (`wlmd -a configuration_file`)
- Replace the data collector
  1. Move the current data collector aside
  2. Install the new data collector in place of the collector you just moved aside
  3. Re-start WLM (`wlmd -a configuration_file`)

## GlancePlus/OpenView Performance Agent and processor sets

**Issue** On systems with multiple processor sets configured, GlancePlus may have incorrect data for the `PRM_SYS` group (ID 0). On these systems, GlancePlus will incorrectly include processes that are outside of the default processor set as belonging to the `PRM_SYS` group (ID 0).

As a result, the WLM `glance_prm` data collection script cannot be used to track application (APP or APP\_PRM) metrics for the `PRM_SYS` group or any PRM group defined based on a PSET.

Only GlancePlus is affected by this issue—WLM properly identifies workload groups (PRM groups and their PRM IDs).

**Workaround** For HP-UX 11i v1 (B.11.11), using GlancePlus C.03.55 or later and installing patch PHKL\_28052 addresses this issue.  
For HP-UX 11i v2 (B.11.23), this issue is fixed in GlancePlus C.03.58.05.

## GlancePlus may not correctly identify processes' PRM groups

**Issue** On some systems, GlancePlus would not correctly identify processes' PRM groups. WLM uses these PRM groups as its workload groups. On these systems, GlancePlus would report all processes as belonging to the `PRM_SYS` group (ID 0). As a result, the WLM `glance_prm` data collection script could not be used to track application (APP) metrics for a PRM group of processes.

**Workaround** This issue is fixed in GlancePlus C.03.35.00. On HP-UX 11i v1, the best way to get this upgrade is to install GlancePlus when installing the 11i Enterprise or 11i Mission-critical Operating Environments.

## **glance Adviser memory consumption increases continually**

- Issue**                    GlancePlus's glance Adviser leaks memory when running continuously. Adviser is used by the WLM data collectors `glance_app`, `glance_gbl`, `glance_prm`, `glance_prm_byvg`, and `glance_tt`.
- Workaround**            This memory leak is fixed in GlancePlus C.03.35.00. On HP-UX 11i v1, the best way to get this upgrade is to install GlancePlus when installing the 11i Enterprise or 11i Mission-critical Operating Environments.

## **WLM enables/disables SLOs at end of interval**

- Issue**                    WLM enables and disables time-based SLOs only at the end of an interval. This interval is 60 seconds by default and can be changed with the `wlm_interval` keyword in your configuration file.
- SLOs are time-based when you set their `condition` or `exception` keyword values in the configuration file.
- If your interval is too long, an SLO may not be enabled as indicated in the configuration. For example, assume the interval is 1800 seconds (30 minutes). Also assume that one SLO is supposed to have the entire system to itself for a short period of time, with WLM enabling that one SLO for 20 minutes while disabling all other SLOs for the same 20 minutes. If these 20 minutes begin at 3pm everyday, but the interval ends at 3:15pm, your configuration does not actually go into effect until 3:15pm. Moreover, it is not changed again until 3:45pm.
- Workaround**            Be aware of how your interval and time-based SLOs interact and adjust them accordingly.

## **No metrics on startup or reconfiguration**

- Issue**                    Metrics have no value on WLM startup or reconfiguration. WLM cannot work toward any SLOs without metrics.
- Workaround**            Data collectors should report metrics as soon after startup or reconfiguration as possible.

## WLM configurations cannot be activated with fewer than 100 Mbytes of memory available

- Issue** When controlling memory, WLM allocates at least a minimum amount to each group. If `extended_shares` is enabled, this minimum is 0.2% of the available memory; otherwise, it is 1%. (Available memory is the amount reported by `prmvavail`; it is the amount that is not reserved for the kernel (`/stand/vmunix`) and its data structures and for nonkernel system processes. Thus, available memory is not the total memory on the system. Available memory varies over time. For more information, see the *Process Resource Manager User's Guide*, available in `/opt/prm/doc`.) WLM requires that this minimum represent at least 1 Mbyte of memory. Thus, when memory control is used, the system should have at least 100 Mbytes of available memory, or at least 500 Mbytes if `extended_shares` is enabled.
- Workaround** Increase the system's memory or decrease the minimum dynamic buffer cache using the kernel parameter `dbc_min_pct`.

## Secure Resource Partitions: Blocked port on a virtual network interface

- Issue** Using the HP-UX feature Security Containment, you can set up a virtual network interface for each secure compartment. A process in one secure compartment can bind to a socket on a virtual network interface associated with a different secure compartment. Although this process will not be able to accept connections or use the socket to send or receive data, it does prevent other processes from binding to that socket.
- Workaround** Be sure your applications that access the network bind only to sockets on the virtual network interface created for their respective secure compartments. For more information, see the following documents:
- For more information, refer to the Security Containment release notes and the *HP-UX 11i Security Containment Administrator's Guide*.

## Reaching the system V semaphore limit

- Issue** If your system has many system V semaphores in use, WLM usage of semaphores may push the total number of semaphores over system limits—especially if the WLM configuration contains a large number of data collectors. The WLM daemon `wlmd` creates two semaphore sets: one containing a single semaphore; the other containing a semaphore for each data collector in the WLM configuration.

If this system limit is reached, `wlmd` prints the following error and exits before the first WLM interval begins:

```
Cannot allocate a system V semaphore set of size x: Increase  
the system-imposed limits.
```

You may also see the following error:

```
WLM--"wm_knob_init, prm_rep_load": PRM--"PRM internal daemon  
binary is missing or has incorrect permissions (PRM-2352)"
```

For additional details relating to this message, look in `syslog`. In some conditions, this message indicates no semaphores are available.

**Workaround** Use the SAM (`/usr/sbin/sam`), SMH (`/usr/sbin/smh`), or `kcweb` (`/usr/sbin/kcweb`, on HP-UX 11i v2 or later) kernel configuration utility to increase the system limits.

On HP-UX 11i v1 and later, increase the following kernel parameters:

|                     |                                                     |
|---------------------|-----------------------------------------------------|
| <code>semnms</code> | Max number of overall semaphores                    |
| <code>semmsl</code> | Max number of semaphores allowed in a semaphore set |

## Configuration wizard requires PRM

**Issue** Starting with the WLM A.03.01 release, WLM no longer includes Process Resource Manager (B3835DA). The WLM configuration wizard requires PRM however. Without PRM installed, the wizard:

- Always sets the initial CPU (core) count to 1 in the pop-up dialog that appears before the wizard itself appears
- Returns a message about the PRM API not being installed when it attempts to validate a configuration

**Workaround** Install PRM A.03.00 or later.

## Processes in transient FSS groups appear unexpectedly in other workload groups

**Issue** A deployed WLM configuration that has `transient_groups` set to 1 and contains FSS transient group candidates from time-to-time contains an FSS group called `_IDLE_`. As needed, WLM moves the jobs of the transient groups to `_IDLE_`, where they get the minimum of CPU and memory resources. The internal identifier for group `_IDLE_` is picked by WLM on the fly (taken from the pool of unused identifiers).

On a redeployment, if the new configuration contains an FSS group that happens to have the identifier WLM selected for the `_IDLE_` group in the deployment that is being replaced, jobs in the `_IDLE_` group will migrate to the FSS group having the same identifier.

The same issue can arise going from:

- A transient deployment to another transient deployment
- A transient deployment to a non-transient deployment

**Workaround** You can prevent this situation by using the `-i` option to `wlmd` when deploying or redeploying a configuration. First, shut down `wlmd` (using the `-k` option), then restart it using the `-i` option.

