Workload Manager Hand's On Tutorial

Plus Global Workload Manager Intro

Dan Herington
Senior Software Architect
Hewlett-Packard

The Adaptive Enterprise

Business and IT synchronized to capitalize on change

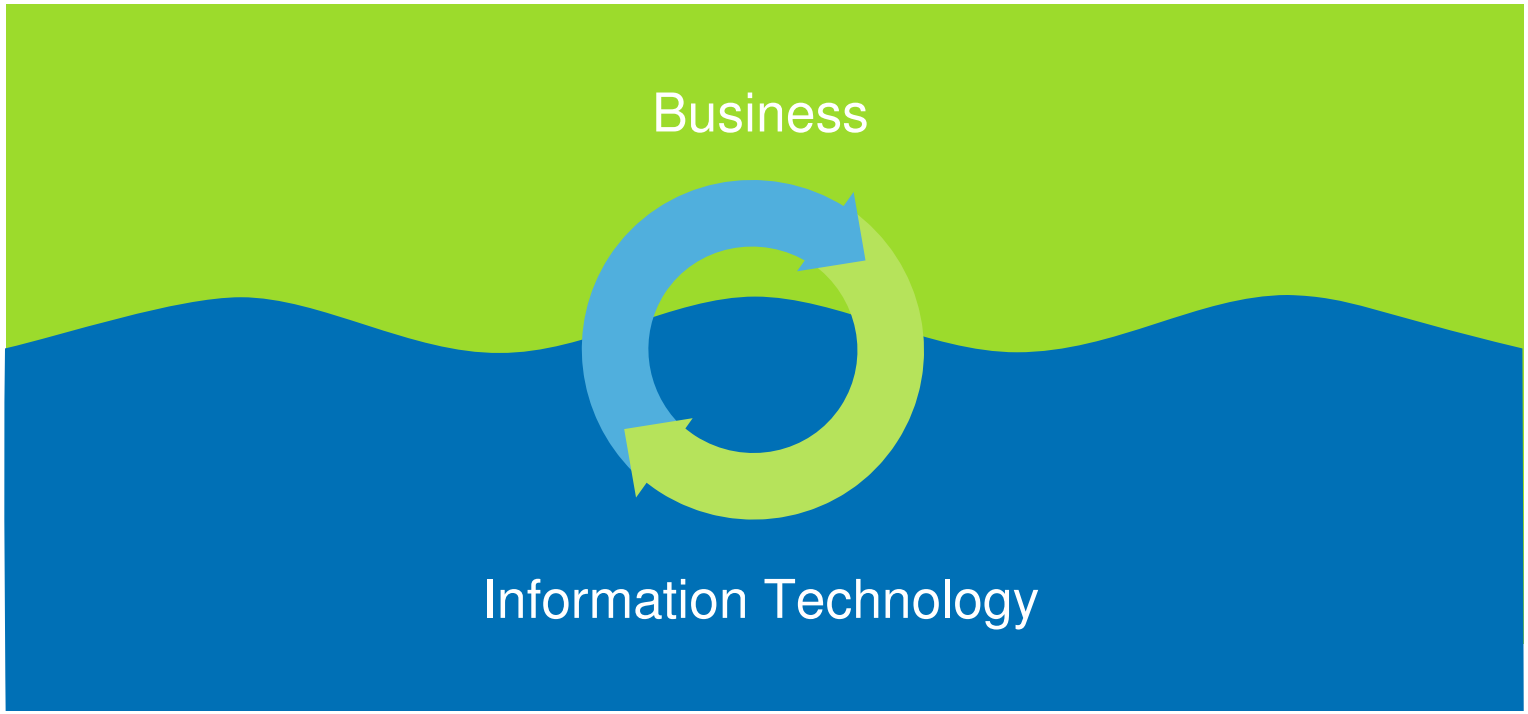

Business benefits: simplicity, agility, value

Why You Need an Adaptive Infrastructure

Tremendous amount of unutilized capacity

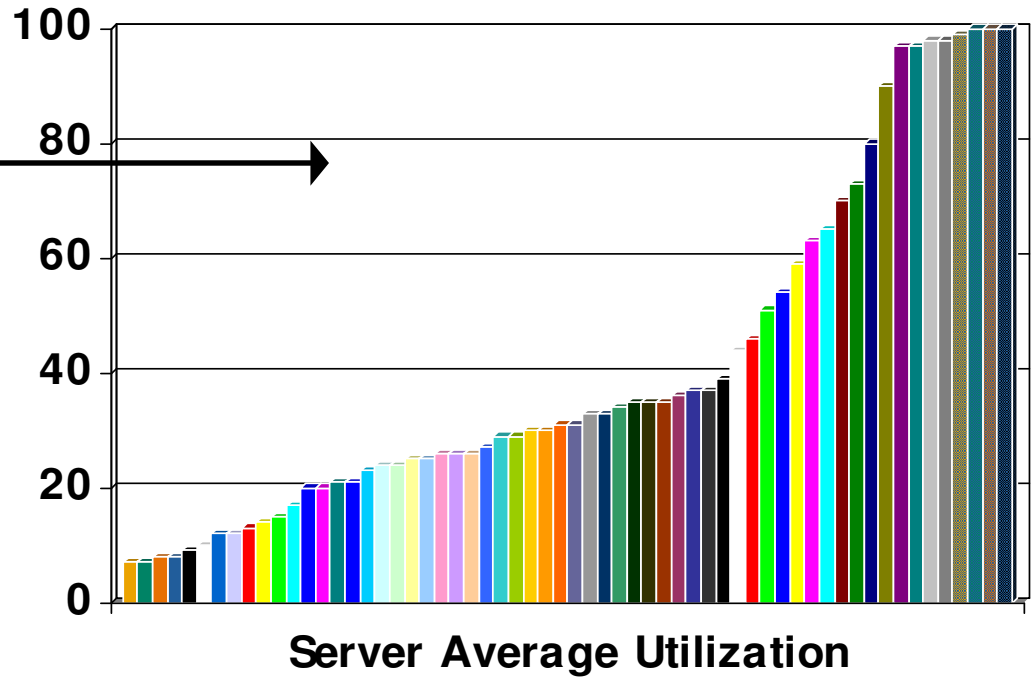

Yet these systems are unable to handle the load

- Most reports put average utilization at approximately 30%
- Some applications still not able to meet performance requirements

Why Utilization is So Low

- Mission Critical applications have peaks in load that require spare capacity

New Alternatives to Overprovisioning

- Build an Adaptive Infrastructure that can dynamically react to application peaks
- Implement Application Consolidation
  - Run many workloads on a small number of servers
- Share Spare Capacity
  - Provide spare capacity for multiple apps on the same system or systems – 40% reduction in spare capacity requirements
- Consolidation requires virtualization solutions to ensure applications are isolated from each other

Adaptive Enterprise vision

Business and IT synchronized to capitalize on change

1. Measure, assess and maintain a dynamic link between business and IT
2. Architect and integrate heterogeneous IT environments
3. Extend and link business processes across suppliers and customers
4. Manage and control business processes, applications and the whole IT environment

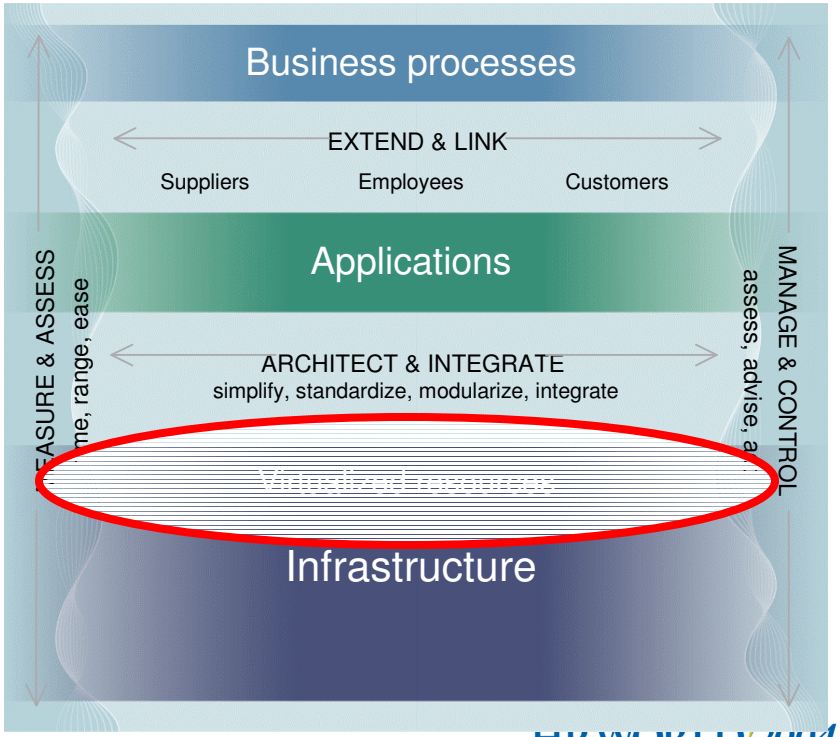

HP's Definition of Virtualization

An approach to IT that pools and shares resources
so utilization is optimized and
supply automatically meets demand

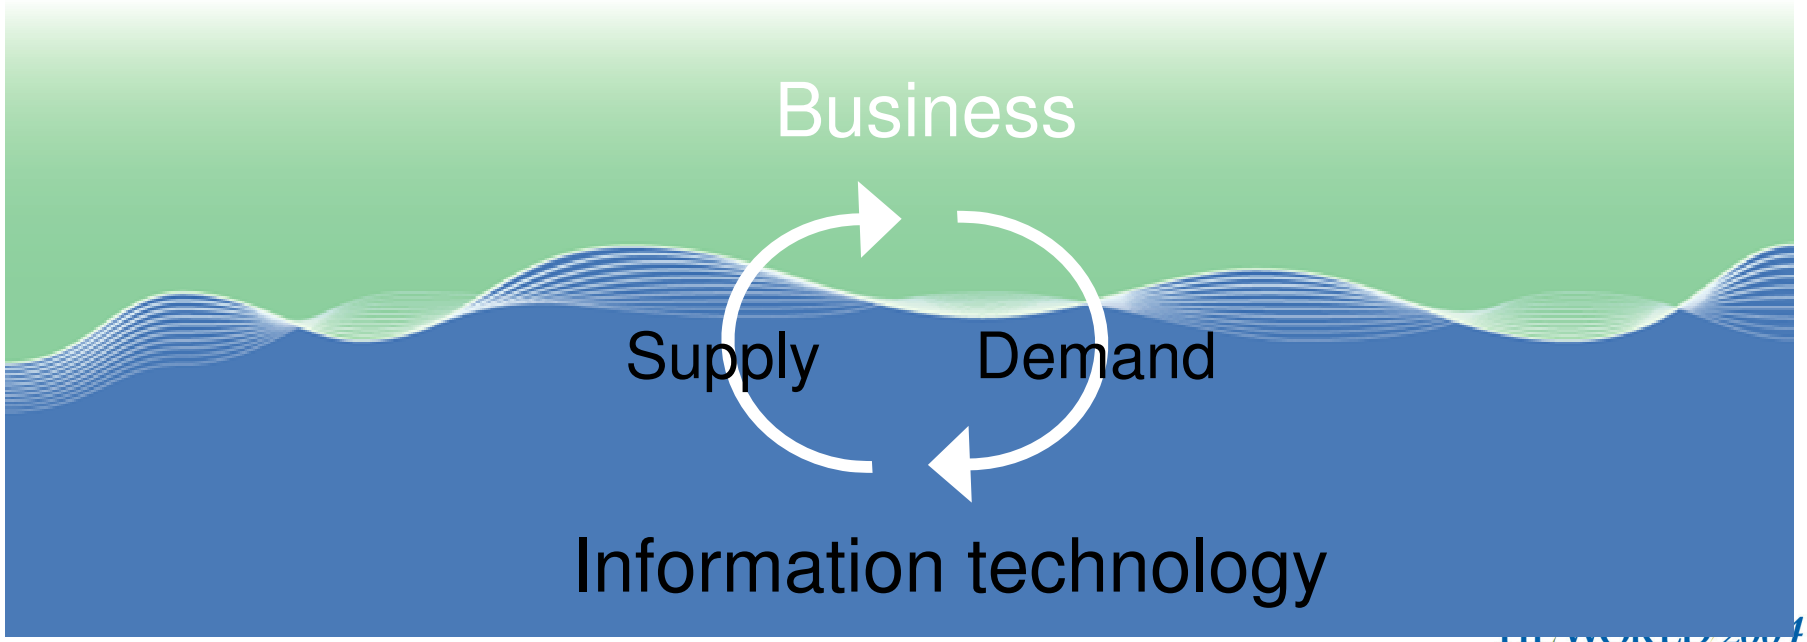

HP Virtual Server Environment for HP-UX 11i

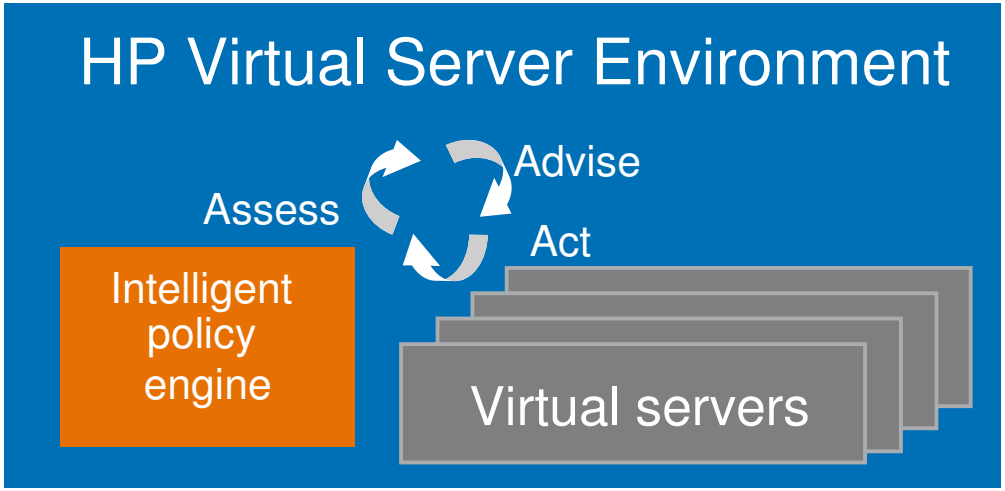

Expands and shrinks virtual servers in real time based on business priorities

- Better RoIT through optimized resource utilization
- Increased business agility through the capability to allocate resources on the fly
- Ensuring service levels through continuous real time assessment, advice, and action

HP Virtual Server Environment Technologies

HP Virtual Server Environment
Adaptive Enterprise Technologies

- Dynamically reconfigurable partitions
  - nPars with iCOD
  - Virtual Partitions
  - Resource Partitions
- Capacity on Demand
  - iCOD
  - iCOD Temporary Capacity (TiCOD)
  - Pay Per Use (PPU)
- VSE Management Tools
  - Systems Insight Manager and ParManager
  - HP-UX Workload Manager
  - Serviceguard

Definition of Partitioning

Partitions are physical or logical mechanisms for
isolating operational environments
within single or multiple
servers to offer the
flexibility of dynamic resizing
while ensuring that
applications can enjoy
protection from
unrelated events
that could otherwise cause disruption,
interruption, or performance degradation.

HP Partitioning Continuum for HP-UX 11i

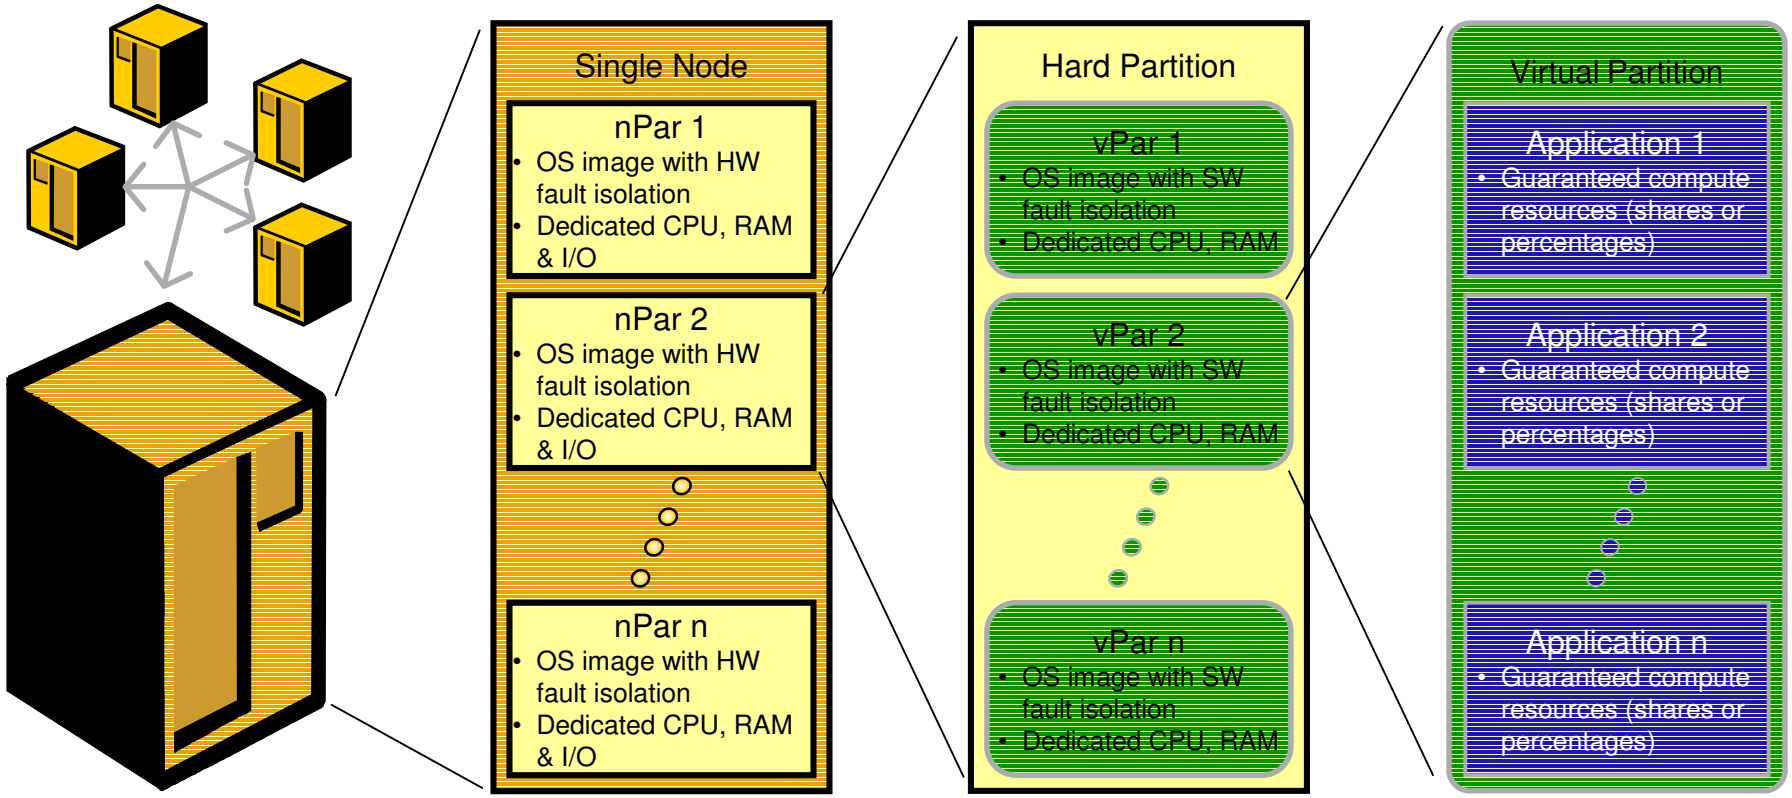

HP-UX WLM (workload manager)
- automatic goal-based resource allocation via set SLOs