## Before modifying any partition managed by WLM, WLM and the global arbiter must be stopped

**Issue** Do not adjust any WLM-managed partition while `wlmpard` is running. This includes using `vparmodify`, `icapmodify`, or `icod_modify` to change the name, configuration, or resources (CPU and memory) associated with the virtual partition or `nPartition` (and this also includes using `parolrad` to modify a cell in a WLM-managed partition, as noted in “Before performing online cell operations on systems where WLM manages partitions, memory, or PSETs, WLM must be stopped” on page 19).

**Workaround** To adjust a partition, you must first shut down WLM—including `wlmpard`—on all partitions that will be affected by the modification, modify the partition, and then restart WLM. Changes to Instant Capacity (iCAP) affect the entire complex; changes to a virtual partition affect the `nPartition` only, unless Instant Capacity is configured on the `nPartition`. For example, if WLM is managing two virtual partitions `vParA` and `vParB`, and you need to migrate memory resources from `vParA` to `vParB`, you must shut down WLM in both virtual partitions. As another example, to change the name of an

nPartition, you must first shut down WLM in every operating system instance across the entire complex, because the name change affects Instant Capacity, and Instant Capacity changes affect every nPartition across the complex.

To stop WLM, stop the `wlmpard` and `wlmd` daemons (use the `-k` option with the corresponding commands).

## Before performing online cell operations on systems where WLM manages partitions, memory, or PSETs, WLM must be stopped

|            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Issue      | If WLM is being used to manage memory records, partitions, or PSET-based workload groups, and you attempt to perform an online cell operation ( <code>parolrad</code> ) while WLM is running, changes made to CPU resources by the operation might not be detected by WLM and can cause problems for WLM management of CPU resources. Error messages will be generated.                                                                                                                                                                         |
| Workaround | <p>Before performing an online cell operation (<code>parolrad</code>) on a system where WLM is managing memory, partitions, or PSETs, you must first stop WLM, perform the operation, and then restart WLM. To stop WLM, stop the <code>wlmpard</code> and <code>wlmd</code> daemons (use the <code>-k</code> option with the corresponding commands). Note that <code>wlmd</code> should be stopped on all partitions managed by WLM.</p> <p>You can check the status of online cell operations by using <code>parolrad -m</code> command.</p> |

## WLM GUI is not compatible with different versions of WLM

|            |                                                                                                                                                                                                                                                                                                                                                                        |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Issue      | <p>If you attempt to use the WLM GUI (<code>wlmgui</code>) with a version of WLM that differs from the version the GUI is associated with, the following message is displayed:</p> <pre>The WLM product running on &lt;hostname&gt; and this tool are incompatible.</pre> <p>The version of the WLM GUI must match the version of the WLM product that it manages.</p> |
| Workaround | Either upgrade WLM to the version of the WLM GUI you are using, or use an earlier version of the WLM GUI that matches the version of WLM that your GUI will manage. Note that multiple versions of the WLM GUI can be installed on a Microsoft Windows PC.                                                                                                             |

## Upgrading or installing PRM before upgrading WLM from C.03.00 or earlier can cause WLM to fail swverify checks

|            |                                                                                                                                                                                      |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Issue      | If you install or upgrade to the latest version of PRM (C.03.02 or later) on a system with WLM A.03.00 or earlier, WLM will fail swverify checks.                                    |
| Workaround | To ensure that WLM works properly on a system with PRM C.03.02 or later, upgrade WLM to A.03.02 or later. When upgrading WLM A.03.00 or earlier, upgrade WLM prior to upgrading PRM. |

---

## Compatibility information and installation requirements

HP-UX WLM A.03.02 (B8843CA) is available for the following environments:

- HP-UX 11i v1 (B.11.11) on HP 9000 servers
- HP-UX 11i v2 (B.11.23) on HP 9000 servers or HP Integrity servers
- HP-UX 11i v1 (B.11.11) and HP-UX 11i v2 (B.11.23) on servers combining HP 9000 partitions and HP Integrity partitions (in such mixed environments, HP-UX 11i v1 supports HP 9000 partitions only)

HP-UX WLM A.03.02.02 (B8843CA) is available for the following environments:

- HP-UX 11i v3 (B.11.31) on HP 9000 servers, HP Integrity servers, and servers combining HP 9000 partitions and HP Integrity partitions

The HP-UX WLM product includes media, manual, and release notes.

Also available: *HP-UX Workload Manager User's Guide* (B8844-90014)

## Disk and memory requirements

WLM A.03.02.xx requires:

- 16 Mbytes of disk space in /opt/wlm/
- 9 Mbytes of disk space in /opt/prm/
- 5 Mbytes of memory (Also, your system must have at least 100 Mbytes of available memory for WLM to properly manage the workloads' memory usage. If `extended_shares` is enabled, WLM memory control requires at least 512 MB of available memory. Available memory is the amount reported by `prmvavail`; it is the amount that is not reserved for the

kernel (/stand/vmunix) and its data structures and for nonkernel system processes. Thus, available memory is not the total memory on the system. Available memory varies over time. For more information, see the *Process Resource Manager User's Guide*, available in /opt/prm/doc.)

## Network operating environment

WLM network interfaces are designed to operate correctly and defend against attacks in a moderate to high threat environment, such as a DMZ. You may use network protections, such as firewalls, to provide an additional level of defense and to give you additional time to react in the event that a security loophole is found.

---

**NOTE** As of A.03.01, WLM enables secure communications by default when you start WLM using the /sbin/init.d/wlm script. You should ensure that the secure mode variables are enabled in /etc/rc.config.d/wlm. You also must distribute security certificates to all systems or partitions being managed by the same WLM global arbiter (wlmprad). For more information on using security certificates and other tasks necessary to enable secure communications, see wlmcert(1M).

---

The WLM wlmprad and wlmcomd daemons use the following port numbers by default:

|         |      |
|---------|------|
| wlmprad | 9691 |
| wlmcomd | 9692 |

Make sure these ports are kept open. To change these port numbers, see wlmprad(1M) and wlmcomd(1M).

## Compatibility with other software

WLM A.03.02.xx was verified using the software listed in the following table. Compatibility with other versions of the software packages is unknown unless explicitly stated.

**Table 1-1** Verified software configurations

| Software package | Version on HP-UX 11i v1 (B.11.11) | Version on HP-UX 11i v2 Update 2 (B.11.23) | Version on HP-UX 11i v3 (B.11.31) |
|------------------|-----------------------------------|--------------------------------------------|-----------------------------------|
| PRM              | C.03.00,<br>C.03.01,<br>C.03.02   | C.03.00,<br>C.03.01,<br>C.03.02            | C.03.02.xx                        |

**Compatibility information and installation requirements****Table 1-1 Verified software configurations (Continued)**

| <b>Software package</b>                                  | <b>Version on HP-UX 11i v1 (B.11.11)</b>                                             | <b>Version on HP-UX 11i v2 Update 2 (B.11.23)</b>    | <b>Version on HP-UX 11i v3 (B.11.31)</b> |
|----------------------------------------------------------|--------------------------------------------------------------------------------------|------------------------------------------------------|------------------------------------------|
| Processor Sets                                           | A.01.00.00.07                                                                        | The PSET version included with HP-UX                 | The PSET version included with HP-UX     |
| Serviceguard                                             | A.11.15 (HP recommends using A.11.17 or later)                                       | A.11.17                                              | A.11.17                                  |
| Servicecontrol Manager; HP Systems Insight Manager       | C.05.00                                                                              | C.05.00                                              | C.05.00                                  |
| Instant Capacity (iCAP) for use in nPartition management | B.05.00, B.11.11.06.50, B.11.11.07.01 (this can be used with vPars A.03.x and later) | B.06.00, B.11.23.06.03, B.11.23.07.00, B.11.23.08.00 | B.11.31.08.01                            |
| Virtual partitions (vPars)                               | A.02.02, A.03.03                                                                     | A.04.01, A.04.02, A.04.03                            | A.05.01                                  |
| Pay per use (PPU)                                        | B.07.00                                                                              | B.07.00                                              | B.08.00                                  |
| EMS                                                      | A.04.00.02, A.04.20                                                                  | A.04.10, A.04.11, A.04.20                            | A.04.20                                  |
| GlancePlus Pak                                           | C.03.86, C.04.50                                                                     | C.03.85, C.03.86, C.04.50                            | C.04.70                                  |
| WLM Toolkits                                             | A.01.10                                                                              | A.01.10                                              | A.01.10.xx                               |
| OpenSSL                                                  | A.00.09.07-d or later                                                                | A.00.09.07-d or later                                | A.00.09.07-d or later                    |

**Table 1-1 Verified software configurations (Continued)**

| <b>Software package</b> | <b>Version on HP-UX 11i v1 (B.11.11)</b> | <b>Version on HP-UX 11i v2 Update 2 (B.11.23)</b> | <b>Version on HP-UX 11i v3 (B.11.31)</b> |
|-------------------------|------------------------------------------|---------------------------------------------------|------------------------------------------|
| HP-UX Bastille          | B.02.01                                  | B.02.01                                           | B.03.00                                  |

### Compatibility with long hostnames

WLM supports long host names. You must install the bundle NodeHostNameXpnd and configure the feature first though. For more information, refer to the section “Installation procedure” on page 26.

### Compatibility with X Windows

The WLM configuration wizard (/opt/wlm/bin/wlmcw) and the WLM monitoring GUI (invoked with wlminfo -i or wlmgui) were verified with the following X Windows variants:

- X Windows on HP-UX
- XFree86, Version 4.1.0-15 (Linux)
- Reflection X for Windows 2000

HP makes no claims with respect to other variants.

### Compatibility with GlancePlus

You can use GlancePlus to monitor PRM, which WLM uses. GlancePlus also provides metrics that you can use to drive SLOs.

WLM is compatible with any HP-UX 11i version of GlancePlus. On HP-UX 11i v1, be sure to install the 11i Enterprise Operating Environment or 11i Mission-critical Operating Environment to ensure you have a GlancePlus version that is fully compatible with WLM.

---

**NOTE** Do not use GlancePlus to change PRM allocations. WLM controls PRM.

---

## **Compatibility of WLM virtual partition management and Instant Capacity / PPU**

If you have Instant Capacity (iCAP) or Pay per use (PPU) software installed, use WLM virtual partitions management only if you have vPars version A.03.01 or later.

If you have a vPars version prior to A.03.01, using WLM virtual partition management may cause an Instant Capacity core to be automatically enabled without customer authorization. If this situation occurs, please contact your HP representative.

With vPars version A.04.01 or later, use Instant Capacity v7 or later.

## **Compatibility of WLM virtual partition management and certain CPU bindings**

Do not use cell-specific CPU bindings or user-assigned CPU bindings on virtual partitions you are going to manage with WLM.

## **Compatibility of WLM with HP Integrity Virtual Machines**

WLM supports HP Integrity Virtual Machines (Integrity VM). You can run WLM both on the Integrity VM Host and in an Integrity VM (guest), but each WLM runs as an independent instance. To run WLM on the Integrity VM Host, you must use a strictly host-based configuration—a WLM configuration designed exclusively for moving cores across HP-UX Virtual Partitions or nPartitions, or for activating Temporary Instant Capacity (TiCAP) cores or Pay per use (PPU) cores. (WLM will not run with FSS groups or PSETs on Integrity VM Hosts where guests are running.) In addition, ensure that the minimum number of cores allocated to a WLM host is greater than or equal to the maximum number of virtual CPUs (vCPU count) assigned to each VM guest. Otherwise, VM guests with a vCPU count greater than or equal to WLM's minimum allocation could receive insufficient resources and eventually crash. For example, if an Integrity VM host has 8 cores and three guests with 1, 2, and 4 virtual CPUs, respectively, your WLM host should maintain an allocation of at least 4 cores at all times. You can achieve this by using the WLM `hmincpu` keyword.

WLM runs inside an Integrity VM but will not support PPU, vPar, and Instant Capacity (iCAP) integration. However, Integrity VM will take advantage of cores added to the Integrity VM Host by PPU, Instant Capacity, and TiCAP. As noted previously, WLM must continue allocating at least as many cores as the maximum number of virtual CPUs in any VM guest on the system. In addition, when running WLM inside an Integrity VM, you should specify a WLM interval greater than 60 seconds. This helps ensure a fair allocation of CPU resources for FSS groups.

For more information on Integrity VM, go to the following location and navigate to the “Solution components” page:

[www.hp.com/go/vse](http://www.hp.com/go/vse)

## **Compatibility of WLM partition management and PSETs**

WLM now supports simultaneous management of partitions (virtual partitions or nPartitions) and PSET-based workload groups. Such support requires the following:

- If Instant Capacity (iCAP) is available on the complex, it must be v7 or later
- If HP-UX Virtual Partitions (vPars) is on the complex, it must be v4 (A.04.01) or later
- If Pay per use (PPU) is on the complex, it must be v7 or later

## **Compatibility of psrset and PSETs**

When WLM is managing PSETs, do not use the `psrset` command to change settings.

## Compatibility with PRM

As of version A.03.01, WLM no longer includes PRM in the bundle. As a result, upgrading WLM no longer triggers a reboot.

WLM will run only with PRM C.03.00 or later. If PRM C.03.00 or later is already on the machine where you are installing or upgrading WLM, you can continue to manage FSS and PSET-based workload groups just as if PRM had been installed with WLM.

If you are installing WLM for the first time on a machine, you can only use a strictly host-based configuration—a configuration that does not include a `prm` structure and is designed exclusively for moving cores across HP-UX Virtual Partitions or `nPartitions`, or for activating Temporary Instant Capacity (TiCAP) or Pay per use (PPU) cores. To manage FSS and PSET-based workload groups, you must install PRM separately. With your purchase of WLM, you are entitled to a PRM license. For information on receiving the PRM license, contact your regional licensing service organization.

When the Servicecontrol Manager (SCM) bundle is built during the WLM installation, it will include the PRM API only if PRM is already installed. The SCM integration bundle will not be able to validate PRM-based configurations without the PRM API.

---

**NOTE** If you plan to use the WLM configuration wizard, install the appropriate version of PRM. To take advantage of the latest updates to WLM, use the latest version of PRM (C.03.02 or later).

To ensure that WLM works properly on a system with PRM C.03.02 or later, upgrade WLM to A.03.02 or later. When upgrading WLM A.03.00 or earlier, upgrade WLM prior to upgrading PRM.

---

You can use WLM to control resources managed by PRM. WLM uses PRM when a `prm` structure is included in the WLM configuration. With such configurations, you can use PRM's informational and monitoring commands such as `prmlist` and `prmmonitor`. You can also use the `prmrun` and `prmmove` commands, among others. If you use the `prmconfig` command, invoke it with no options or the `-u` (unlock) option—do not use the `-r` (reset) option.