HP-UX = Broadest Partitioning Portfolio

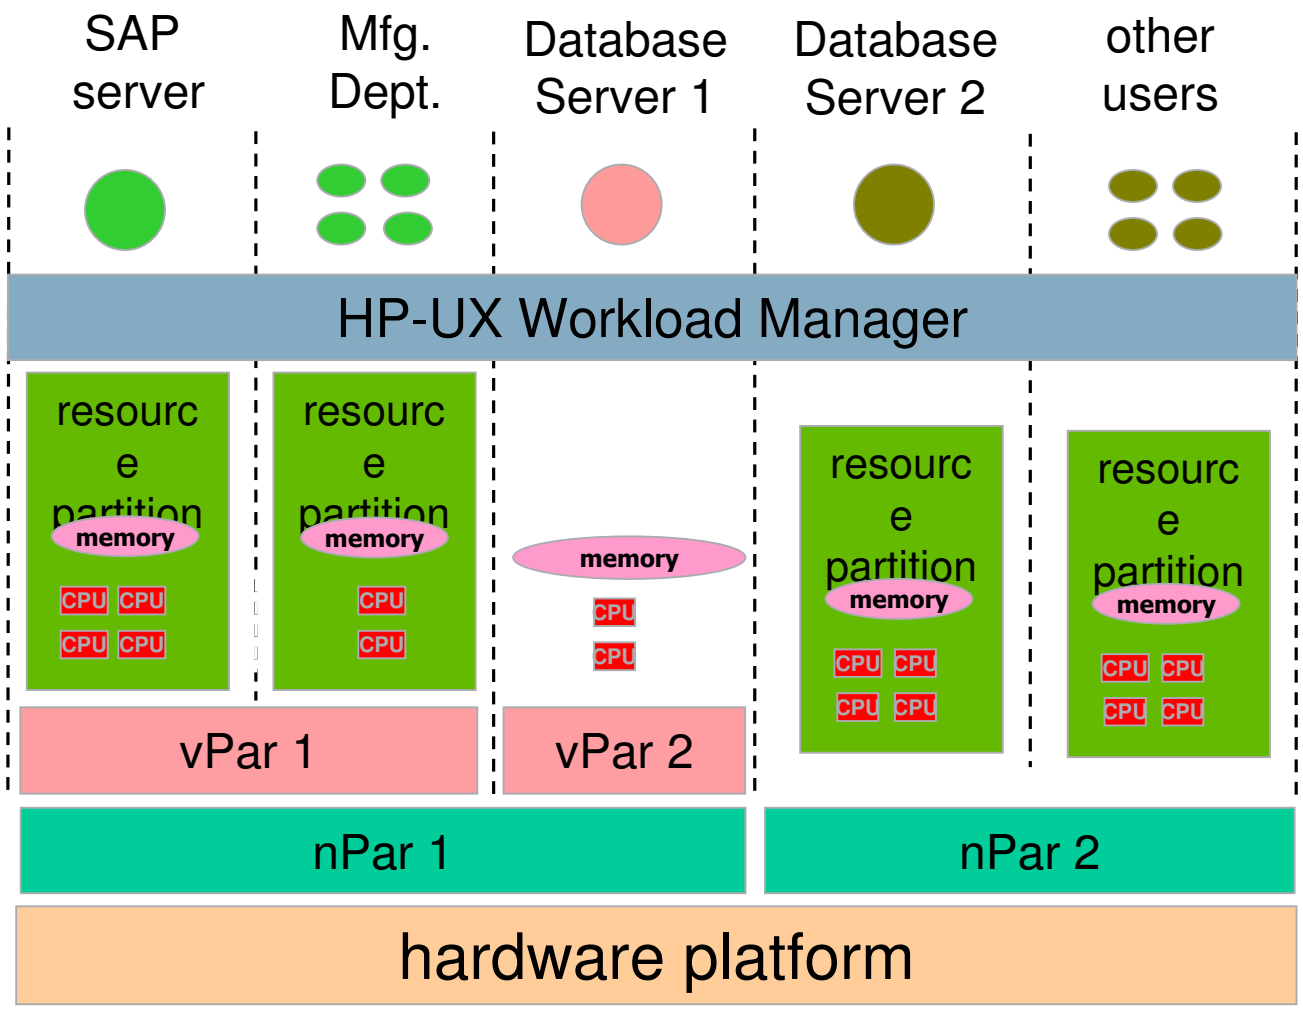

application space

WLM to automatically resize partitions

Adaptive Resource Partitions inside of an OS image

vPars

nPars

hp 9000

HP-UX On-Demand Technologies

- Instant Capacity on Demand (iCOD)
  - Activate new permanent capacity when needed
- Instant Capacity on Demand – Temporary Capacity (TiCOD)
  - Activate/deactivate new temporary CPU capacity when needed
- Pay-per-Use Utility Computing – (PPU)
  - Lease systems based on CPU utilization

VSE Management Tools

- HP Systems Insight Manager
  - Single Management Interface for HP-UX, Windows and Linux systems management
- HP-UX Workload Manager
  - Automated reallocation of resources based on business priorities and real-time workload performance or load
- HP Serviceguard
  - Industry-Leading High-Availability Clustering Solution

nPars

nPartitions

Multiple applications on the same server with full electrical isolation between partitions

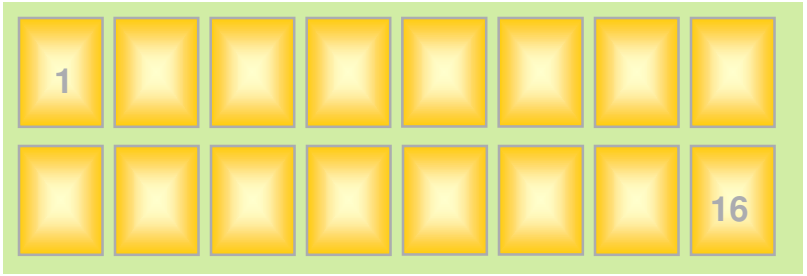

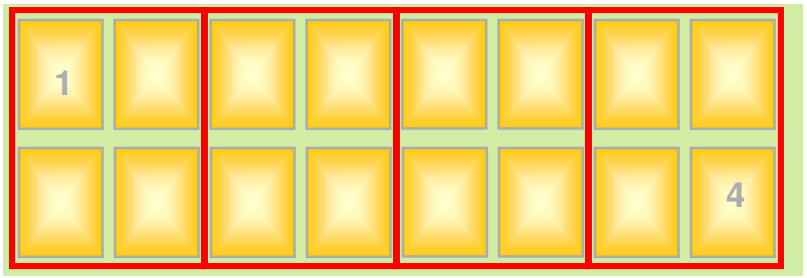

- **Increased system utilization**
  - partitioning Superdome into physical entities: up to16 nPartitions
- **Increased Flexibility: Multi OS**
  - Multi OS support: HP-UX, Linux (*), Windows (*)
  - Multi OS version support
  - Multiple patch level support
- **Increased Uptime**
  - hardware and software isolation across nPartitions
  - MC/ServiceGuard support (within Superdome or to another HP 9000 server)

hp's cellular architecture is very flexible

8 Socket system

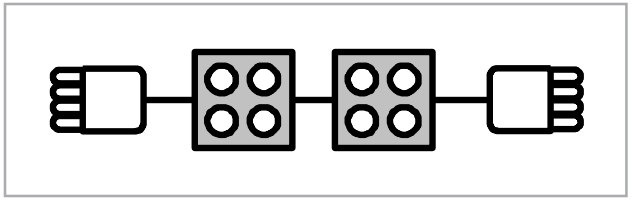

16 Socket system

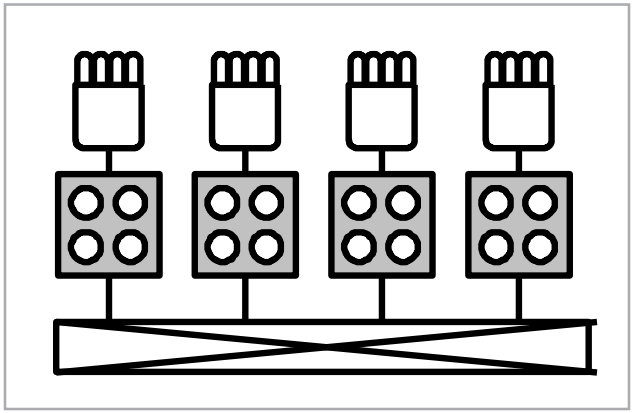

64 Socket system

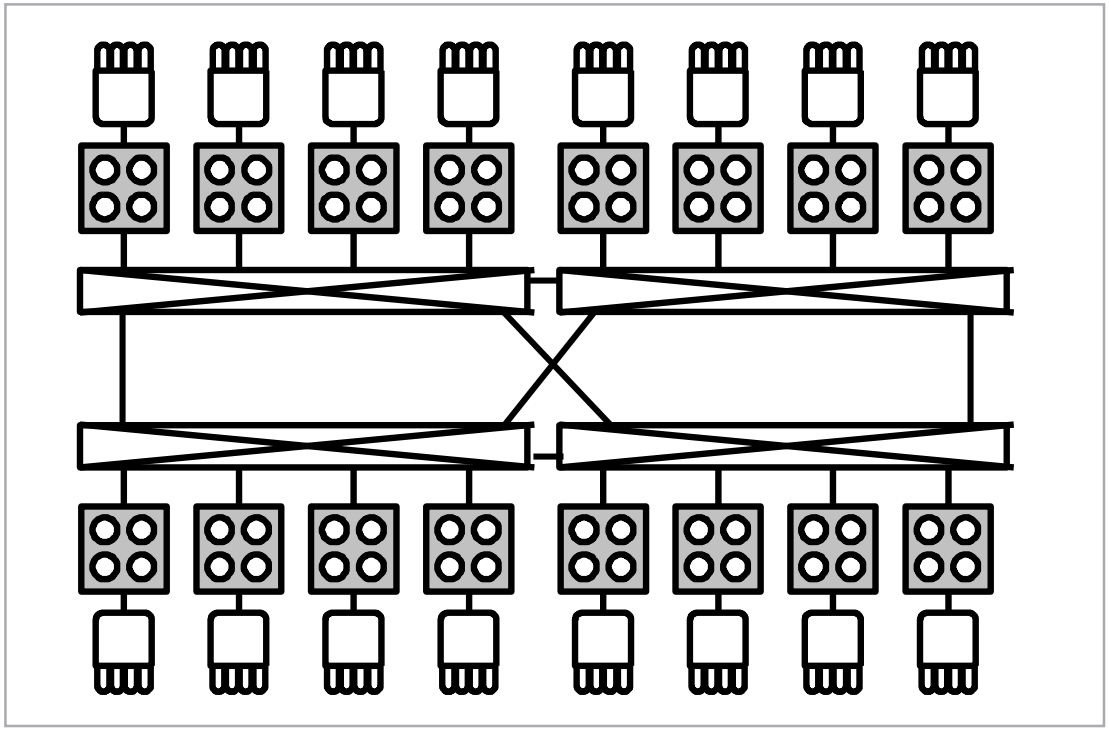

Legend

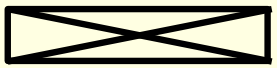

Two crossbar switches

4 Socket Cell

I/O Backplane

MX2 System Daughtercard

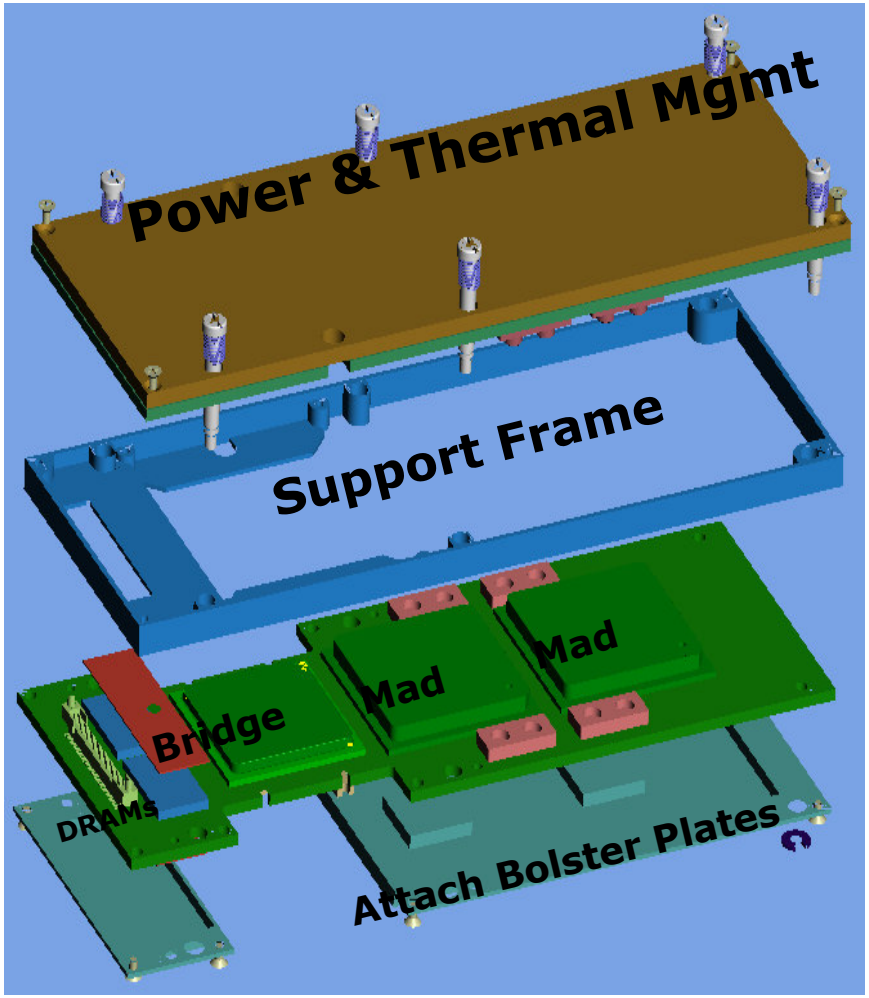

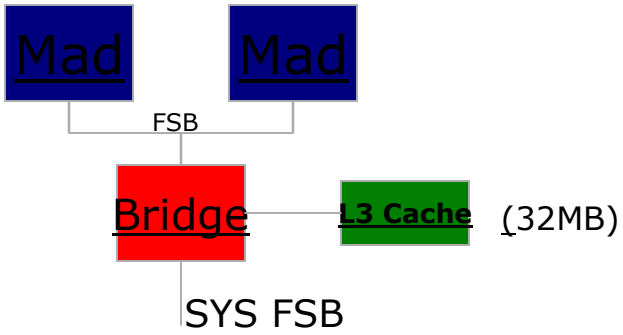

Partition Manager New Features

Significant Changes from ParManager on HP-UX 11i

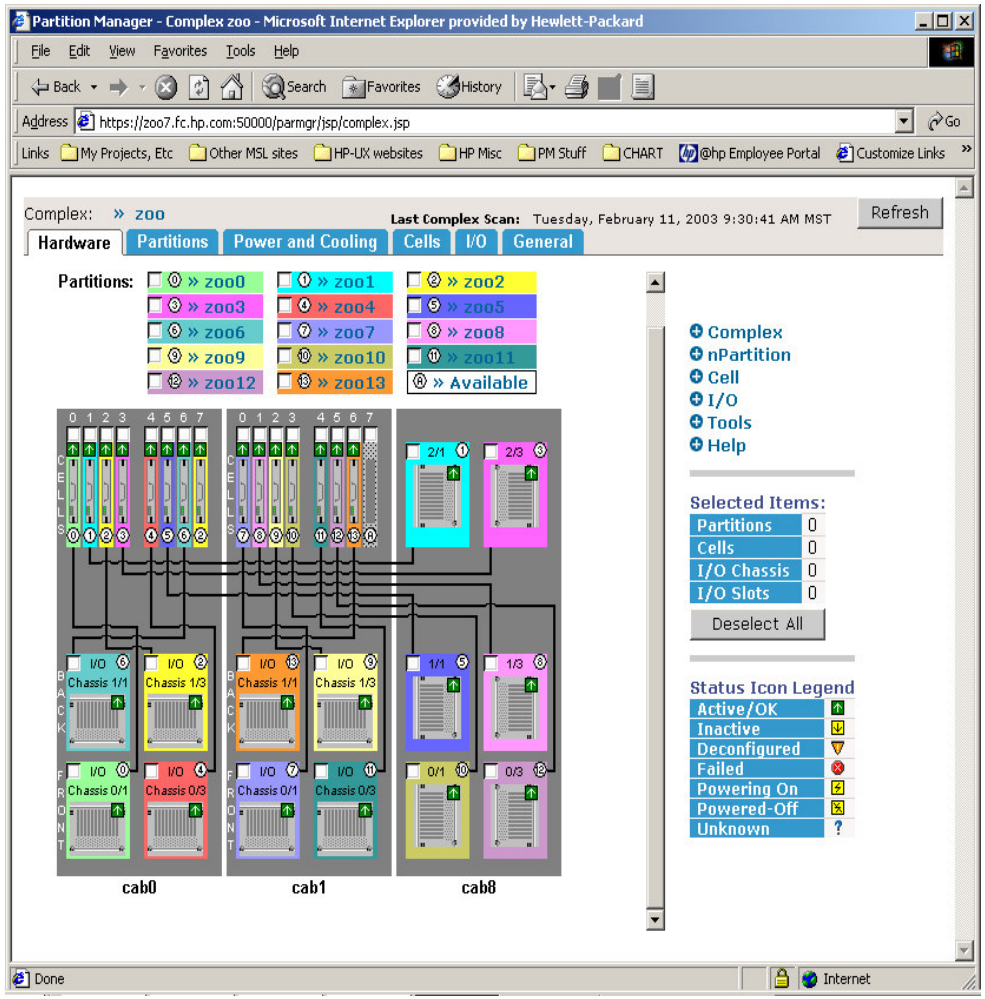

- ✓ New web interface
- ✓ Graphical “big picture” views of
  - nPars
  - Hardware in complex
- ✓ Supports new OS/HW features
  - Cell local memory for HP-UX 11i v.2 partitions
  - nPartition configuration privilege
- ✓ Remote admin of Superdome complex
- ✓ Compatible with iCOD/pay-per-use
- ✓ Native on Windows (2H 2004)

vPars

HP-UX Virtual Partitions

Multiple HP-UX instances running on the same system or in the same nPar

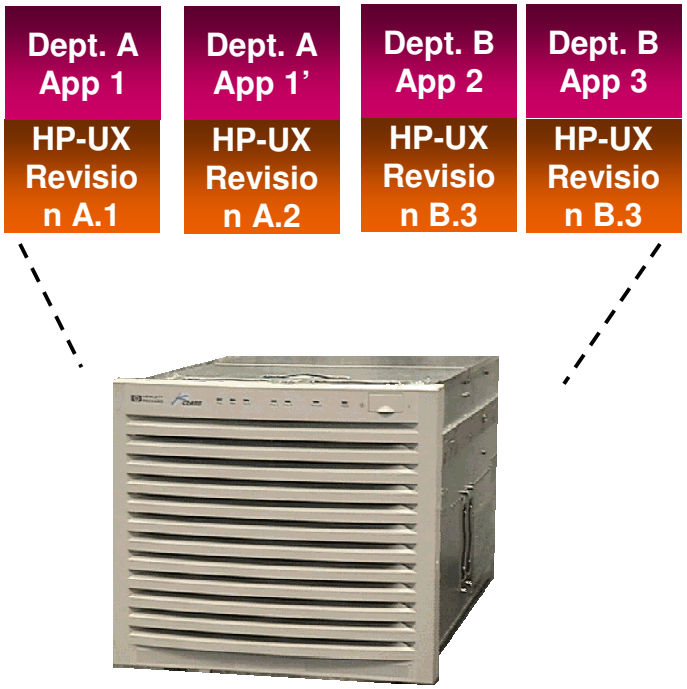

Increased system utilization

- partitioning a single physical server or hard partition into multiple virtual partitions for rp5405,rp5470, rp7400, Superdome, rp8400, rp7410, rp8420, rp7420