Ordinarily, WLM and PRM should not be used to manage resources on the same system at the same time. Using both WLM and PRM to control resources at the same time might cause inconsistent behavior and undesirable performance. However, you can use both products at the same time if the PRM configuration uses FSS groups only (no PSET-based groups) and the WLM configuration is strictly host-based. (A strictly host-based configuration is one that does not include a `prm` structure; it is designed exclusively for moving cores across HP-UX Virtual Partitions or `nPartitions`, or for activating Temporary Instant Capacity (TiCAP) cores or Pay per use (PPU) cores.) You might want to use both products to take advantage of certain features of PRM that are not included with the latest release of WLM, such as PRM's

CPUCAPOFF mode, enabled with the `prmconfig -M CPUCAPOFF` command. (In this mode, a PRM group's upper bound for CPU resource consumption is determined by the CAP value, available on HP-UX 11i v3 and later. For more information, see the *HP Process Resource Manager User's Guide* or `prmconfig(1M)`.)

## Compatibility with gWLM

WLM and HP Global Workload Manager (gWLM) should not be used to manage the same system at the same time. Using these products to control resources at the same time may cause inconsistent behavior and undesirable performance.

## Compatibility with Java

To run the WLM GUI (`wlmgui`) or wizard (`wlmcw`), Java Runtime Environment version 1.4.2 or later is required. For PRM-based configurations, these utilities require the PRM C.03.00 or later. To take advantage of the latest updates to WLM, use the latest version of PRM (C.03.02 or later).

## Installation procedure

If you are installing WLM from the quarterly Application Release (AR) CD-ROM, use the procedure in this section to minimize system reboots.

To install your software, use the SD-UX `swinstall` command. For information on installing using a GUI, refer to “Using the `swinstall` Graphical User Interface (GUI)” on page 41.

---

**NOTE**      If you are upgrading your version of WLM, after installation, you must explicitly restart HP-UX WLM with the `wlmd -a` or `wlmd -A` commands.

---

The following steps give an overview of the HP-UX WLM A.03.02.xx installation process:

- Step 1.** Verify that HP-UX is up-to-date
- Step 2.** (Optional) Configure long hostnames
- Step 3.** Install OpenSSL, if necessary
- Step 4.** Determine what patches are needed
- Step 5.** (Optional) Create a software depot to minimize system reboots
- Step 6.** Disable PRM from starting automatically
- Step 7.** (Optional) Disable PRM's remote configuration daemon

**Compatibility information and installation requirements**

**Step 8.** Shut down WLM if it is already running

**Step 9.** Install patches and install/upgrade WLM

**Step 10.** Check the file `/var/adm/sw/swagent.log`

**Step 11.** (Optional) Ensure you have the proper version of GlancePlus

**Step 12.** (Optional) Install files for integration with HP Systems Insight Manager

**Step 13.** (Optional) Merge new `/etc/rc.config.d/wlm` file with your modified file

These steps are discussed in detail below.

**Step 1.** Verify that HP-UX is up-to-date

HP suggests that as part of your systems' ongoing maintenance and support, the systems be kept up-to-date through the HP IT Resource Center's Support Plus program. This program delivers quarterly patch bundles via the web and as part of the quarterly AR media. To register, go to the following location:

<http://itrc.hp.com>

Select the patch/firmware database link and register to create an account and log in, as instructed. After logging in, scroll down to the "find standard patch bundles" section and select "HP-UX patch bundles". The bundles of interest are "Quality Pack" and "Support Plus".

For HP-UX 11i v1 (B.11.11), the recommended bundles are described in the following table:

**Table 1-2 HP-UX 11i v1 (B.11.11) general support patches**

| <b>HP-UX 11i v1 patch</b> | <b>Description</b>                                                                                                                                                                |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GOLDQPK11i                | Defect-fix patches for HP-UX 11i v1 on HP 9000 workstations and servers<br><br>(When placed in your local depot, this bundle appears with the names GOLDBASE11i and GOLDAPPS11i.) |
| HWEnable11i               | Required hardware enablement patches for all HP-UX 11i v1 workstations and servers                                                                                                |

If you are using HP-UX 11i v1 (B.11.11) and are interested in using processor sets (PSETs), you can download software (to support processor sets) from the following location:

<http://www.hp.com/go/wlm>

Select the Patches/support link and search for "processor sets". This software also corrects several problems (see section "System panic when PRM is enabled; install failure in absence of PRM when certain kernel patches are present" on page 9).

---

**NOTE** Install GOLDQPK11i before installing the processor sets software—unless you are using a depot to install them at the same time.

---

**Compatibility information and installation requirements**

Processor sets are included with HP-UX 11i v2 (B.11.23) and HP-UX 11i v3 (B.11.31).

For HP-UX 11i v2 (B.11.23), install the following patches:]

**Table 1-3 HP-UX 11i v2 (B.11.23) patches for WLM**

| <b>HP-UX 11i v2 patch</b> | <b>Description</b>                                                                                                                                                                                                                                                |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BUNDLE11i patch bundle    | Bundle of patches (Installing this bundle updates your system to HP-UX 11i v2 Update 2.)                                                                                                                                                                          |
| PHKL_35174                | Patch required when using WLM for management of nPartitions or virtual partitions on rx8640 or rx7640 Integrity servers. For more information on this patch, go to the HP IT Resource Center Web site:<br><br><a href="http://itrc.hp.com">http://itrc.hp.com</a> |

- a.** Determine which patch bundles are already on the system

Run the `swlist` command, which will list the bundle if found:

```
# /usr/sbin/swlist -l bundle patch_bundle
```

where `patch_bundle` is GOLDBASE11i, GOLDAPPS11i, or HWEEnable11i

- b.** Download any missing patch bundles from the following location:  
[http://www.hp.com/go/softwareinfo/SUPPORT\\_PLUS/](http://www.hp.com/go/softwareinfo/SUPPORT_PLUS/) to /tmp

---

**NOTE** Do not install these patch bundles at this point.

---

Related support sites include:

- HP Software Depot (<http://www.hp.com/go/softwaredepot>)  
Refer to the “enhancement releases” link.
- HP IT Resource Center (<http://itrc.hp.com>)  
Refer to the “maintenance/support” link.
- HP Technical Documentation (<http://docs.hp.com>)  
Refer to the HP-UX 11i v1 topics “Installing and Updating” and “Patch Management”.

**Step 2.** (Optional) Configure long hostnames

Ensure the NodeHostNameXpnd bundle is installed if you want to use long host names on HP-UX 11i v2 or v3. Configure long host names as described below.

Set the kernel tunable to allow long hostnames:

```
# kctune expanded_node_host_names=1
```

Set the node and hostname, with the changes being persistent across boots:

```
# /sbin/set_parms hostname
```

where `hostname` is the literal string `hostname`—not the desired hostname. You will be prompted for the hostname.

For configuration information, read the paper “Node and Host Name Sizes on HP-UX,” available in `/usr/share/doc/NodeHostNameSize.pdf` on systems with the NodeHostNameXpnd bundle installed.

**Step 3.** Install OpenSSL, if necessary

If you download software from the Web and create a software depot of your own, OpenSSL A.00.09.07-d or later must be installed on the system on which you are installing WLM. You can obtain the OpenSSL software from the HP Software Depot (<http://www.hp.com/go/softwaredepot>). If you are installing WLM from the quarterly AR CD-ROM, the required OpenSSL software is installed with WLM.

**Step 4.** Determine what patches are needed

**a.** Check the following data to determine what patches are suggested for WLM:

If you plan to use the WLM configuration wizard, the `wlmgui` command, or `wlminfo`'s interactive mode GUI (`-i`), all of which depend on Java (Java Runtime Environment version 1.4.2 or later is required), refer to the following location for patch information:

<http://www.hp.com/go/java>

In addition, running these utilities with PRM-based configurations requires PRM C.03.00 or later.

On HP-UX 11i v1, apply patch PHCO\_29109 to address a pthread library issue.