Increased Flexibility

- multiple independent instances of HP-UX
- dynamic CPU migration across virtual partitions

Increased Isolation

- application isolation across virtual partitions
- OS isolation
- individual reconfiguration and reboot

vPars logical overview

- multiple applications or multiple instances or versions of the same application
- provides name space and resource isolation

- creates illusion of many separate hardware platforms
- manages shared physical resources
- monitors health of operating system instances

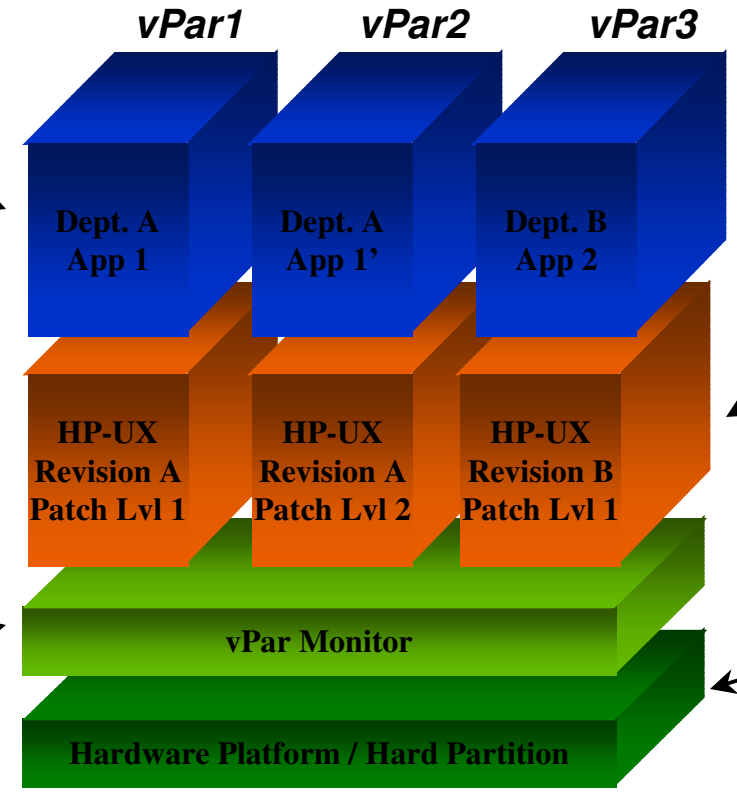

- each operating system instance tailored specifically for the application(s) it hosts
- operating systems instances are given a user-defined portion of the physical resources
- provides name space and resource isolation

- supported on rp5470, rp7400, Superdome, rp8400, rp7410, rp8420, rp7420 systems
- no additional platform support required

Partitionable Resources

- CPUs may be "bound" to a single partition or allowed to "float" among partitions
- bound CPUs require a partition reboot to be reassigned among partitions
- unbound CPUs may be dynamically reassigned among partitions

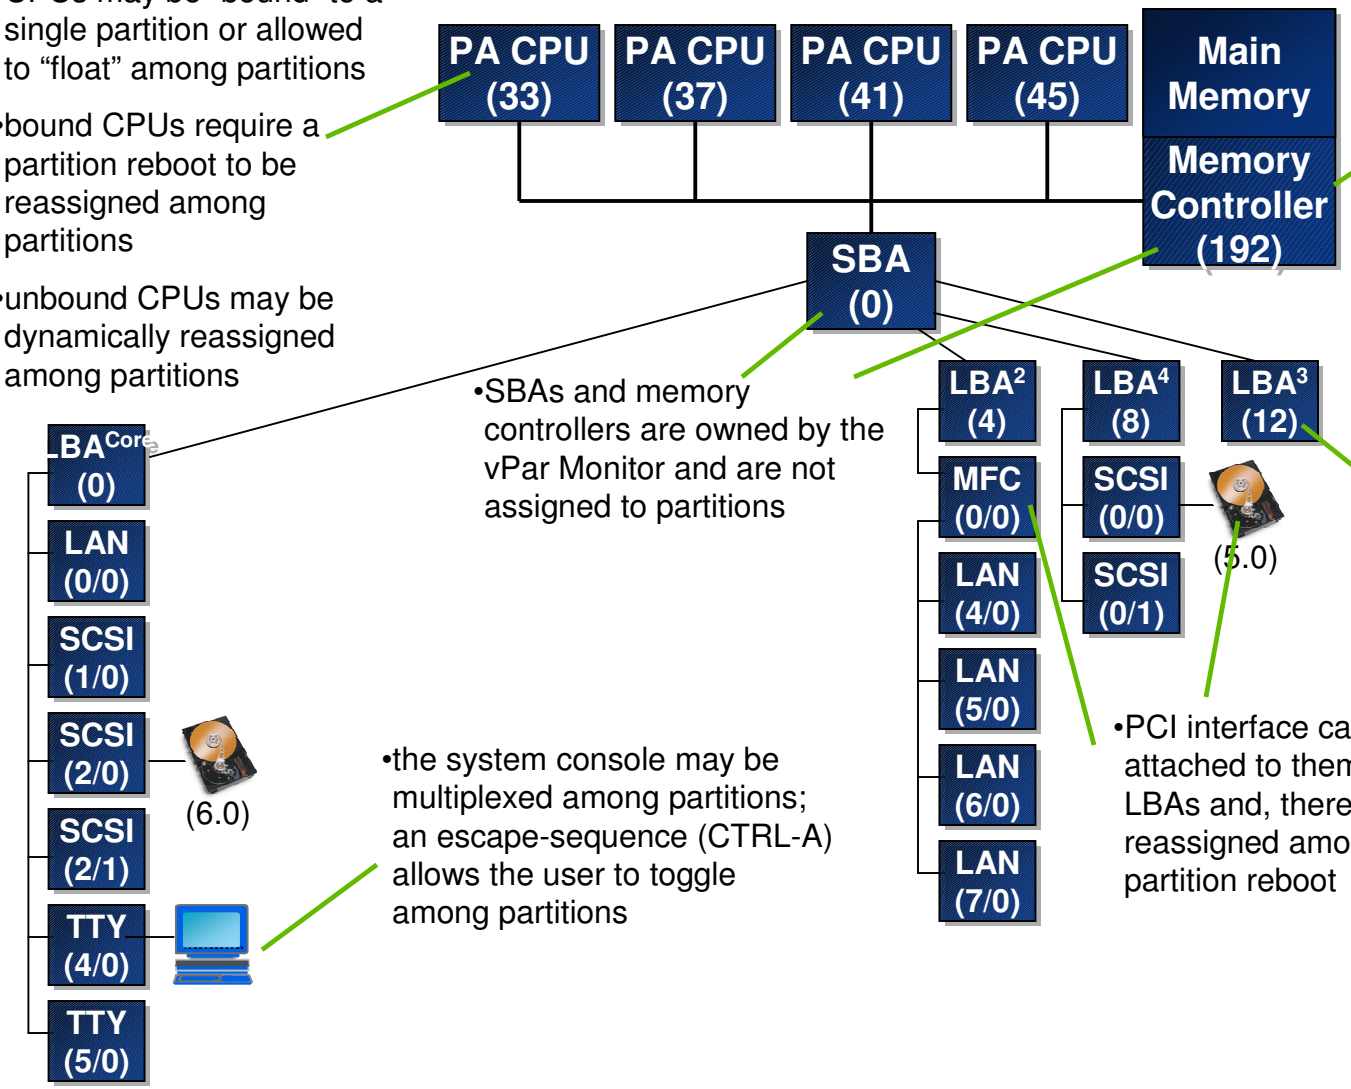

- main memory is allocated to partitions in multiples of 64MB ranges
- adding or removing memory to or from a partition requires a partition reboot

- SBAs and memory controllers are owned by the vPar Monitor and are not assigned to partitions

- LBAs are bound to a single partition
- adding or removing LBAs to or from a partition requires a partition reboot

- the system console may be multiplexed among partitions; an escape-sequence (CTRL-A) allows the user to toggle among partitions

- PCI interface cards and the devices attached to them are connected through LBAs and, therefore, cannot be logically reassigned among partitions without a partition reboot

Resource Partitions

Resource Partitions

Apps are running in the same OS, but have separate process schedulers and separate memory managers

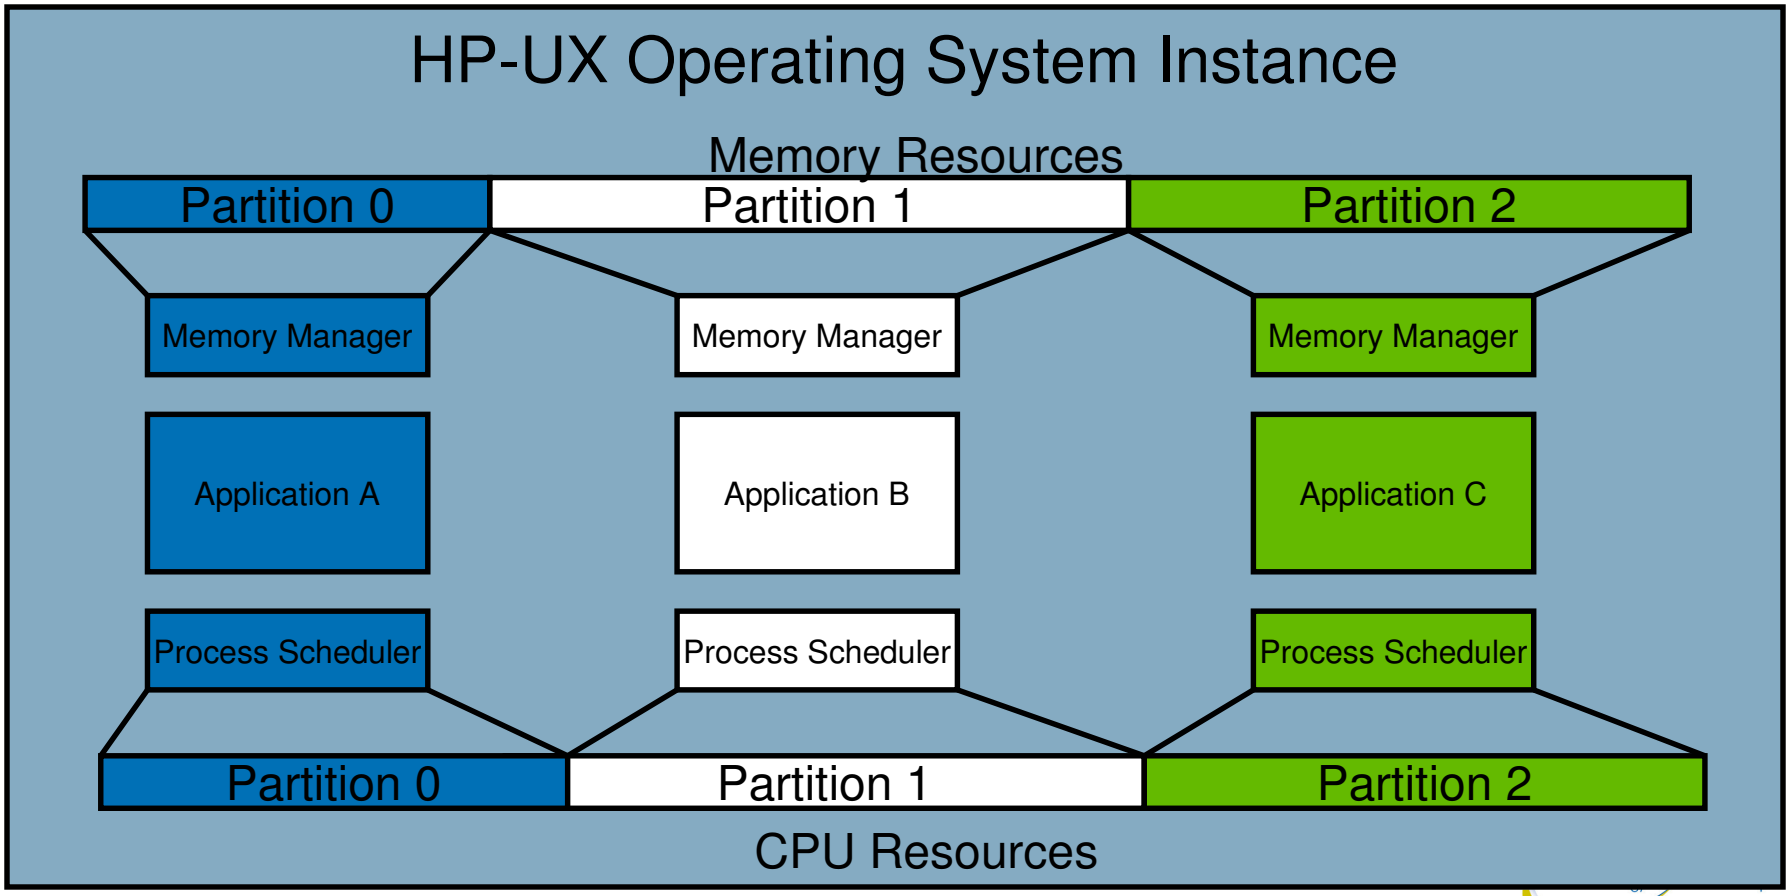

Resource Partitioning Features

- Resource controls:
  - CPU Controls
    - CPU allocation by percentage, shares, or whole processors
    - Optional capping in FSS partitions
    - Concurrent FSS and PSETs
  - Real memory controls
    - Each partition gets a separate memory manager in 11i
  - Disk bandwidth
    - Both LVM and Veritas VxVM Volume Groups
  - Automatic process assignment to partition
    - Users/Groups
    - Executable path/Process name
    - Children automatically run with parent by default

PRM CPU Scheduling

Fair Share Scheduler

- Shares become slots in a carousel
- Each slot is 10ms CPU tick
- FSS allows Standard Unix scheduler to schedule processes from current partition during tick

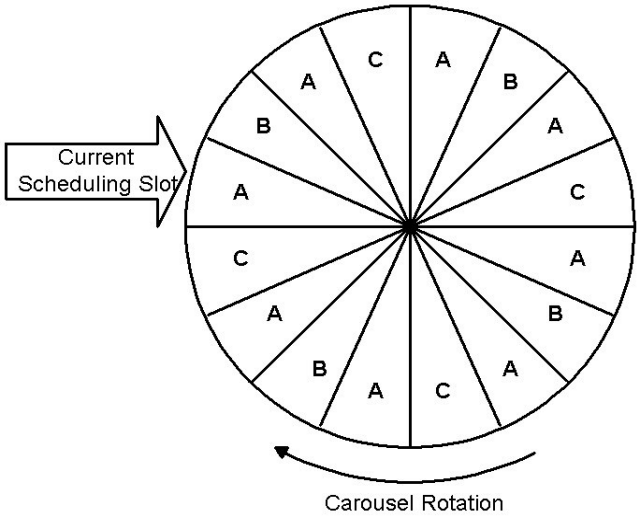

Processor Sets

- CPU is allocated on whole CPU boundaries
- Standard Unix scheduler allocates CPU within a PSET
  - Separate process schedulers for each partition
  - Processes in the partition are scheduled on the CPUs in the PSET assigned to the partition

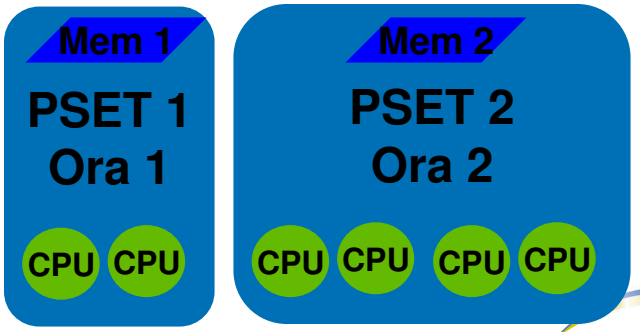

HP-UX Resource Partitions
Key Features and Benefits

- Controls allocation of CPU, real-memory and disk I/O bandwidth based on user-specified policies
- Applications do not require modification to work with PRM
- Configuration can be changed at any time - even under load
- Supports resource policies based on users and applications
- Supports the configuration of both FSS and PSET based partitions on the same system
- Tight integration with HP-UX systems management and workload management tools

On-Demand

HP Instant Capacity On Demand (iCOD)

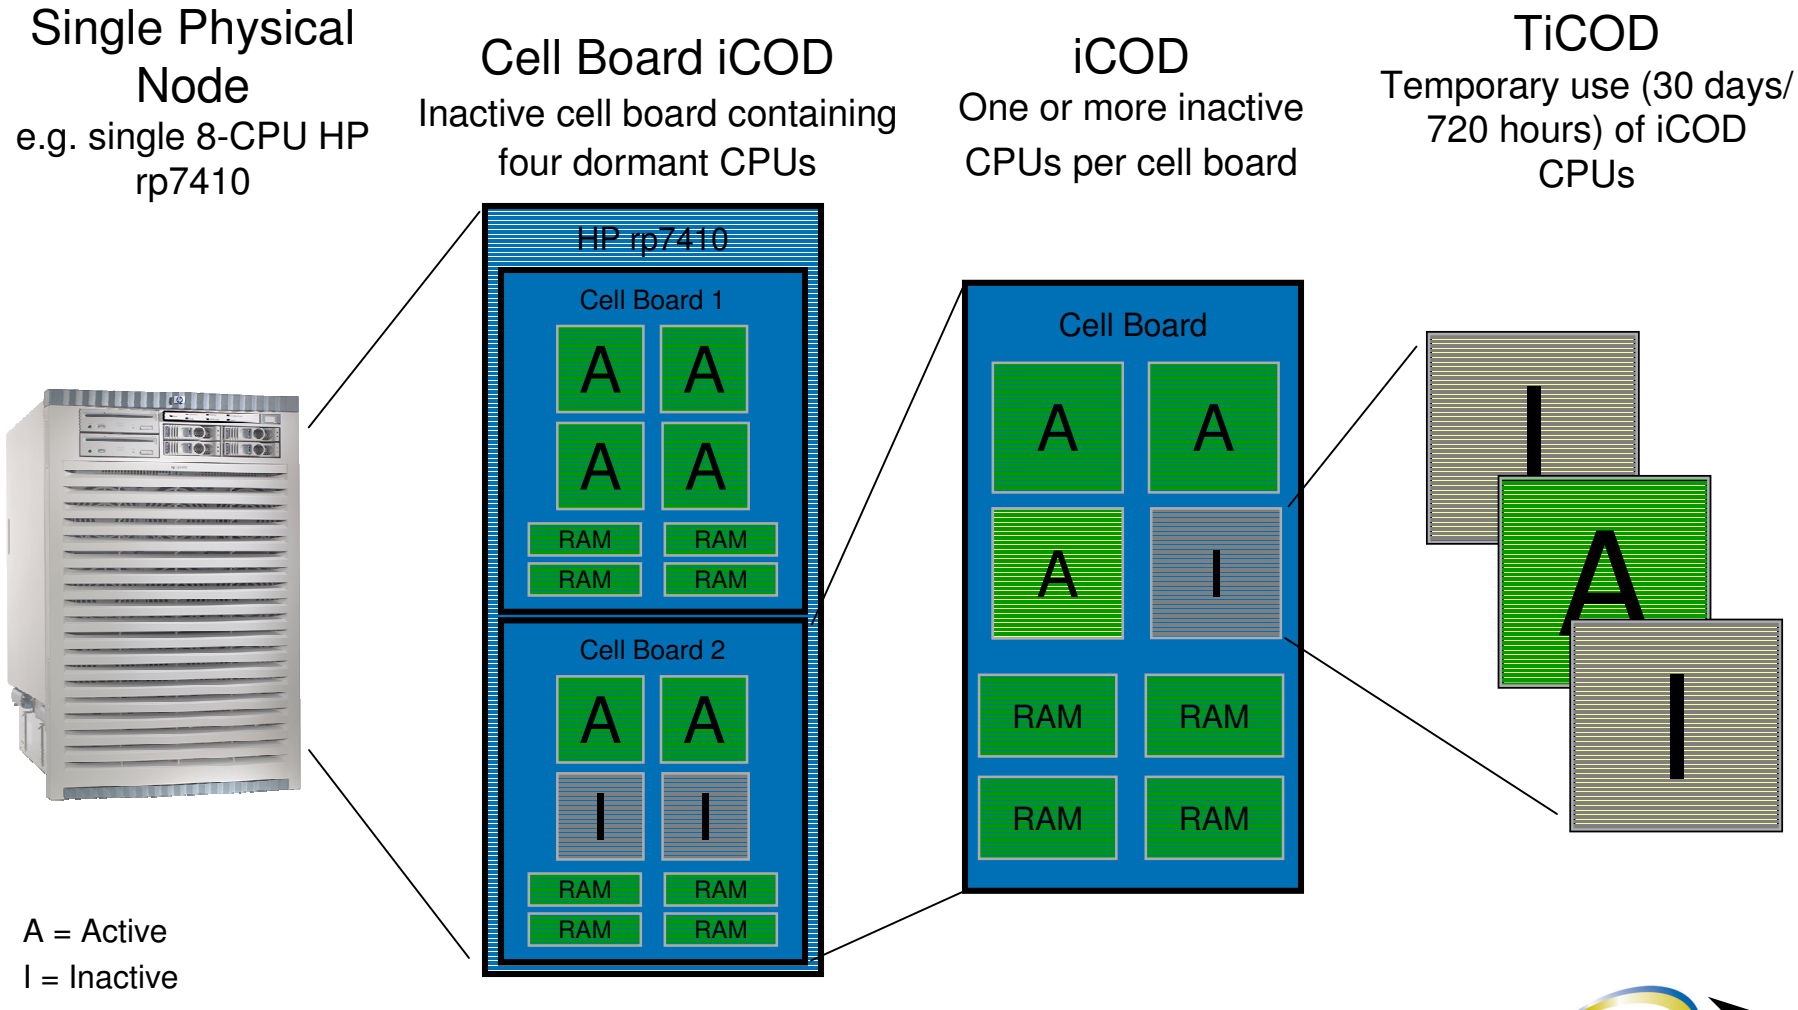

Instant Capacity on Demand (iCOD)

- System acquired with inactive processors
- Processors are paid for when they are activated
  - Price paid is current price when activated
- CPUs can be activated on-line – no reboot required
- Excellent solution for expected growth

- iCOD is licensed for an entire complex
- CPUs can be deactivated in one nPar and activated in another – without reboot

iCOD Temporary Capacity (TiCOD)

- Alternative purchasing model for iCOD processors
- Temporary Capacity is purchased in 30 Day increments
  - 30 CPU-Days = 43,200 CPU-Minutes
- Any number of iCOD CPUs can be activated
- Activating processors causes the iCOD software to deduct minutes from the “bank”
- Deactivating the iCOD processors stops the deductions

- Excellent solution for:
  - Short term peaks in application load
  - Lower cost failover server - Activation of additional capacity upon failover of a large workload onto a failover server

Pay-per-Use Utility Computing (PPU)

- Type of lease
- Acquire a system with peak capacity required
- Monthly charge based on base payment plus a variable payment based on actual resource usage
- 2 utilization measurement models
  - Active CPU – CPUs are activated/deactivated and variable payment is based on how long CPUs were active
  - Percent Utilization – All CPUs are active and the system is monitored for CPU utilization – variable payment is based on average utilization of all CPUs
- Excellent solution for highly variable loads, especially revenue generating loads because costs are in line with revenues

Workload Manager

HP Workload Manager

- HP WLM is a state-of-the-art dynamic workload manager for HP-UX servers
  - It automatically adapts the partition configuration based on the loads on the applications running in those partitions and your business priorities
  - Supports:
    - Resource partitions and vPars and nPars with iCOD
    - Automatic activation/deactivation of iCOD and pay-per-use CPUs
    - Resource partition memory reallocation when workloads are activated/deactivated due to failover or batch job activation
- WLM helps you comfortably increase utilization while still ensuring that your mission critical applications maintain their performance requirements

WLM Service Level Objectives

SLO's use goals, constraints, and conditions.

An SLO consists of:

- **A workload (partition)**
- **Constraints (min, max cpu)**
- **A goal**
- **Priority**
- **Conditions (time of day, event, etc)**

Group A

Min CPU: 20%

Max CPU: 50%

Group A receives 3 shares for each additional user.

Policy applies 9am to 5pm AND

when ServiceGuard Package XYZ

WLM goal types

- Any of the following can be used to allocate resources to a workload:
  - resource utilization
    - CPU entitlement based on utilization of current entitlement
    - Easiest to configure – no data required
  - direct measurement of the performance of the workload
    - response time
    - throughput
  - measurement of load on application
    - number of users/processes
    - queue length

WLM Major Features

- 11i and 11iV2 support on PA and Integrity
- Support for Resource Partitions, vPars and nPars with iCOD
- Support for PPU and TiCOD
- Out of the box support of any workload with CPU Utilization controls
- Goal-based Service Level Objective Support based on any available data value
- Out of the box toolkits to support Oracle, Weblogic, Glance, Apache, Job duration/SAS, Serviceguard
- Command line and GUI utilities for configuration and monitoring
- Auditing (billing) utilities – utilities that accumulate the actual usage of resources by each workload over time
- Advisory mode - to allow customers to monitor their workloads without turning on active WLM controls

WLM support for Hard Partitions (nPars)

When the workload in nPar1 is busy, WLM will deactivate CPU's in nPar2 and activate the available iCOD processors in nPar1.

Since the total number of active processors on the system has not change, this does not incur any costs for activation of the iCOD processors.

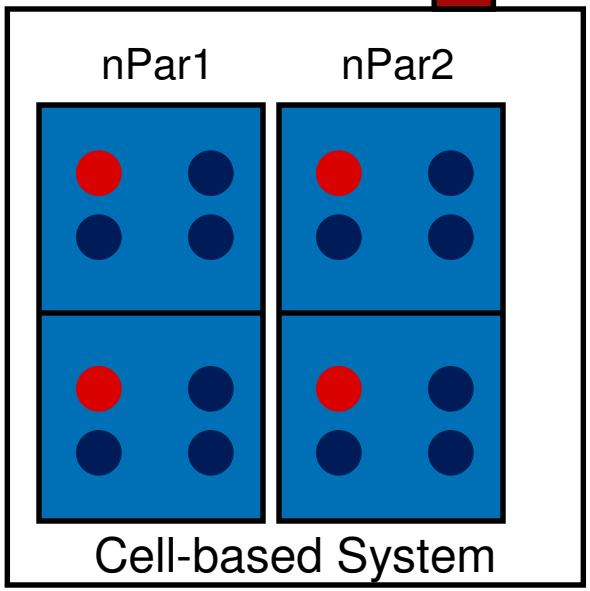

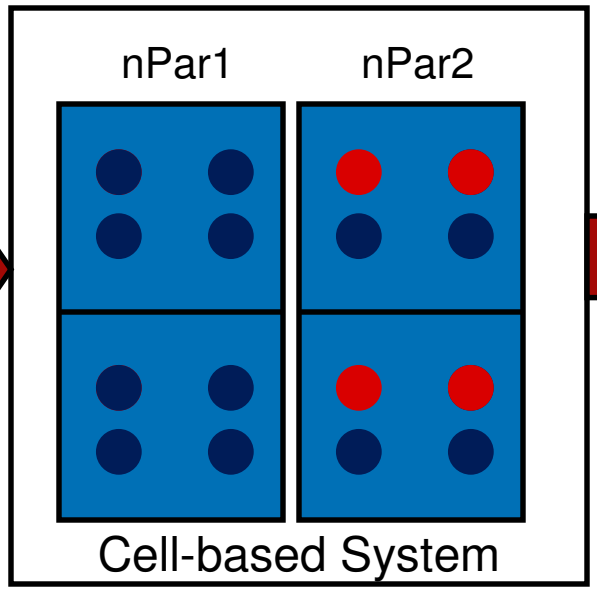

This is a 4-cell 16-CPU server with 2 nPars – each with 2 cells.

12 active processors on the system and 4 available iCOD processors

WLM will use goal-based Service Level Objectives to determine which nPar the 12 active processors should be running in.

When the workloads in nPar2 get busy, WLM can deactivate CPUs in nPar1 and activate them in nPar2.

This allows each nPar to scale from 4 to 8 CPUs depending on the status of the workloads running in each nPar.