For the latest information on WLM patches, visit the “Patches / support” page available from the following location:

<http://www.hp.com/go/wlm>

## Compatibility information and installation requirements

For HP-UX 11i v1, the following patches are available:

**Table 1-4 HP-UX 11i v1 (B.11.11) patches for WLM**

| HP-UX 11i v1 patch*                                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PHKL_33372                                                             | Memory Resource Groups (MRG) patch (supersedes PHKL_32619)                                                                                                                                                                                                                                                                                                                                                                                    |
| PHKL_30034,<br>PHKL_30035,<br>PHKL_31993,<br>PHKL_31995,<br>PHKL_32061 | <p><b>NOTE:</b> Install all these patches at the same time.</p> <p>These patches:</p> <ul style="list-style-type: none"> <li>• Preserve allocations in cases where a workload group with a large allocation and just enough jobs to get that allocation could be outperformed by a smaller allocation group with more jobs</li> <li>• Improve performance of I/O-based applications when the Fair-Share Scheduler (FSS) is enabled</li> </ul> |

\* These patches might have been superseded. You can use these listed patches or their superseding patches.

For HP-UX 11i v2 (B.11.23), install the following patches:]

**Table 1-5 HP-UX 11i v2 (B.11.23) patches for WLM**

| HP-UX 11i v2 patch*       | Description                                                                                                                                                                                                                    |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BUNDLE11i patch bundle    | Bundle of patches (Installing this bundle updates your system to HP-UX 11i v2 Update 2.)                                                                                                                                       |
| PHKL_32518                | Allows you to create more than 64 workload groups—assuming you have already installed BUNDLE11i.                                                                                                                               |
| PHKL_33052                | Prevents performance degradation when using workload groups that are based on FSS groups in a single-processor virtual partition, nPartition, or system                                                                        |
| PHKL_33604,<br>PHKL_33605 | <p><b>NOTE:</b> Install these patches at the same time.</p> <p>These patches help ensure each job within a FSS-based workload group gets a more equal share of the group's allocation than it would under HP-UX semantics.</p> |

\* These patches might have been superseded. You can use these listed patches or their superseding patches.

**b.** Determine which patches are already on the system

If a recent Quality Pack is installed, some of these patches may already be on your system. To determine if these patches are installed, use `swlist` as follows:

```
# /usr/sbin/swlist [-l patch] patch
```

where *patch* is the patch's tag. For example, to check that patch PHKL\_32061 is installed, use the following command line:

```
# /usr/sbin/swlist -l patch PHKL_32061
```

Similarly, if you downloaded patch bundles in Step 1, check the bundles for the patches:

```
# /usr/sbin/swlist -d patch @ /tmp/bundle_name
```

If a given patch is not listed, check HP's IT Resource Center (<http://itrc.hp.com>) for a superseding patch, then use `swlist` to list the superseding patch.

**c.** Download any missing patches to /tmp.

If any of the required patches are missing (not installed or not in a bundle you downloaded), download the patches from HP's IT Resource Center:

Americas/Asia-Pacific

<http://itrc.hp.com/>

Europe

<http://europe.itrc.hp.com/>

---

**NOTE** Do not install these patches at this point.

---

## Compatibility information and installation requirements

### Step 5. (Optional) Create a software depot to minimize system reboots

Creating a depot is beneficial when you are completing two or more of the following tasks:

- Installing patch bundles (GOLDQPK11i or HWEnable11i)

The patch bundles include kernel patches. Installation of a patch bundle requires a system reboot.

- Installing PHKL\_\* patches

The PHKL\_\* patches are HP-UX kernel patches. Installation of these patches requires a system reboot.

- Installing/upgrading WLM

The installation or upgrade of WLM requires a kernel rebuild and system reboot.

Completing two or more of the above tasks will cause multiple reboots unless you create a software depot. With a depot including all the patches, the patch bundles, and WLM, you can combine the installation steps into a single operation, thereby limiting the number of installation reboots that must occur.

If you do not mind multiple reboots or your system will not be subject to multiple reboots because its software is up-to-date, proceed to Step 6, “Disable PRM from starting automatically” on page 36.

In order to install these items with a single `swinstall(1M)` command, all of the related filesets must be collected into a single software depot distribution directory.

The following steps outline how to create a software depot. The steps use examples showing how to include the following items in your depot:

- Quality Pack GOLDQPK11i (also known as a patch bundle)
- Patch PHKL\_32619
- perl
- WLM

---

**NOTE** To satisfy SD-UX, it may be necessary to copy dependency filesets into this common depot.

---

- a. Ensure that adequate disk space is available. The size of each patch can be found in its README text.

- b. Create a directory for the depot:

```
# mkdir /tmp/wlm_depot
```

- c. Place any patch bundles you downloaded in Step 1 (GOLDQPK11i or HWEnable11i) in /tmp/ if they are not already there.

For more information on these patch bundles, refer to Table 1-2.

- d. Check each of the patch bundles.

```
# /usr/sbin/swlist -d @ bundle_path
```

For example, check GOLDQPK11i as follows:

```
# swlist -d @ /tmp/GOLDQPK11i_11.11.depot
```

The output should include lines similar to the following lines:

```
GOLDAPPS11i    B.11.11.0312.4 Gold Applications Patches for HP-UX  
11i v1, December 2003
```

```
GOLDBASE11i   B.11.11.0312.4 Gold Base Patches for HP-UX 11i v1,  
December 2003
```

If your output is not similar, please download the file again.

- e. Add the patch bundles to the depot:

```
# /usr/sbin/swcopy -s bundle_path bundle_name @ depot_path
```

where *bundle\_path* and *depot\_path* are fully specified path names, and *bundle\_name* is the patch bundle's tag. For example:

```
# swcopy -s /tmp/GOLDQPK11i_11.11.depot GOLDAPPS11i @ /tmp/wlm_depot
```

```
# swcopy -s /tmp/GOLDQPK11i_11.11.depot GOLDBASE11i @ /tmp/wlm_depot
```

This step automatically registers the wlm\_depot. Therefore, there is no need to register it using the swreg command.

- f. Verify the depot:

```
# /usr/sbin/swverify -d \* @ depot_path
```

where *depot\_path* is a fully specified pathname. For example:

```
# swverify -d \* @ /tmp/wlm_depot
```

**Compatibility information and installation requirements**

- g.** Place any missing individual patches that you downloaded in Step 4 in `/tmp/`.

For HP-UX 11i v1 (B.11.11) patches, refer to Table 1-4.

For HP-UX 11i v2 (B.11.23) patches, refer to Table 1-5.

- h.** Add each patch file to the depot:

```
# swcopy -s patch_file_path \* @ depot_path
```

where `patch_file_path` and `depot_path` are fully specified path names. For example:

```
# swcopy -s /tmp/PHKL_32619.depot \* @ /tmp/wlm_depot
```

- i.** Verify the depot:

```
# swverify -d \* @ depot_path
```

where `depot_path` is a fully specified pathname. For example:

```
# swverify -d \* @ /tmp/wlm_depot
```

- j.** Ensure perl is available at `/opt/perl/bin/perl`

If you intend to use the:

- `wlmaudit` utility to display audit reports
- `wlmoradc` or `smooth` utilities from the WLM Oracle Database Toolkit (ODBTK)
- `expsmooth` utility from the WLM BEA WebLogic Server Toolkit (WebLogicTK)
- `wlmcert` utility and secure network communications

and perl is not available at `/opt/perl/bin/perl`:

- If perl is not already installed, add it to your depot from the quarterly Application Release (AR) CD-ROM using the following `swcopy` line:

```
# swcopy -s mount_path perl \* @ /tmp/wlm_depot
```

where `mount_path` is the path to the quarterly AR CD-ROM.