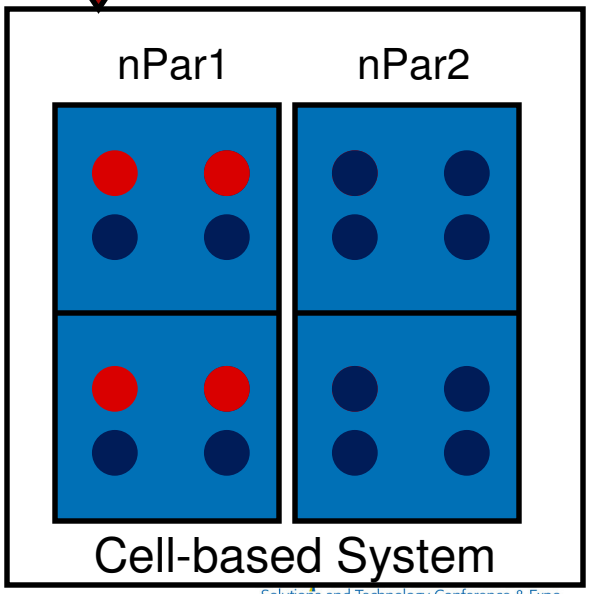

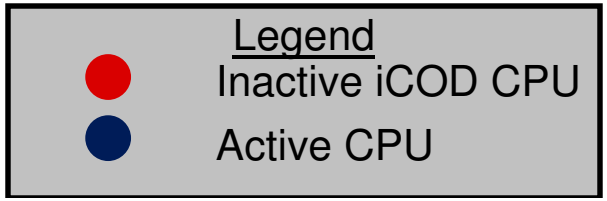

Remote Monitoring GUI

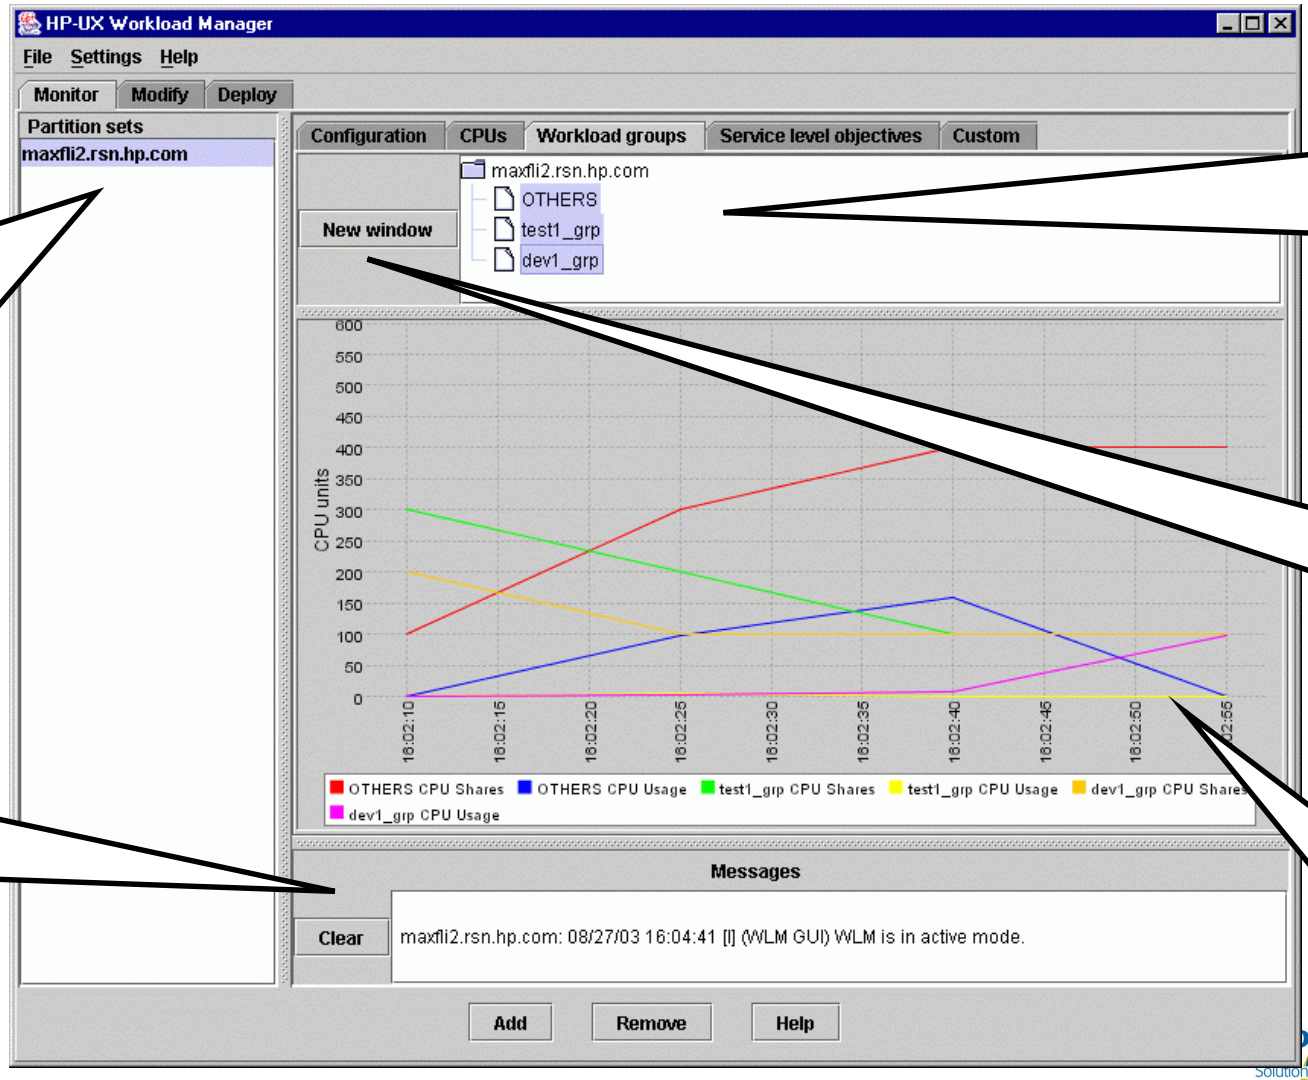

User Defined Sets of HP-UX partitions running WLM

User selected partitions for each graph

Opens the graph in a separate window

Remote WLM messages displayed here

Graph data displayed here

Remote Monitoring GUI

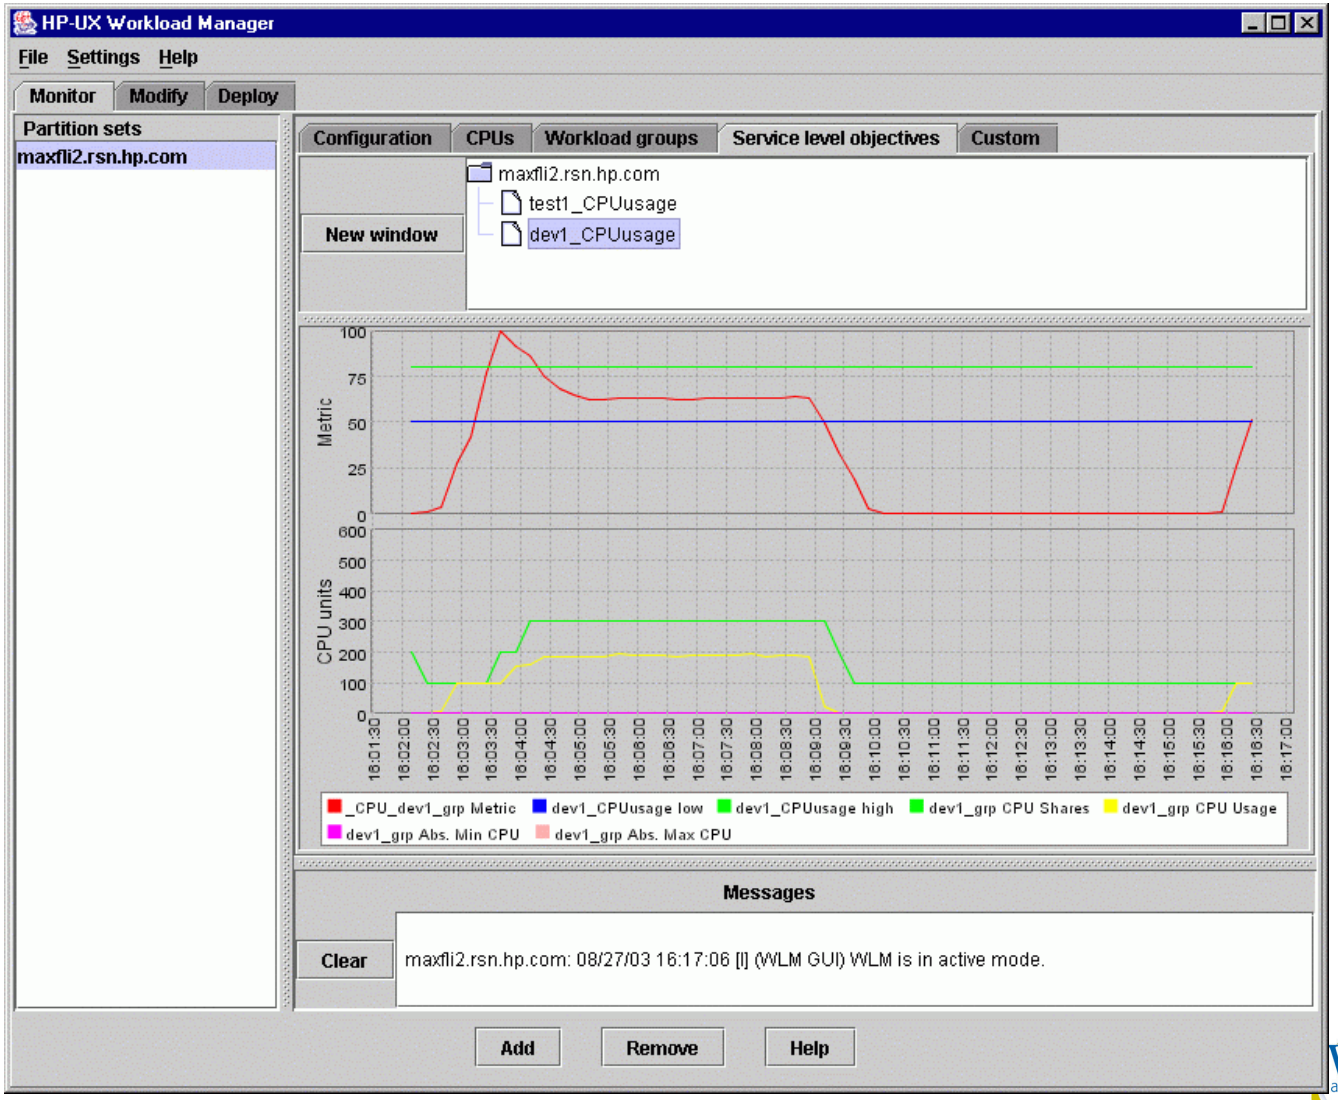

Serviceguard

HP Serviceguard

Foundation high availability clustering
Loosely-coupled cluster

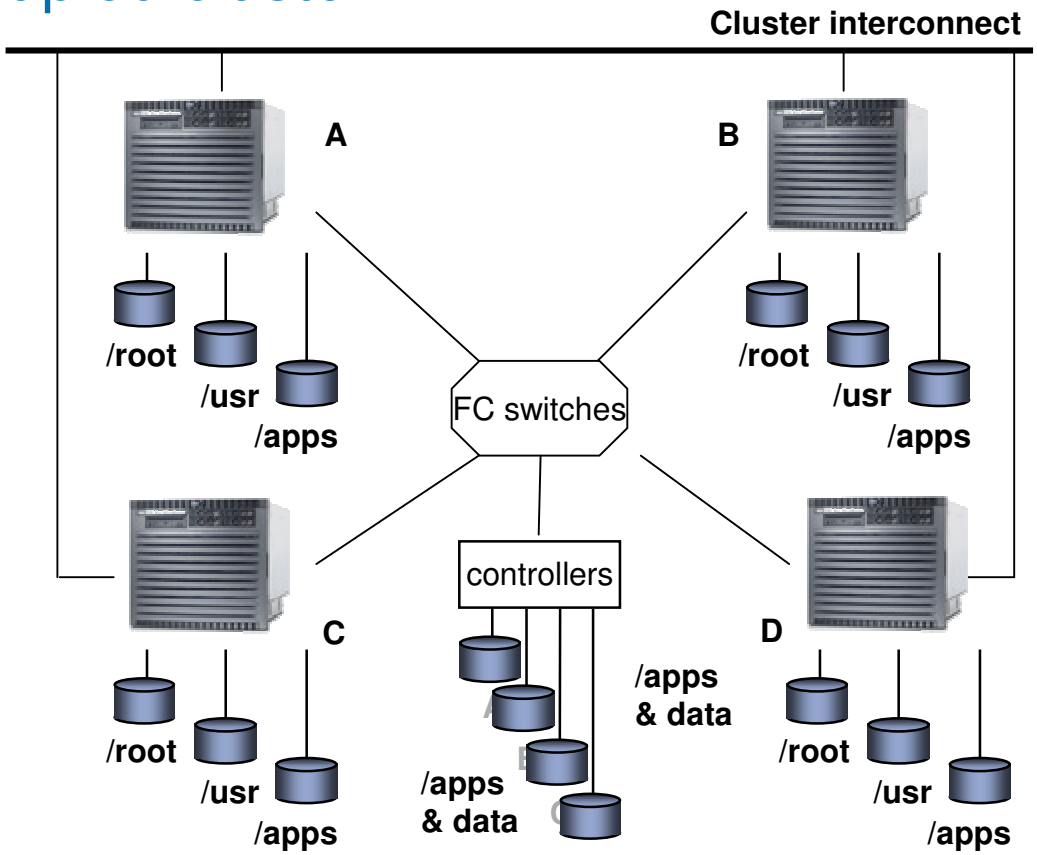

Up to 16 nodes

Each node has its own root file system.

Applications may be installed once or on each node.

Shared centralized or shared distributed storage
file systems (exclusive access) or raw volumes (shared access with SGeRAC)

HP Serviceguard Manager

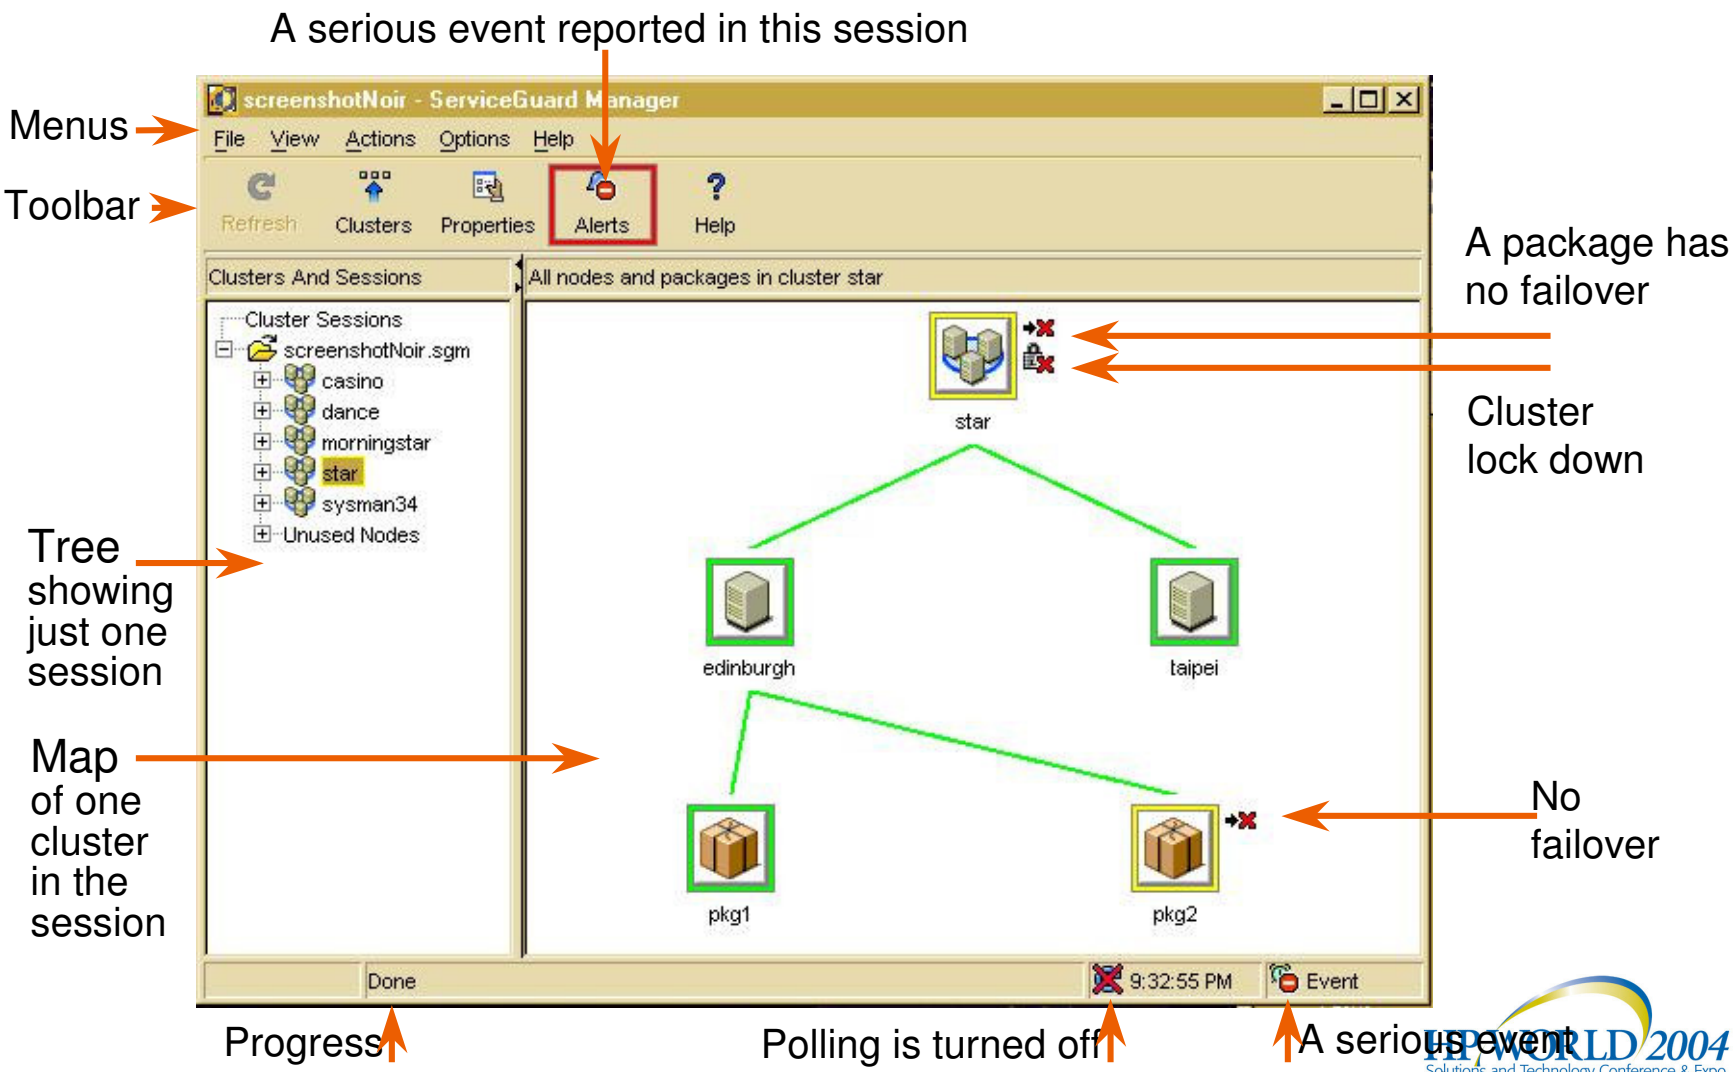

The Virtual Server Environment in Action

HP Virtual Server Environment in action: Optimized utilization in a clustered environment

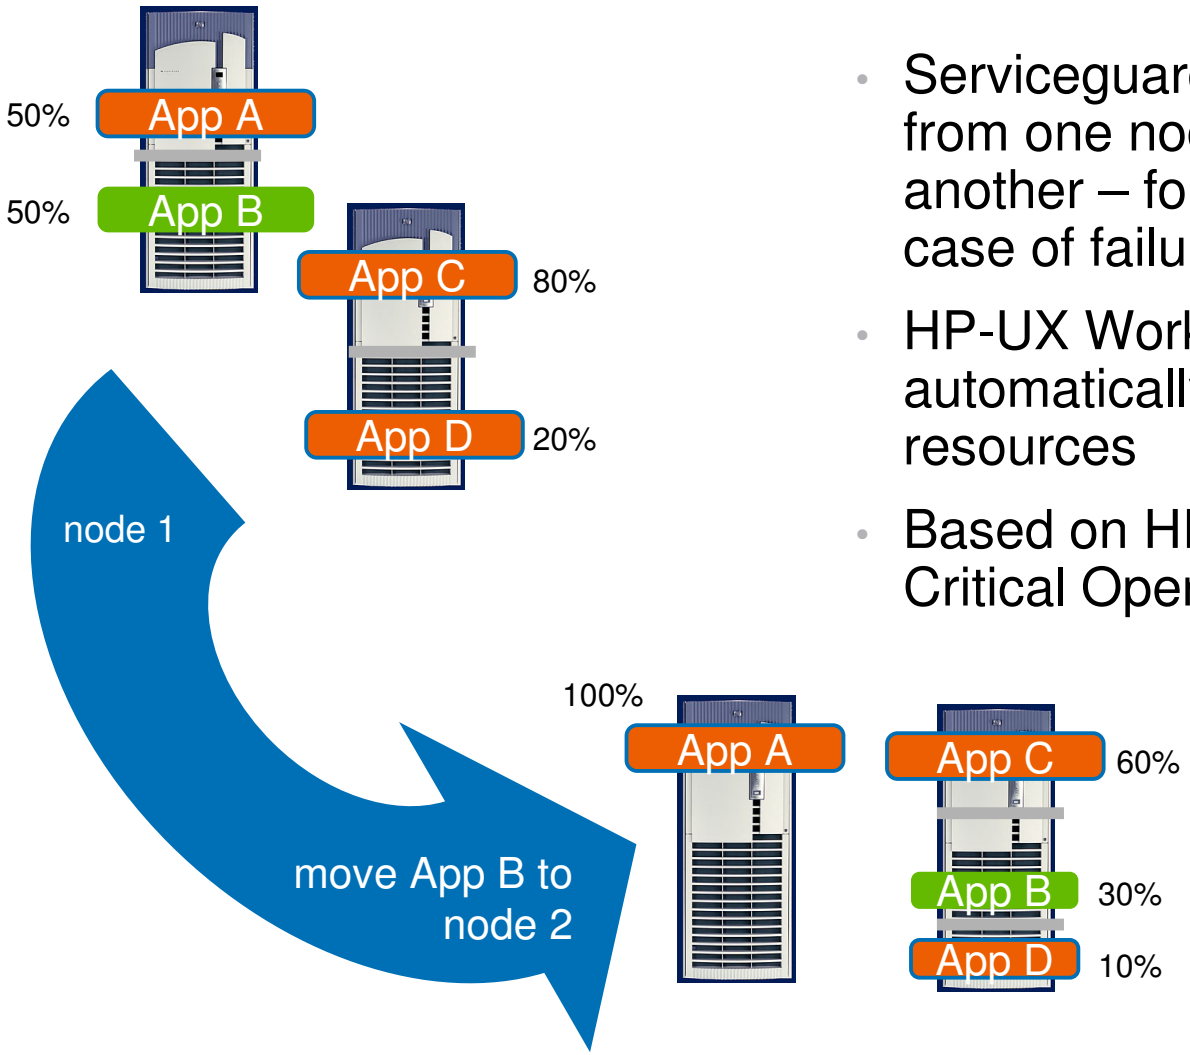

- Serviceguard package transfers from one node or partition to another – for maintenance or in case of failure
- HP-UX Workload Manager automatically reallocates resources
- Based on HP-UX 11i Mission Critical Operating Environment

Resource management of your adaptive infrastructure

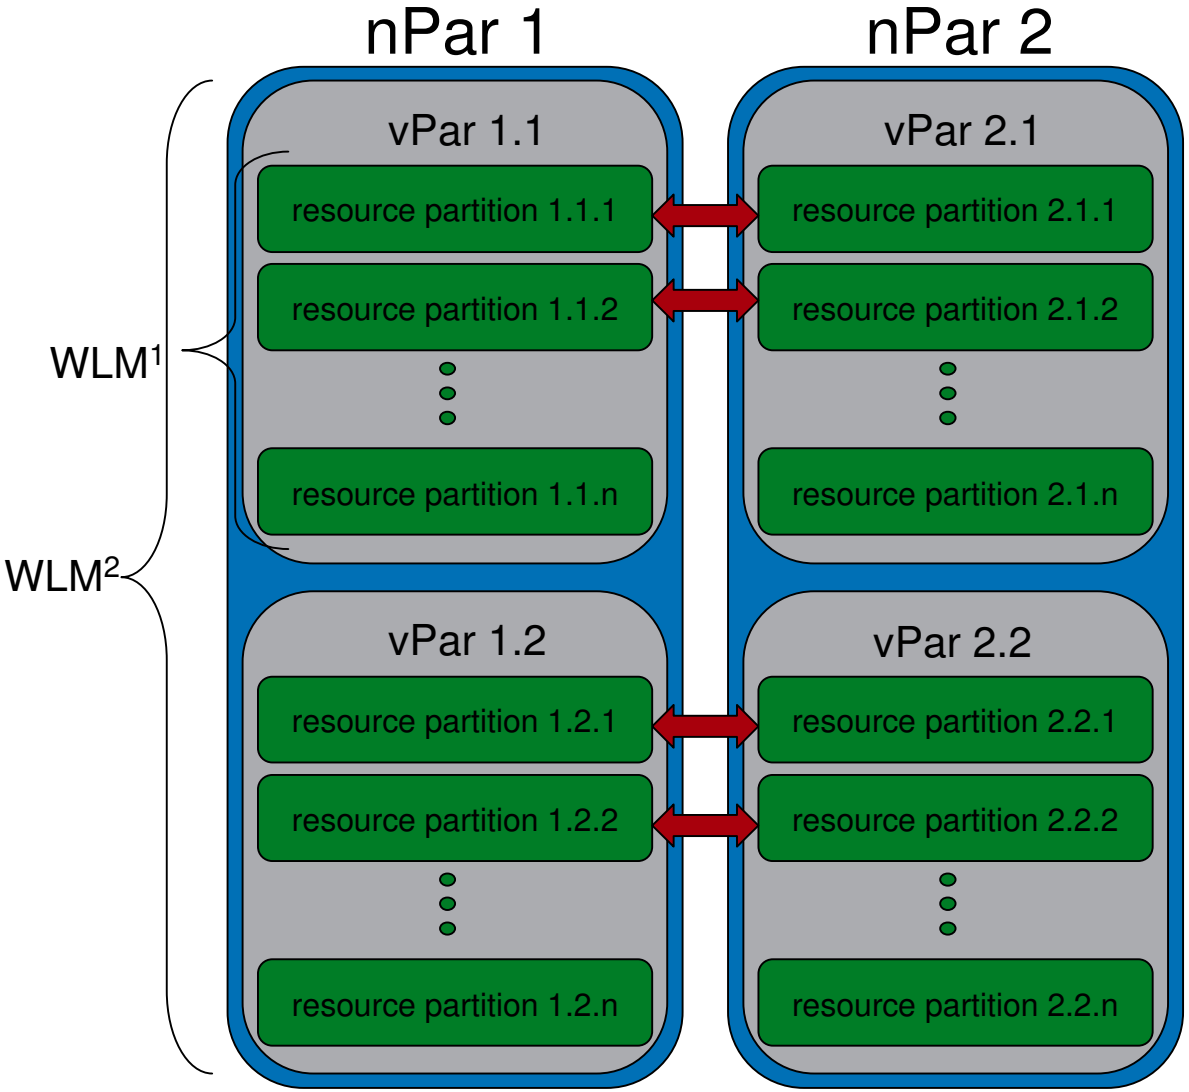

- 2 nPars provides
  - hardware fault isolation
- 2 vPars within each nPar provides
  - software fault isolation
  - OS version isolation
- Any number of resource partitions (one for each major application, or group of same priority minor applications) in each vPar provides:
  - resource isolation
- WLM[1] automatically allocates CPU resources as needed to resource partitions
- WLM[2] automatically allocates CPUs as needed to vPars
- Failover across nPar boundary (indicated by ⟷) provides
  - HA for both hardware and software faults
  - WLM will reallocate resources upon failover

The Next Generation

Introduction to the Global Workload Manager (gWLM)

gWLM Vision

gWLM enables utility computing by supporting centralized resource-sharing policies (definition, enforcement and monitoring) on HP servers.

- Across different partitioning and virtualization approaches
- Across multiple systems – datacenter wide
- Supporting different operating systems (HP-UX, Linux, Win)
- For scale up and scale out utilities

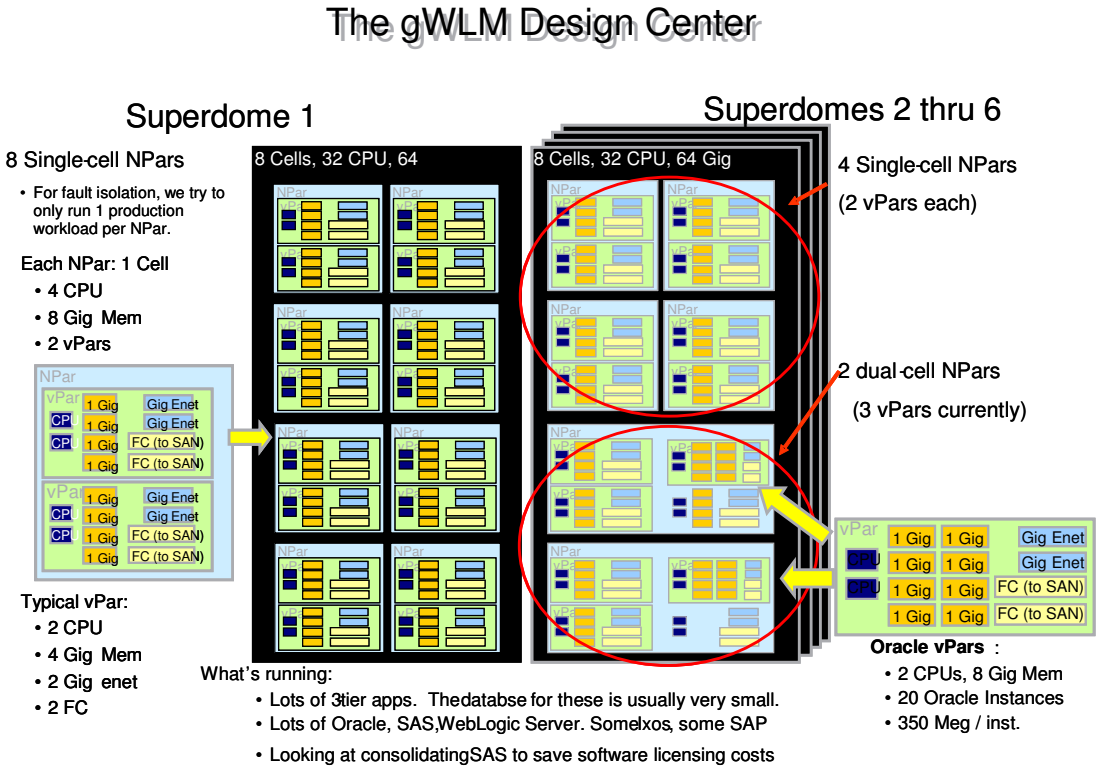

- Simplicity through pre-defined policies, central management, ease of use
- Increased agility by dynamically allocating server resources on-the-fly to meet customer's changing demands
- Delivering additional value by enabling more efficient utilization of shared servers

Ease of Use

- Easy to understand & configure
  - Discovery of the infrastructure
  - Out of the box support for most common use cases
    - Fixed, OwnBorrow, Utilization Policies
    - Easier to understand how much CPU each workload will get
  - UI to visualize what it's doing and why
  - Configuration Wizard
- Quick time-to-value
  - Factory supplied policies
  - Canned reports
- Easy to manage ***very*** large numbers of workloads
  - A single policy can be applied to a large number of workloads
  - Most sites will only need a handful of policies
  - Role-based access allows Sr. staff to define policies and other sys admins can apply and monitor them

Dynamism with Confidence

- Features to help you ensure that the tool is configured as planned and will behave as expected
  - Advisory mode
  - Realtime Reporting
  - Historical reports
- Help application owners trust dynamism
  - CPU resources can be guaranteed
    - Fixed Policy
    - OwnBorrow (PolicyMin,owned)
    - Utilization (PolicyMin)
  - Audit reports to prove each workload got what was guaranteed

gWLM Out of the Box Policies

- OwnBorrow
  - Specify minimum, maximum and “owned” resources
  - gWLM will ensure there is enough resource to satisfy all workloads at the “owned” level
  - Workloads can share unused resources, but are guaranteed access to the owned amount if needed
- Fixed
  - Specify a fixed amount of resource for this workload
  - No sharing will be allowed
- Utilization
  - Specify minimum, maximum and a utilization target
  - If actual utilization is higher than target, policy will request more, if below target policy will request less

gWLM Canned Reports

- Troubleshoot a poorly performing workload
- Get periodic capacity and performance report
- Produce a resource audit report for internal customers
- Police my internal customers – identify resource hogs
- Find capacity for a new workload
- Right-size a workload's entitlement

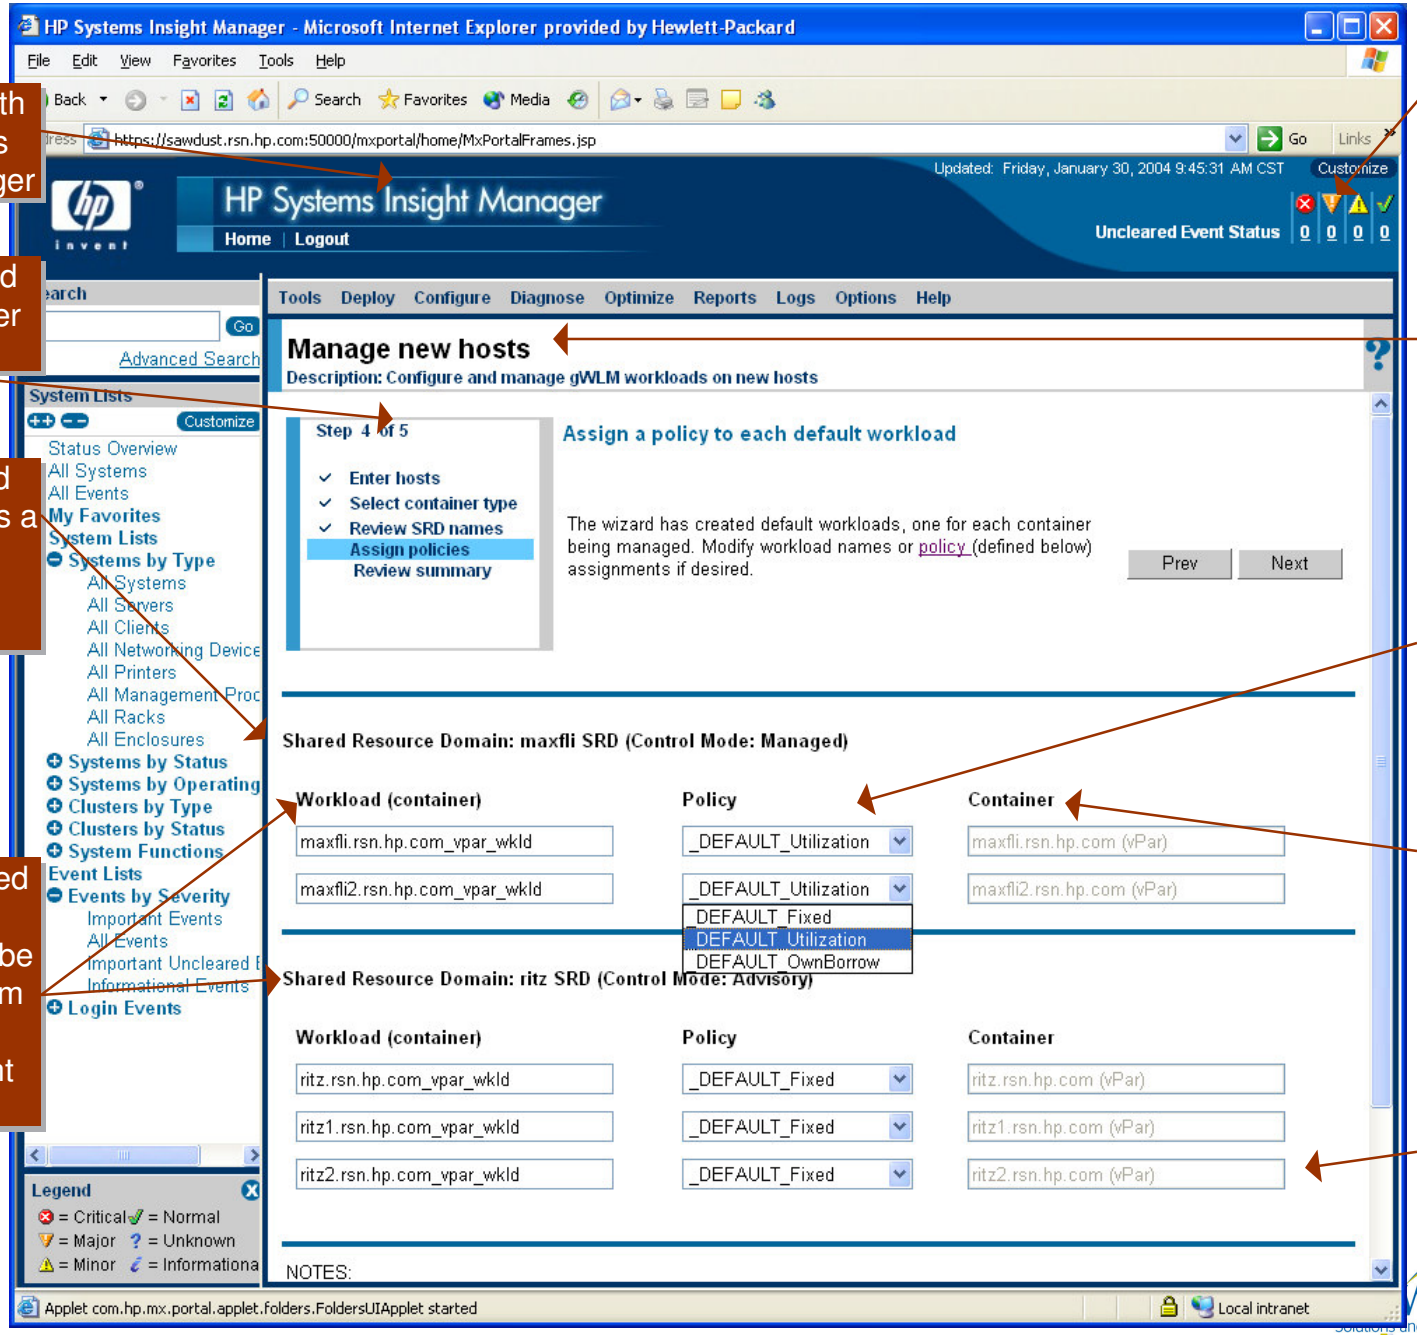

Integrated with HP Systems Insight Manager

Task oriented graphical user interface

Each shared environment is a **Shared Resource Domain**

Multiple shared resource domains can be managed from the central management station

Integrated event system to track SLO violations

Configuration panel, monitoring panels also available

Factory-defined policies simplify configuration of shared domains

Abstraction of different types of containers, vpars, psets, …

Auto-discovery of containers

gWLM – utilization monitoring

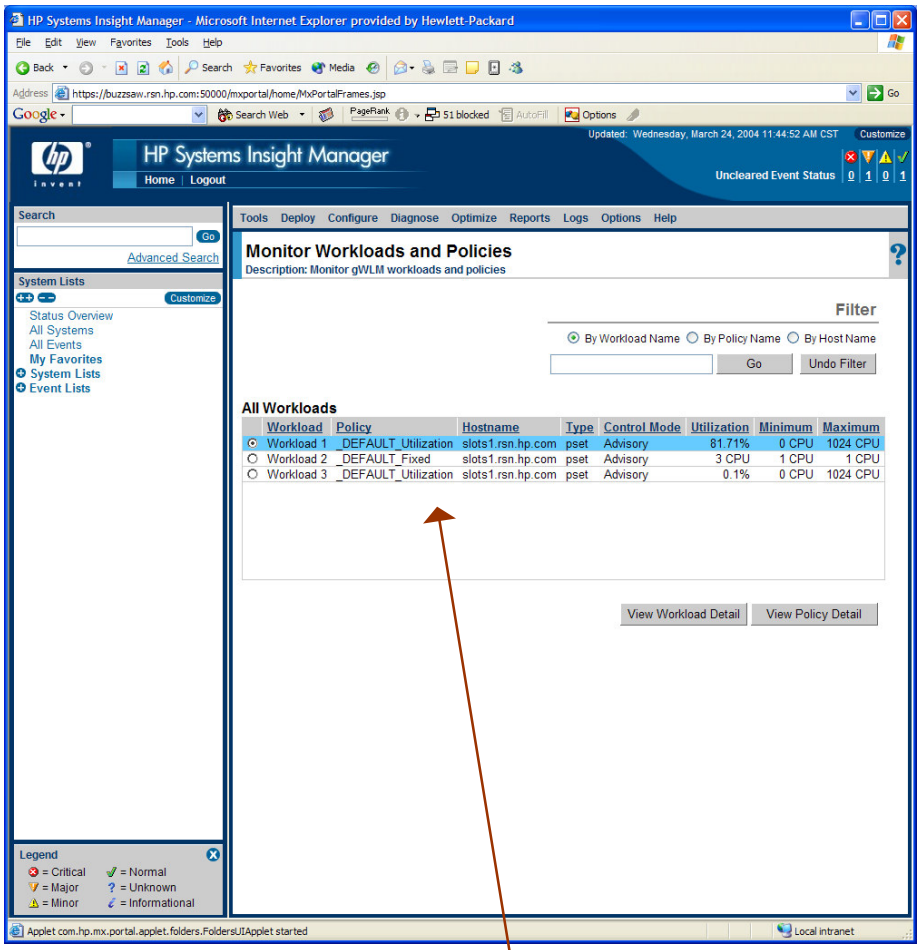

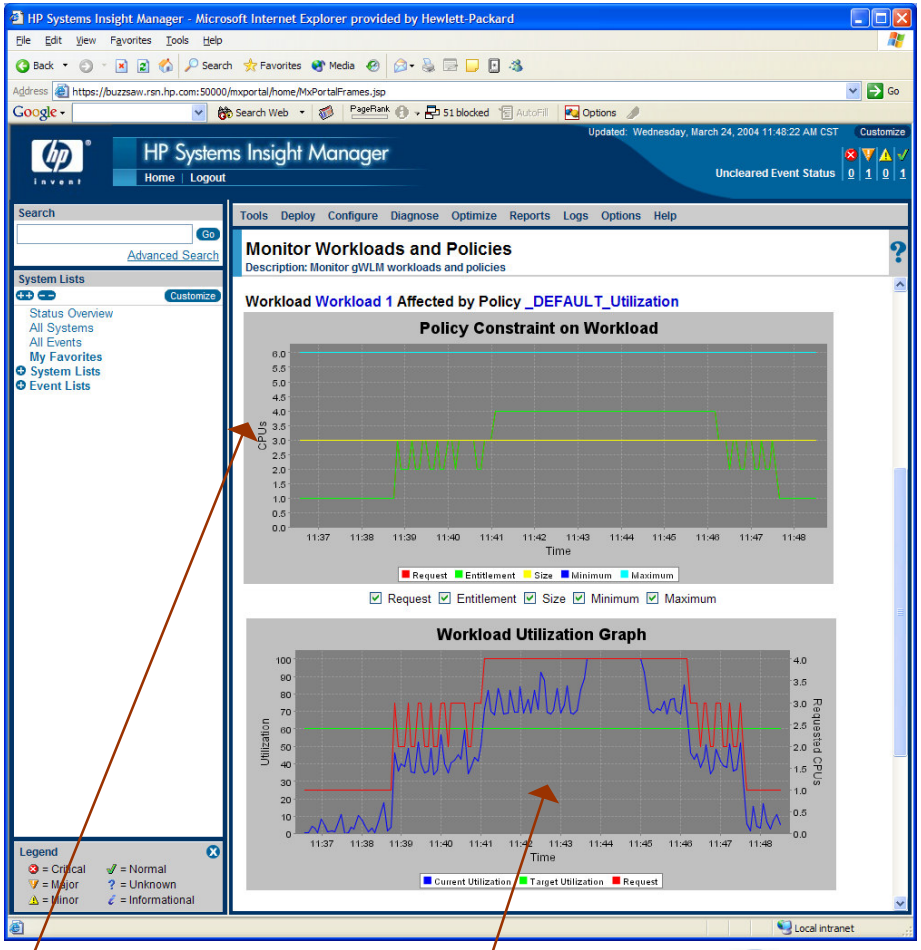

Selection screen for monitoring capabilities with gWLM

Graphs with performance against policy

Real-time and historical utilization graphs for each workload

gWLM Demo

The Demo Scenario

- IT manages a server utility with 20 servers each set up with 3 vPars
- They offer the infrastructure as a managed service to their business units
- They want to increase utilization so they are allowing the sharing of unused resources by the workloads on each server
- Each workload owns a certain number of CPUs and that is how the business units are charged for the service

The Demo System

- 3 vPar-based Workloads on an 8 CPU server
  - A Production Sales Application
    - Owns 3 CPUs – Max 6
  - A Production Finance Application
    - Owns 4 CPUs – Max 6
  - A Dev/Test Environment
    - Owns 1 CPU – Max 6
- We will show 2 demos:
  - Configuring gWLM to manage these workloads
  - Monitoring gWLM's allocation of CPUs as the loads on the applications changes

Getting Started with Workload Manager

Objectives

Learn to:

- Set up initial configuration using the Configuration Wizard
- Use **`wlmd`** to enable WLM
- Monitor WLM

WLM Configuration Wizard

- The Configuration Wizard is useful for creating an initial WLM configuration if you prefer not working directly with the configuration file
- Any editing done to the initial configuration produced by the wizard can be done manually or with the WLM GUI