- If perl is installed, but in a location other than `/opt/perl/bin/perl`, ensure there is a symbolic link from `/opt/perl/bin/perl` to your perl.

- k.** Locate and mount the WLM installation media. For assistance, consult `mount_cdfs(1M)`, or use SAM, SMH, or the SD-UX Interactive User Interface for `swcopy`, which you can use to “find” the CD-ROM.

1. Add WLM (product B8843CA) to the depot.

---

**NOTE** The WLM bundle B8843CA requires a `customer_id/codeword` pair to unlock it from the quarterly AR CD-ROM. Purchase of WLM entitles you to this codeword, which can be found on the Software Certificate you received from HP. The WLM codeword is needed to complete the next step.

---

Use a line similar to the following to add WLM to your depot:

```
# /usr/sbin/swcopy -x customer_id=cust_id      \  
                  -x codeword=codeword       \  
                  -s mount_path              \  
                  B8843CA @ depot_path
```

where `mount_path` is the path to the CD-ROM and `depot_path` is the same fully specified depot as in previous steps. For example:

```
# swcopy -x customer_id=xyzCorp -x codeword=123456789101bcdf \  
        -s /cdrom B8843CA @ /tmp/wlm_depot
```

- m. Verify the depot:

```
# swverify -d \* @ depot_path
```

where `depot_path` is a fully specified pathname. For example:

```
# swverify -d \* @ /tmp/wlm_depot
```

## Step 6. Disable PRM from starting automatically

If you have previously installed WLM or PRM, ensure that PRM is disabled from starting automatically by setting the variable `PRM_CONFIG` to 0 in the file `/etc/rc.config.d/prm`:

```
PRM_CONFIG=0
```

## Step 7. (Optional) Disable PRM's remote configuration daemon

The PRM remote configuration daemon is not useful with WLM. It need not be started. It is recommended that it be disabled.

If you have previously installed WLM or PRM, ensure that PRM's remote configuration daemon is disabled from starting automatically by setting the variable `PRM_RMTCONF` to 0 in the file `/etc/rc.config.d/prm`:

```
PRM_RMTCONF=0
```

## Compatibility information and installation requirements

### Step 8. Shut down WLM if it is currently running

Run the following command on every system or partition where you will install/upgrade WLM:

```
# /opt/wlm/bin/wlmd -k
```

Also shut down any running instances of the WLM global arbiter:

```
# /opt/wlm/bin/wlmpard -k
```

### Step 9. Install patches and install/upgrade WLM

---

**NOTE** The procedure below explains installation using the command-line version of `swinstall`. For information on using the `swinstall` GUI, refer to the section “Using the `swinstall` Graphical User Interface (GUI)” on page 41.

---

WLM is bundled as B8843CA, which contains the products Workload-Mgr, WLM-Monitor, and WLM-Dev tools. If desired, you can use `swlist` to list the product’s filesets:

```
# swlist -l fileset -s depot_path B8843CA
```

Installing WLM ensures that the following product is also installed:

- WLM Toolkits A.01.10.xx (WLMToolkits)

If the system already has the correct version installed, it will not be modified.

The WLM product B8843CA replaces the WLM product B8843BA from WLM Version 1.0. If you install using `swinstall`’s `match_target` option, the new B8843CA product will match your installed B8843BA product.

Workload Manager is installed in `/opt/wlm/`. The manpages are also installed in `/opt/wlm/`.

- a. Plan for system downtime. This should include generation of a current system backup and notification to all users.
- b. Shut down MeasureWare if it is running:

```
# /usr/sbin/init.d/mwa stop
```

c. Install WLM

There are two methods for installing WLM. You can:

- Install WLM from the depot you created in Step 5
- Install WLM directly from the CD-ROM

These options are discussed below.

**Installing WLM from a depot**

If you are installing WLM from your local depot, use `swinstall` as follows:

```
# /usr/sbin/swinstall -x patch_match_target=true      \  
                    -x autoreboot=true              \  
                    -s depot_path                    \  
                    B8843CA [patch_bundles] [individual_patches]
```

where `depot_path` is the fully specified path of the depot just created. For example:

```
# swinstall -x patch_match_target=true      \  
           -x autoreboot=true              \  
           -s /tmp/wlm_depot                \  
           B8843CA GOLDAPPS11i GOLDBASE11i PHKL_32619
```

SD-UX will attempt to build a new kernel and reboot the system.

---

**NOTE** Starting with WLM A.03.01, WLM no longer requires a reboot. The `autoreboot` option need no longer be set to `true` to satisfy requirements of the HP PRM kernel-building fileset PRM-Sw-Krn.PRM-KRN. Consequently, the SD-UX Interactive User Interface no longer prompts for an OK to reboot.

---

**Compatibility information and installation requirements****Installing from the CD-ROM**


---

**NOTE** Installing without the benefit of a depot may cause unnecessary reboots. For information on reducing reboots, see Step 5 on creating a software depot.

---

If you are installing WLM from the CD-ROM:

- i. Install the general support patch bundles and software from Step 1

If your system is missing any general HP-UX patch bundles or software mentioned in Step 1, install them now.

Be sure to read the patch bundle documentation for any special installation instructions.

- ii. Install WLM

---

**NOTE** The WLM bundle B8843CA requires a customer\_id/codeword pair to unlock it from the AR CD-ROM. Purchase of WLM entitles you to this codeword, which can be found on the Software Certificate you received from HP. The WLM codeword is needed to complete the next step.

---

Use a line similar to the following to install WLM:

```
# /usr/sbin/swinstall -x customer_id=cust_id      \
                    -x codeword=codeword        \
                    -x autoreboot=true          \
                    -s mount_path B8843CA
```

where *mount\_path* is the path to the CD-ROM. For example:

```
# swinstall -x customer_id=xyzCorp      \
            -x codeword=123456789101bcdf \
            -x autoreboot=true          \
            -s /cdrom B8843CA
```

- iii. Install WLM-specific patches from Step 4

If your system is missing patches from Step 4, install them now.

Be sure to read the patch documentation for any special installation instructions.

---

- d. Restart MeasureWare if it was running before the WLM upgrade/installation:

```
# /usr/sbin/init.d/mwa start
```

**Step 10.** Check the file `/var/adm/sw/swagent.log`

Check the file `/var/adm/sw/swagent.log` for messages relating to additional tasks that may need attention following the `swinstall` of WLM bundle B8843CA.

**Step 11.** (Optional) Ensure you have the proper version of GlancePlus

If you plan to use WLM data collectors that are based on Glance Adviser (such as `glance_app` or `glance_tt`), be sure to install and use GlancePlus C.03.35.00 or later to avoid a memory leak that occurs when Glance Adviser runs continuously. If you are using HP-UX 11i v1, install GlancePlus when installing the 11i Enterprise or 11i Mission-critical Operating Environments. If you are using processor sets, install GlancePlus C.03.58.05 or later.

**Step 12.** (Optional) Install files for integration with HP Systems Insight Manager

If you plan to use WLM with HP Systems Insight Manager (or Servicecontrol Manager), install the fileset `CMSConfig.WLMB-CMS-Tools` on the CMS. This fileset is available from the depot `/var/opt/mx/depot11` on the host where WLM has been installed.

If you are installing this fileset on a CMS that has WLM installed, the installation will fail if the fileset is not compatible with (does not have the same revision string as) the installed WLM.