Starting the Configuration Wizard

```
# /opt/wlm/bin/wlmcw
```

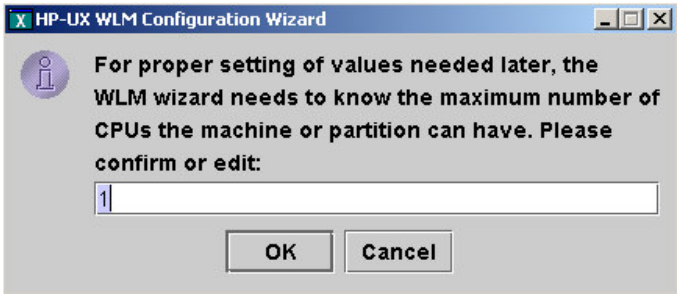

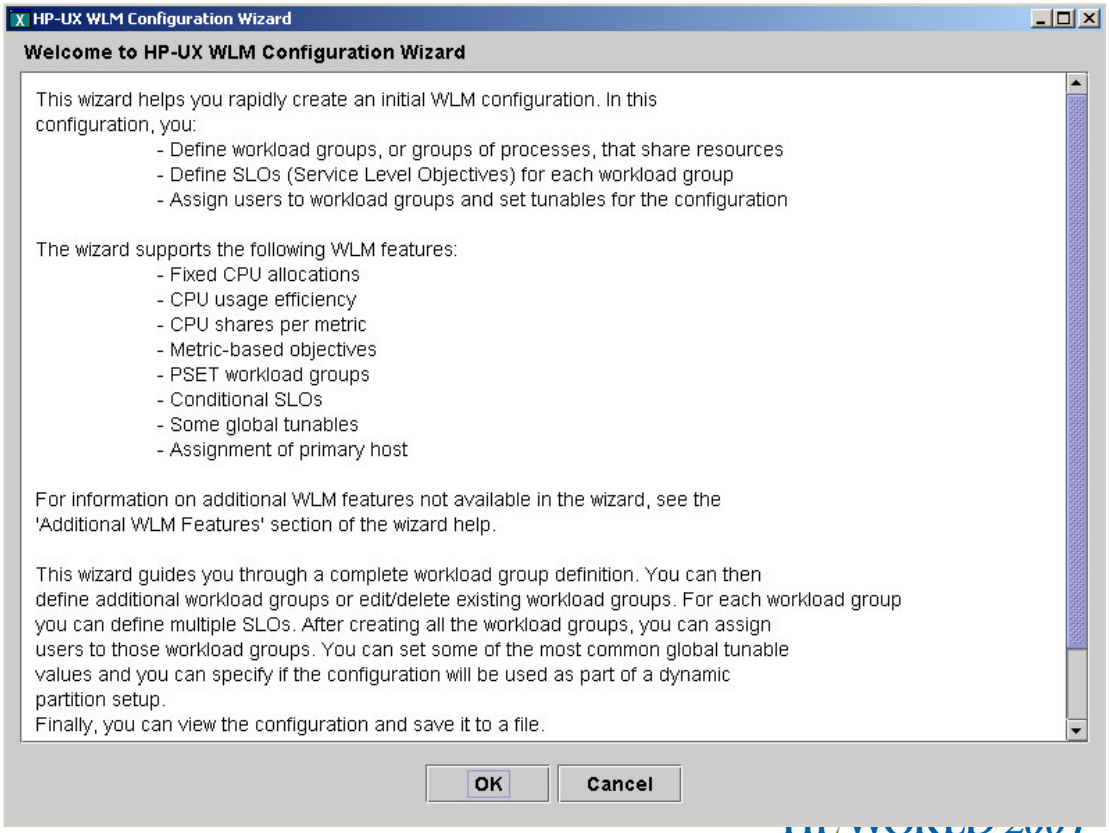

Configuration Wizard steps

1. Define the Workload Group
2. Define SLO
3. Define additional or Edit/Delete existing workload groups or Commit current workload group
4. Define Additional or Edit/Delete Existing Workload Group
5. Assign Users
6. Set Global Tunables
7. Set Primary Host
8. View and Save Configuration

Specifying workload groups

- There are two types of workload groups: FSS and PSET

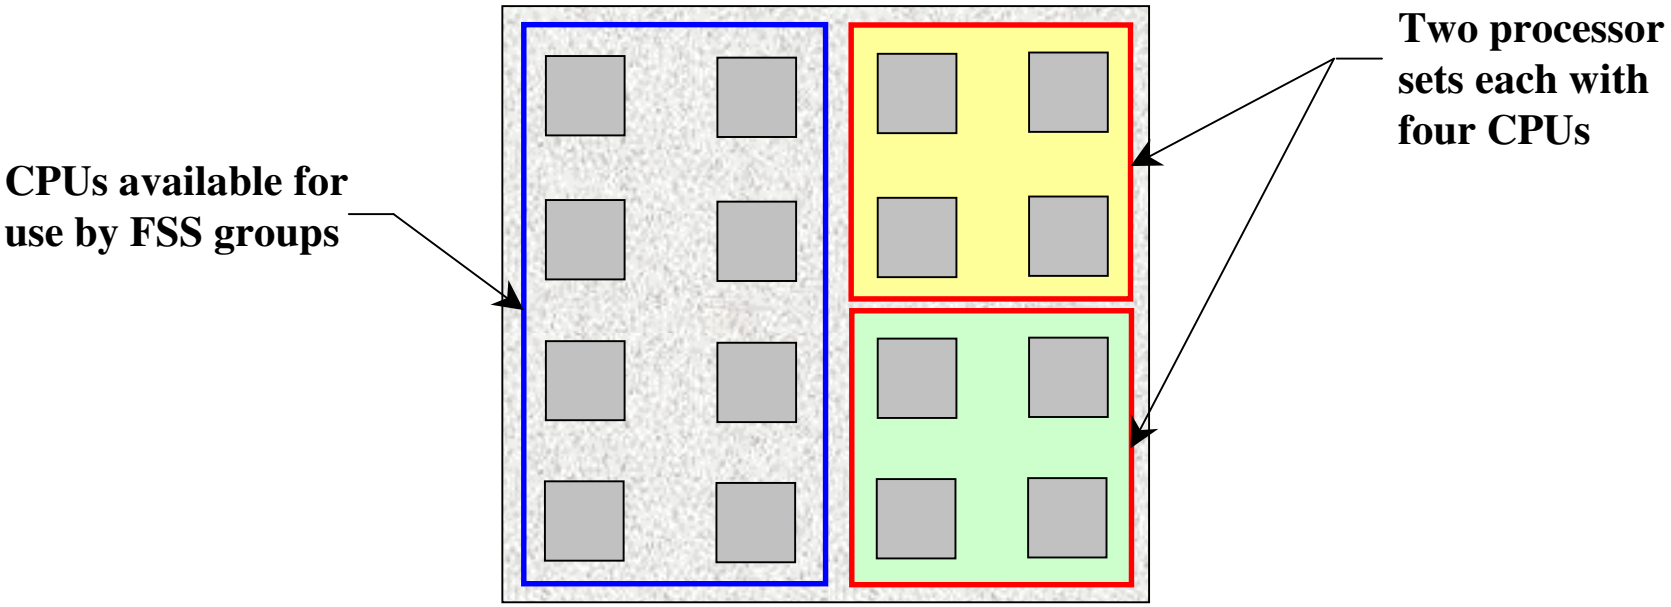

16 CPU Muli-processor system

Reserved workload groups

`PRM_SYS` (ID 0)

- Default workload group for system processes
- Can automatically use up to 50% CPU
- The CPU cycles that **`PRM_SYS`** doesn't use make up the 100% that is divided among workload groups according to their allocations
- **`PRM_SYS`** is created automatically but can be explicitly defined in the WLM configuration file

`OTHERS` (ID 1)

- Default workload group for non-system processes
- Receives any leftover CPU resources by default
- Created automatically but can be explicitly defined
- Must be defined in the WLM configuration file if memory or disk bandwidth allocations are specified

Defining the workload group

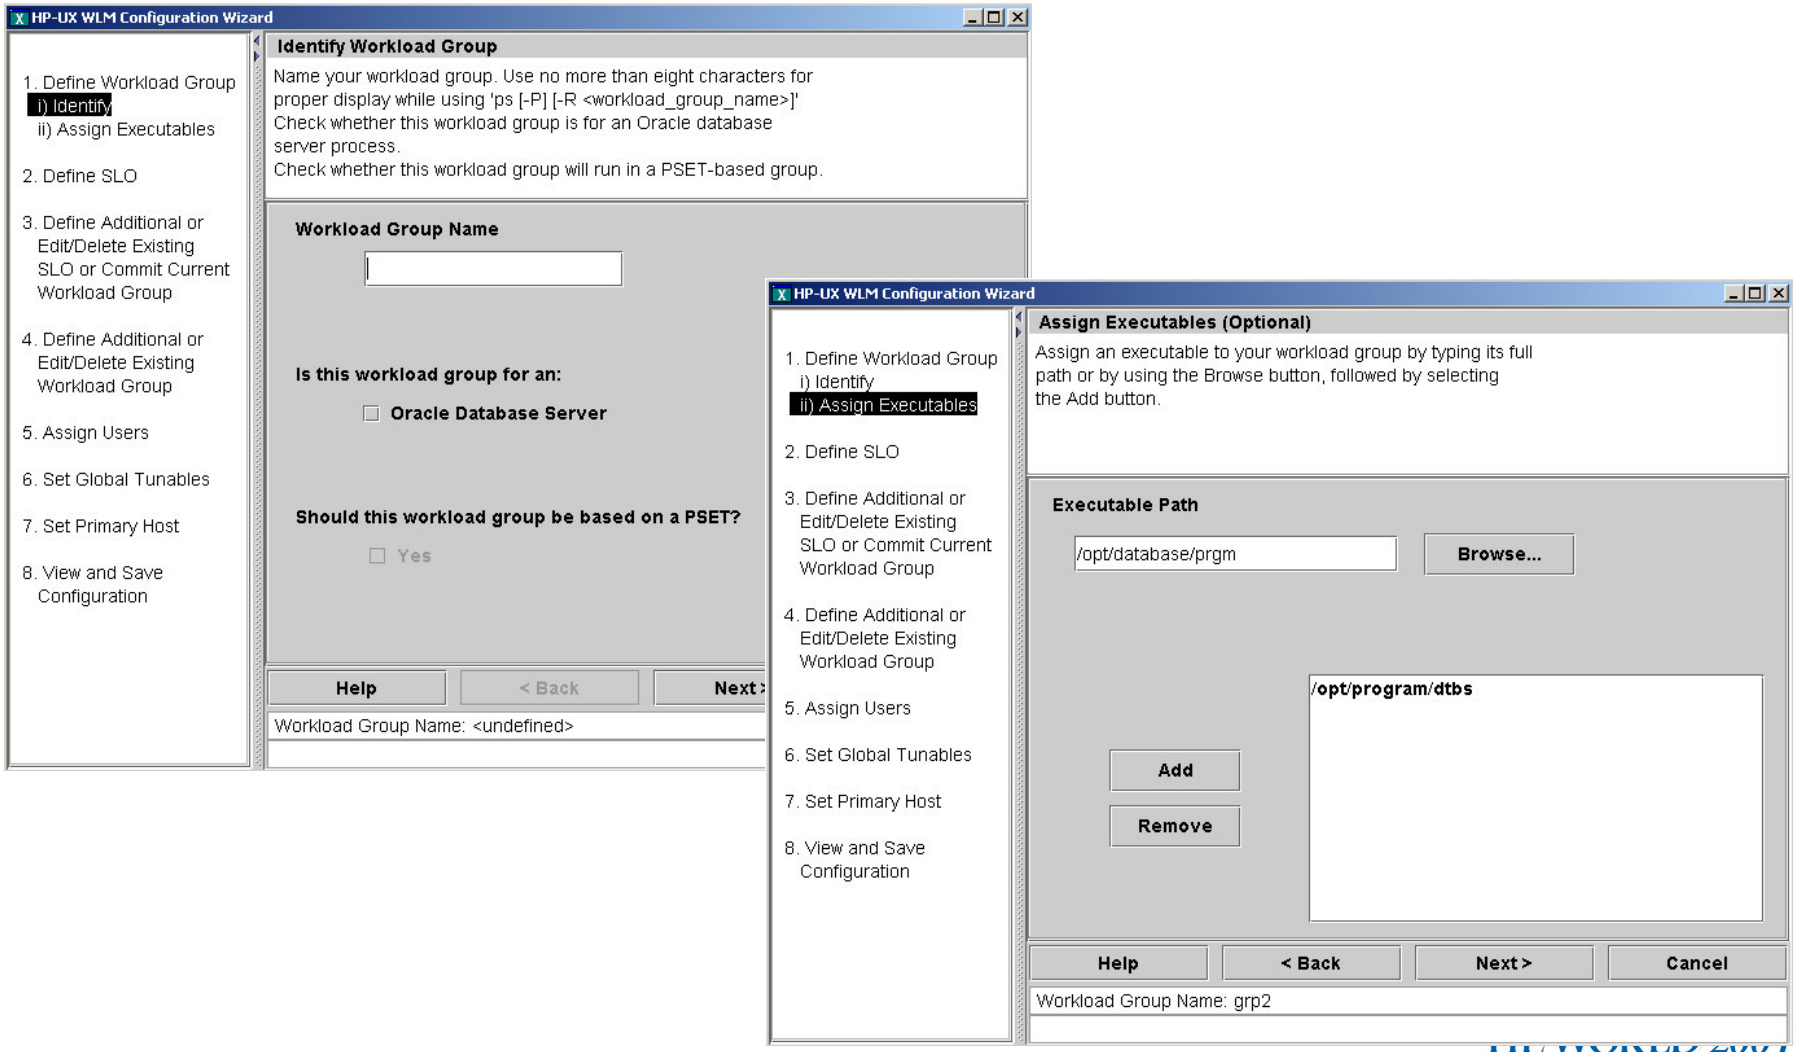

Defining the SLO

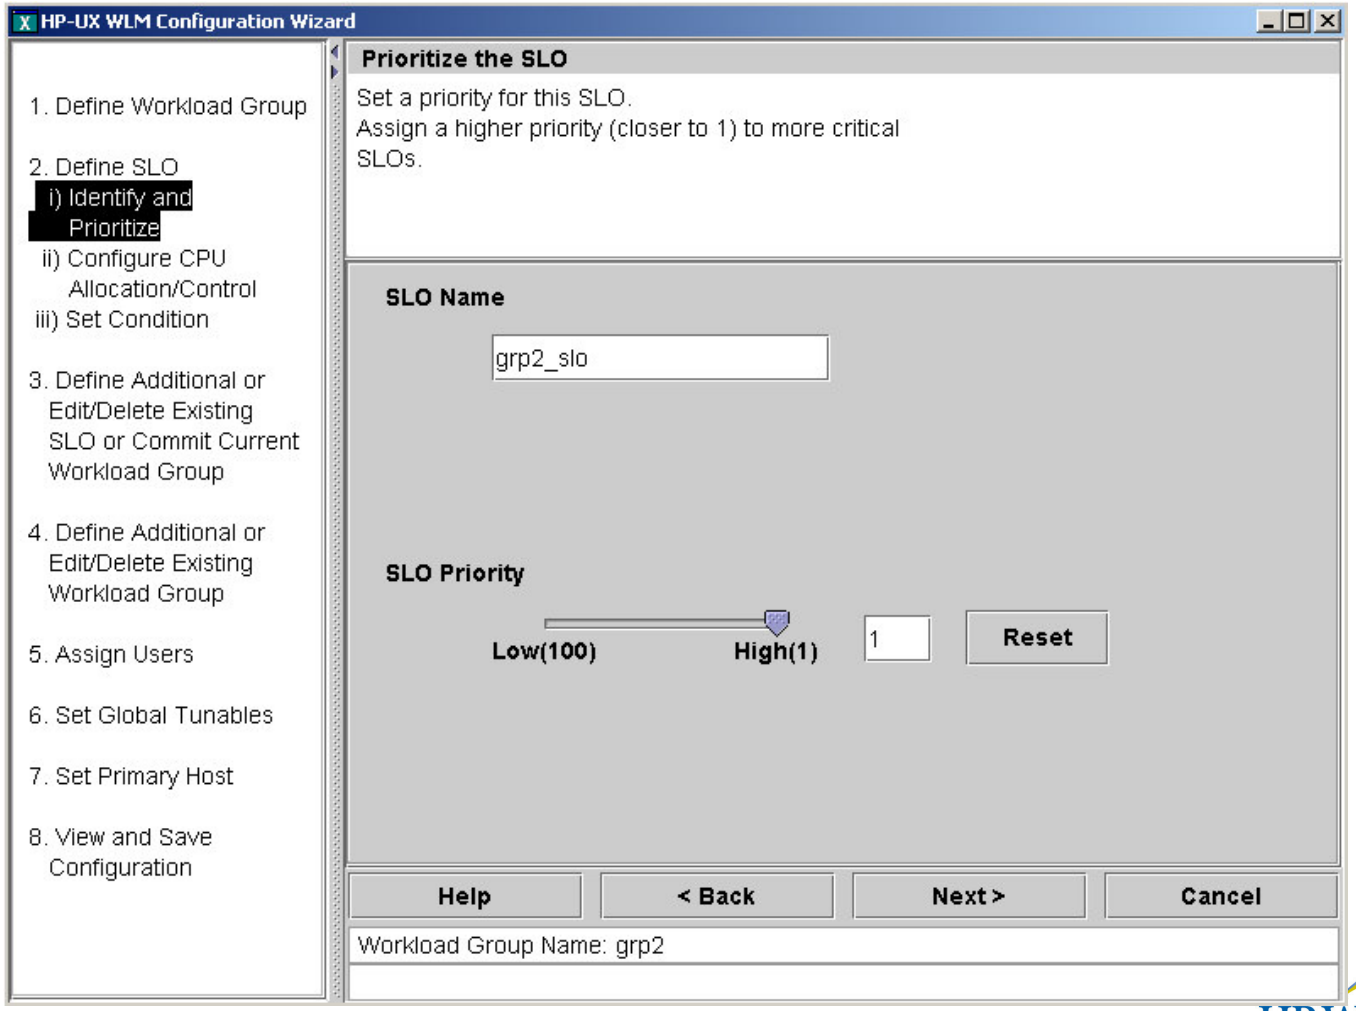

Non-Metric Based CPU allocation (1/2)

- Fixed Allocation:
  - Has explicit allocation request
  - No goal is specified
- CPU usage:
  - Usage goals specify a CPU utilization range
  - Has maximum and minimum CPU request bounds
  - WLM internally tracks the workload group's actual CPU usage versus its CPU allocation
  - Useful when you want to automatically size the allocation based on what a workload needs

Non-Metric Based CPU allocation (2/2)