**Step 13.** (Optional) Merge new `/etc/rc.config.d/wlm` file with your modified file

If you have made any modifications to your `/etc/rc.config.d/wlm` file by setting any of the variables ending in “`_ENABLE`” to 1, merge the file `/opt/wlm/newconfig/etc/rc.config.d/wlm` with your `/etc/rc.config.d/wlm` file.

As of WLM A.02.02, the `/opt/wlm/newconfig/etc/rc.config.d/wlm` file includes the variables `WLMCOMD` and `WLMCOMD_ENABLE`.

---

**NOTE**

If this was the first time you installed WLM and you have previously installed PRM, be sure to edit the file `/etc/rc.config.d/prm` to set the following variables as indicated:

```
PRM_CONFIG=0
```

```
PRM_RMTCONF=0
```

**Compatibility information and installation requirements**

These settings prevent PRM and its remote configuration daemon from starting automatically at reboot. (The PRM remote configuration daemon is not useful with WLM and does not need to be started.)

---

## Using the swinstall Graphical User Interface (GUI)

Using the swinstall GUI can simplify installation. The steps below illustrate how to use the GUI to install patches and WLM from a depot (Step 9 above):

**Step 1.** Start the swinstall GUI:

```
# /usr/sbin/swinstall -s depot_path
```

where *depot\_path* is the path for the depot you created for the patches, patch bundles, and WLM.

**Step 2.** Select “Manage Patch Selection” under the Actions menu. (Do not select the menu item labeled “Match What Target Has”. That item is for loading products, not patches.)

**Step 3.** From the resulting submenu, select:

“Automatically select patches for software installed on the target”.

Because this is not the default selection, you may need to unmark the other option before proceeding.

The swinstall command will now scan the depot, and mark patches to be loaded.

**Step 4.** Return to the Actions menu and select any other individual patches, patch bundles, and WLM (B8843CA) for installation once the marking phase is complete.

**Step 5.** Choose “Install (analysis)...” from the Actions menu.

**Step 6.** Select OK to install the software after the analysis phase is completed.

Because these products contain kernel filesets, SD will prompt you that a reboot is necessary.

---

### NOTE

If this was the first time you installed WLM and you have not previously installed PRM, be sure to edit the file `/etc/rc.config.d/prm` to set the following variables as indicated:

```
PRM_CONFIG=0
```

```
PRM_RMTCONF=0
```

These settings prevent PRM and its remote configuration daemon from starting automatically at reboot. (The PRM remote configuration daemon is not useful with WLM and does not need to be started.)

---

## Patches and fixes in this version

WLM version A.03.02.xx includes the following fixes.

### **wlmcert does not truncate an existing certificate file when a new certificate with the same name is installed**

In previous versions of WLM, if a new certificate is installed where an existing certificate of the same name exists, `wlmcert` fails to truncate the existing certificate file as it should when that file is bigger than the new one. For example, assume the certificate named `host1.pem` is already installed at the following location:

```
/var/opt/wlm/certstore/truststore/host1.pem
```

If you install a new certificate for `host1` and the new certificate file `host1.pem` is shorter than the existing one, `wlmcert` does not truncate the old file, leaving remnants of the old file in the new file.

This problem has been fixed so that the certificate file is truncated properly.

This fix addresses CR# JAGag04768.

### **CPU usage might increase significantly when using `sg_pkg_active` script with Serviceguard 11.16 or earlier**

In previous releases of WLM, using the WLM `sg_pkg_active` script with versions of Serviceguard earlier than 11.17 could result with relatively high CPU usage in complexes that include a high number of Serviceguard packages.

This problem has been fixed.

### **WLM Temporary Instant Capacity 15-day threshold too limiting**

Temporary Instant Capacity (TiCAP) activates capacity in a temporary “calling-card fashion” such as in 30-day increments (where a day equals 24 hours for one core). With Temporary Instant Capacity on the system, any number of Instant Capacity cores can be activated as long as your prepaid temporary capacity time has not expired. With earlier versions of WLM, if fewer than 15 processing days of temporary capacity are available, WLM stops activating Temporary Instant Capacity resources. This threshold can be too limiting for certain environments.

Beginning with this release of WLM, a new global arbiter configuration keyword (`utility_reserve_threshold`) enables you to increase or decrease this threshold. The default is 15 days. The following example shows the threshold set to 5 days, causing WLM to continue allocating temporary capacity until 5 or fewer days of temporary capacity resources are available:

```
par {  
    utility_reserve_threshold = 5;  
}
```

You can set the keyword to 0 to cause WLM to always activate temporary capacity resources as long as any number of temporary capacity resources are available. Note that it is recommended that some quantity of temporary capacity always be kept in reserve. Having a buffer of temporary capacity allows you to avoid delays activating additional cores when a purchase of additional capacity is necessary. For more information, see the *HP-UX Workload Manager User's Guide* or `wlmparconf(4)`. For information on additional restrictions involving activation of temporary capacity resources, see the Instant Capacity (iCAP) documentation.

This fix addresses CR# JAGaf97368.

## **prmmmonitor does not display memory “Upper Bound” for WLM groups**

In a WLM configuration that includes workload groups for which `gminmem`, `gmaxmem`, and `memweight` are specified, the `prmmmonitor MEM` command fails to display the current PRM memory cap in the “Upper Bound” column. This problem occurs when using `prmmmonitor` to check the memory entitlement and usage for each group.

A new “MB Max” column has been added to the `prmmmonitor MRG` command display, showing the maximum amount of memory available to the group. Because this enhancement is introduced in PRM C.03.02, that version of PRM (or later) must be installed on your system in order to get the new `prmmmonitor MRG` display.

Note that starting with WLM A.03.02, the `wlminfo group` command has been enhanced to display memory utilization of all groups in the current deployed configuration, and used with the `-v` option, displays each group's `gmincpu`, `gmaxcpu`, `gminmem`, and `gmaxmem` values, if they are available. The new `-v` option is ignored if live data is not being displayed (for example, when the `-o` option is being used). If memory management is not being used, a dash (-) instead of a zero is displayed in the ‘Mem Shares’ column. If a group's `gmincpu`, `gmaxcpu`, `gminmem`, or `gmaxmem` value is not assigned in the current configuration, a dash (-) is displayed in the corresponding column.

This enhancement addresses CR# JAGaf86413.

## **If WLM cannot increase a partition's CPU resources, the `wlminfo par` command displays the wrong CPU (core) count**

If WLM is not able to increase a partition's CPU resources (number of cores), the `wlminfo par` command erroneously displays the same value for both the intended and actual core count. For example, if WLM fails to raise the count on a partition from 2 to 3, `wlminfo par` should display the intended core count as 3 and the actual core count as 2; however, the command displays both the intended and actual count as 3.

The `wlminfo par` command has been fixed to display the intended and actual core counts correctly.

This fix addresses CR# JAGaf88801.

## **WLM issues `vparstatus` failure warnings that can be ignored**

If the `vparstatus` command is used immediately after you use the `vparmodify` command, the `vparstatus` command might fail. The failure occurs because the virtual partition database is still locked by the virtual partition monitor. As a result, WLM issues the following warning:

```
vPar command failures seen while gathering partition status information.
```

This message can be ignored unless any of the following conditions occur:

- WLM issues error messages indicating that modifications of cores are failing
- System users notice that expected migrations of cores fail to happen

Beginning with WLM A.03.02, WLM is less likely to issue the command failure warning, as it now ignores `vparstatus` failures up until the last `vparstatus` call during modification of a virtual partition.

This fix addresses CR# JAGaf86468.

## **`wlmpard` ignores Temporary Instant Capacity (TiCAP) capacity changes**

Starting with WLM A.03.00, when `utilitypri` is enabled (in the global arbiter configuration file), `wlmpard` fails to notice temporary capacity changes once all hosts have connected. It only detects changes when a new host connects. As a result, when the amount of reserve capacity falls below the threshold (15 processing days by default, or the number of days specified with the `utility_reserve_threshold` keyword), `wlmpard` might continue to activate temporary capacity.

Beginning with WLM A.03.02, `wlmpard` will detect such changes and stop allocating reserve capacity after the threshold is reached.

This fix addresses CR# JAGag06301.

### **wlmpmconf might generate invalid configurations**

The script for converting PRM configuration files to WLM configuration files, `wlmpmconf`, might create configurations with invalid syntax when the PRM configuration file involves the `PRM_SYS` group and secure compartment records. Error messages indicate that unexpected syntax is found in various user and secure compartment statements in the `prm` structure.

Beginning with WLM A.03.02, `wlmpmconf` now creates valid WLM configuration files under these circumstances.

This fix addresses CR# JAGag03890.

### **WLM daemon (wlmd) memory consumption increases continually**

Memory use of `wlmd` might increase over time, especially when PRM API errors are being logged in the WLM message log.

Beginning with WLM A.03.02, this memory leak has been corrected.

This fix addresses CR# JAGag06034.

### **WLM configuration wizard does not allow primary\_host with PSET-based workloads**

When using the WLM configuration wizard (`wlmcw`) to create a configuration that includes PSET-based workloads and specifies the `primary_host` keyword for partition management, the wizard generates an error message.

As of WLM A.03.01, partition management is supported with configurations that include PSET-based workloads. The configuration wizard has been fixed to allow such configurations.

This fix addresses CR# JAGag07511.

---

## **Software availability in native languages**

The manpages are available only in English. The release notes and user's guide are available in English and Japanese.

---

## Security

This section highlights security items you should be aware of.

### Relationship between host name and SSL certificates

SSL certificates are created when you install WLM. This enables WLM to run in secure mode within the system on which it is installed. As of Version A.03.01, when you start WLM using the “/sbin/init.d/wlm start” script, the script uses secure mode by default. This requires that you distribute security certificates to all systems or partitions being managed by the same WLM global arbiter (`wlmpard`). In addition, if you upgrade WLM and the `/etc/rc.config.d/wlm` script had been modified prior to the upgrade, you must check that the following variables in `/etc/rc.config.d/wlm` are enabled (set to 1):

```
WLMD_SECURE_ENABLE
```

```
WLMPARD_SECURE_ENABLE
```

```
WLMCOMD_SECURE_ENABLE
```

The name of each certificate created when you install WLM is based on the name of the host where the certificate is generated. Thus, on `host1`, the certificate is named `host1.pem`. This makes it easier for you to identify trusted systems. If you have not yet assigned a host name to the system where WLM is being installed, the certificate is given the default name `loopback.pem`. When you assign a name to the host, security will continue to work even if the host name differs from the certificate name. To achieve a match between the host and certificate names, you can use the `wlmcert` command to remove the current certificate and then to reset the certificates so that the host and certificate names match. For more information on security certificates and the `wlmcert` command, see `wlmcert(1M)`. This and other WLM manpages are also available at the following location:

<http://www.hp.com/go/wlm>

When using WLM to manage partitions, each partition must have in its truststore the certificate of every other partition with which it is being managed.

---

#### NOTE

If you use Serviceguard on the system running `wlmpard`, any systems to which `wlmpard` might fail over must have the same certificates installed in their truststores as does the primary `wlmpard` node. Therefore, be sure to install the certificates from the systems managed by that `wlmpard` on any systems to which `wlmpard` might fail over. Also, install the certificates from all failover systems to the systems being managed by that `wlmpard`.

---

## Data collectors

Data collectors invoked by WLM run as root and can pose a security threat. Hewlett-Packard makes no claims of any kind with regard to the security of data collectors not provided by Hewlett-Packard. Furthermore, Hewlett-Packard shall not be liable for any security breaches resulting from the use of said data collectors.

## wlmgui and wlmcmd

WLM and the WLM GUI allow you to set up secure communications as described in `wlmcert(1M)`. If you choose not to use secure communications, here are several security tips:

- Do not use `wlmgui` over the Internet. Use `wlmgui` and `wlmcmd` only on trusted LANs where you trust all the users: All data exchanged between `wlmcmd` and `wlmgui`, including the user's password, is transmitted without encryption over the network.
- Restrict communications between `wlmcmd` and `wlmgui` to only authorized users to improve security.
- Rely—only to the extent that you trust your intranet—on the monitoring information from `wlmgui` to decide on a course of action.
- The WLM GUI sends data to `wlmcmd` over the network without verifying the recipient.
- Each connection to `wlmcmd` represents a separate process on the system. As such, each connection consumes resources, such as open file descriptors, a process ID, memory, and so forth. A large number of connections could result in denial of service. You can restrict connections by deploying `wlmcmd` on systems behind a firewall that blocks access to the port being used.

## Partitions

WLM manages virtual partitions and nPartitions through a global arbiter. WLM's global arbitration uses non-secured communications. A rogue user could manipulate the communications, resulting in one or more partitions being granted an incorrect number of cores. Use global arbitration only on trusted local area networks.

By default, `wlmpard` communicates to the partitions on a system through port 9691.

If the partitions use a firewall or if you are using the HP-UX Bastille product on the partitions, it is likely that communications on this port are being blocked. To use `wlmpard` in your environment, specifically allow port 9691 or another port to be open to incoming connections. If you use a port other than 9691, be sure to restart `wlmpard` to communicate on the new port.

If you use Bastille or the Install-Time-Security Levels to configure the IPFilter firewall, you may want to put the rules regarding which port to leave open in the following file:

```
/etc/opt/sec_mgmt/bastille/ipf.customrules
```

After that, run `bastille -b` to load the rules and make sure that Bastille does not remove them later during subsequent runs/lockdowns.

---

## Available manuals

Several related manuals are available. These manuals can be accessed at <http://docs.hp.com>. They are also available for order.

Specific WLM information can be found in:

- *HP-UX Workload Manager User's Guide* (B8844-90014)

The release notes for each version of WLM can be found at <http://docs.hp.com>.

The user's guide and release notes are available in `/opt/wlm/share/doc/`.

WLM Toolkits information can be found in:

- *HP-UX Workload Manager Toolkits User's Guide* (T1302-90040)

The WLMTK documentation is available in `/opt/wlm/toolkits/doc/`.

PRM information can be found in:

- *HP Process Resource Manager User's Guide* (B8733-90019)

Information about HP-UX system administration and Logical Volume Manager can be found in:

- *Managing Systems and Workgroups* (B2355-90950)

---

## WLM toolkits

Hewlett-Packard provides integration toolkits to simplify integrating Workload Manager with major third-party applications.

The toolkits are made available on the web as they are developed. Also, they are bundled with the subsequent Workload Manager release. To download toolkits, visit the Workload Manager web page:

<http://www.hp.com/go/wlm>

---

## Providing feedback

- Email your feedback to the WLM and WLM Toolkit development teams at the following address:  
[wlmfeedback@rsn.hp.com](mailto:wlmfeedback@rsn.hp.com)
- For a forum with other WLM users, visit the IT Resource Center's forum for HP-UX Workload/Resource Management:  
<http://forums.itrc.hp.com/cm/>
- For the latest patch information, white papers, and documentation, visit the Workload Manager web page:  
<http://www.hp.com/go/wlm/>

## **Training**

HP offers a course in HP-UX resource management using WLM. For information, including a course outline, visit:

<http://www.hp.com/education/courses/u5447s.html>