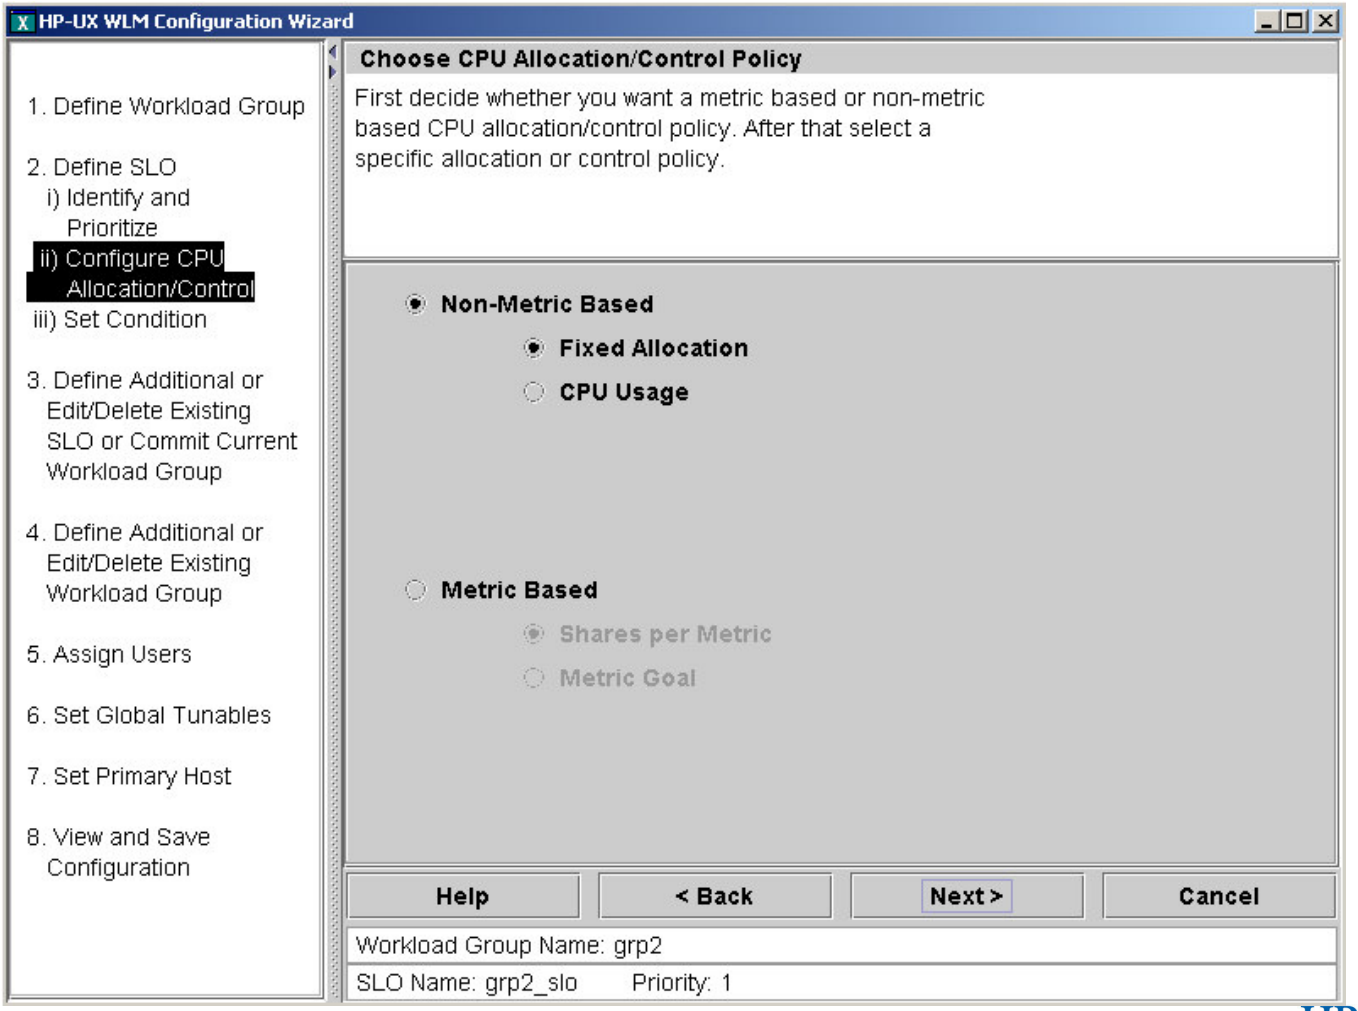

Metric Based CPU allocation (1/2)

- Shares per Metric
  - Specifies how many CPU shares a workload group gets per some metric
  - The CPU shares value is later multiplied by the metric to determine the workload group's allocation
- Metric Goal
  - Specifies a metric for the workload group and some value that the metric should be less than or greater than
  - WLM then manages the workload group's CPU allocation to keep the metric less than or greater than the value

Metric Based CPU allocation (2/2)

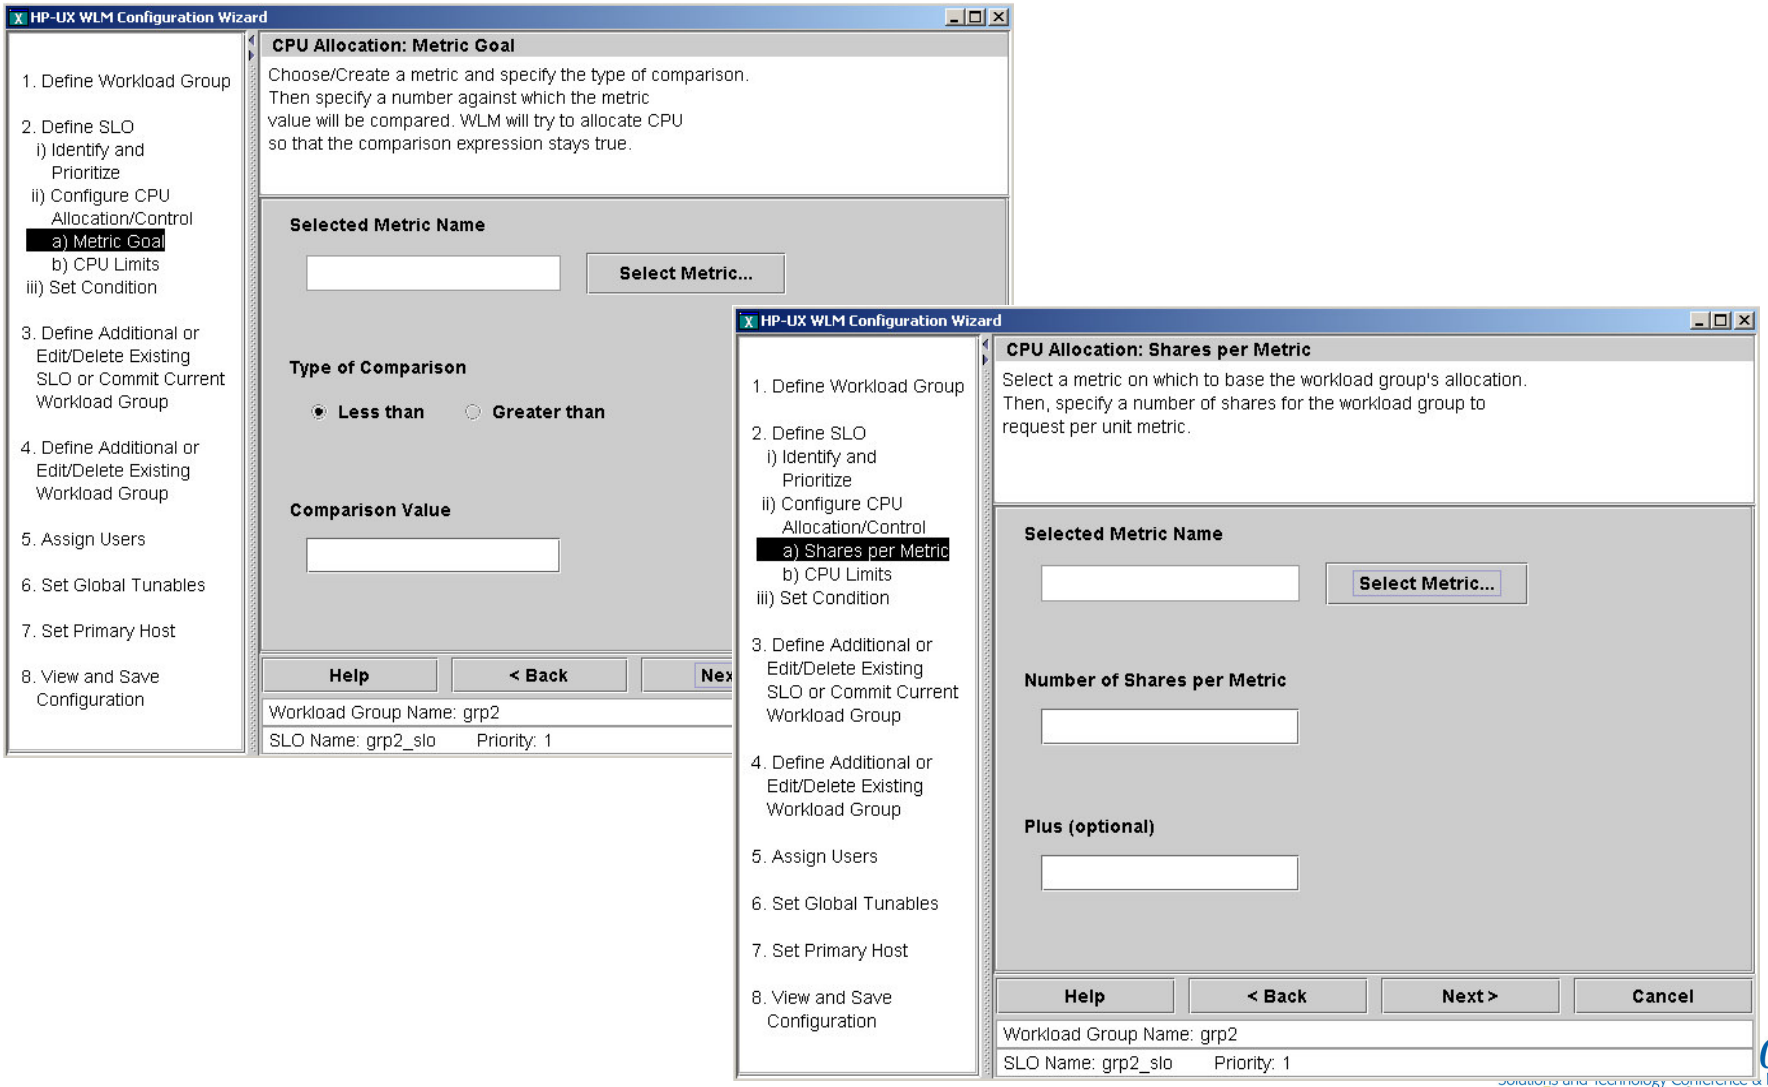

Metric Sub-wizard

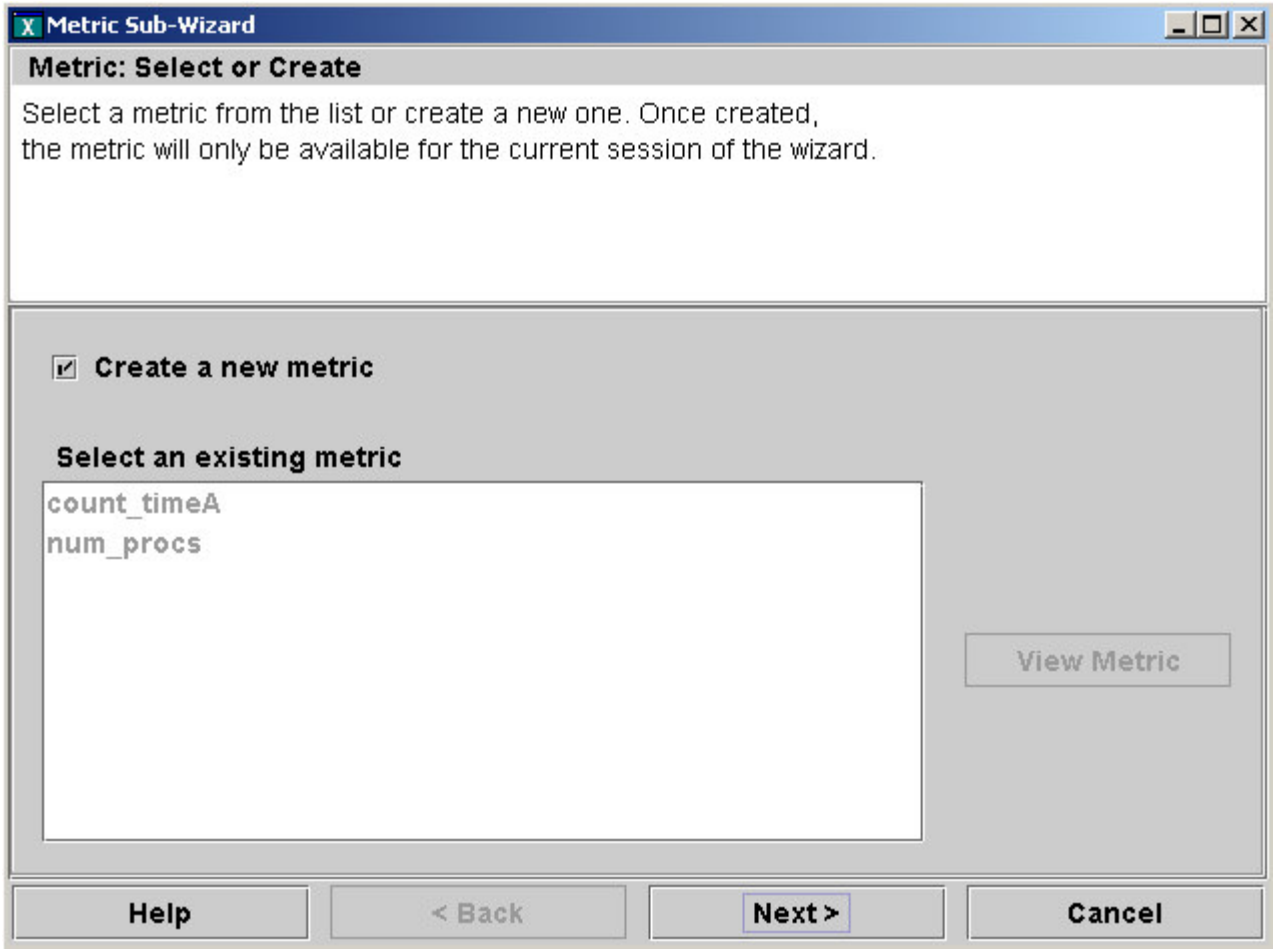

Specifying the metric

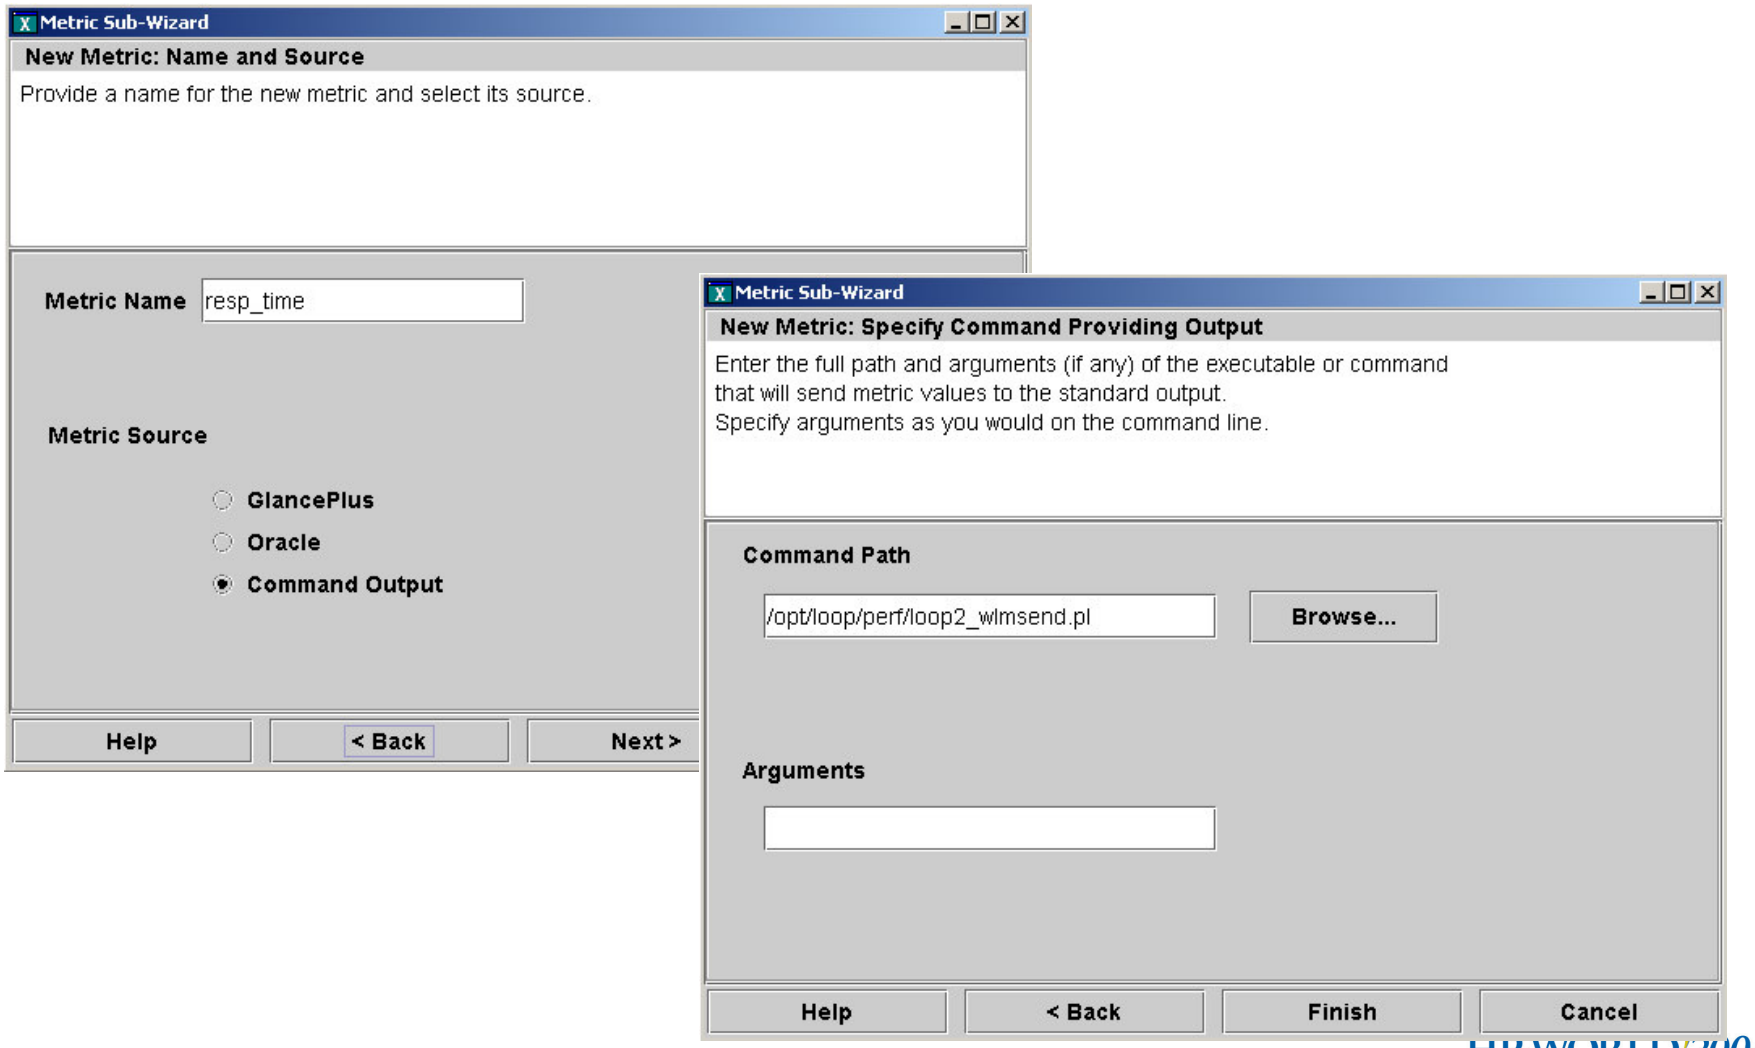

Specifying a condition or Serviceguard package

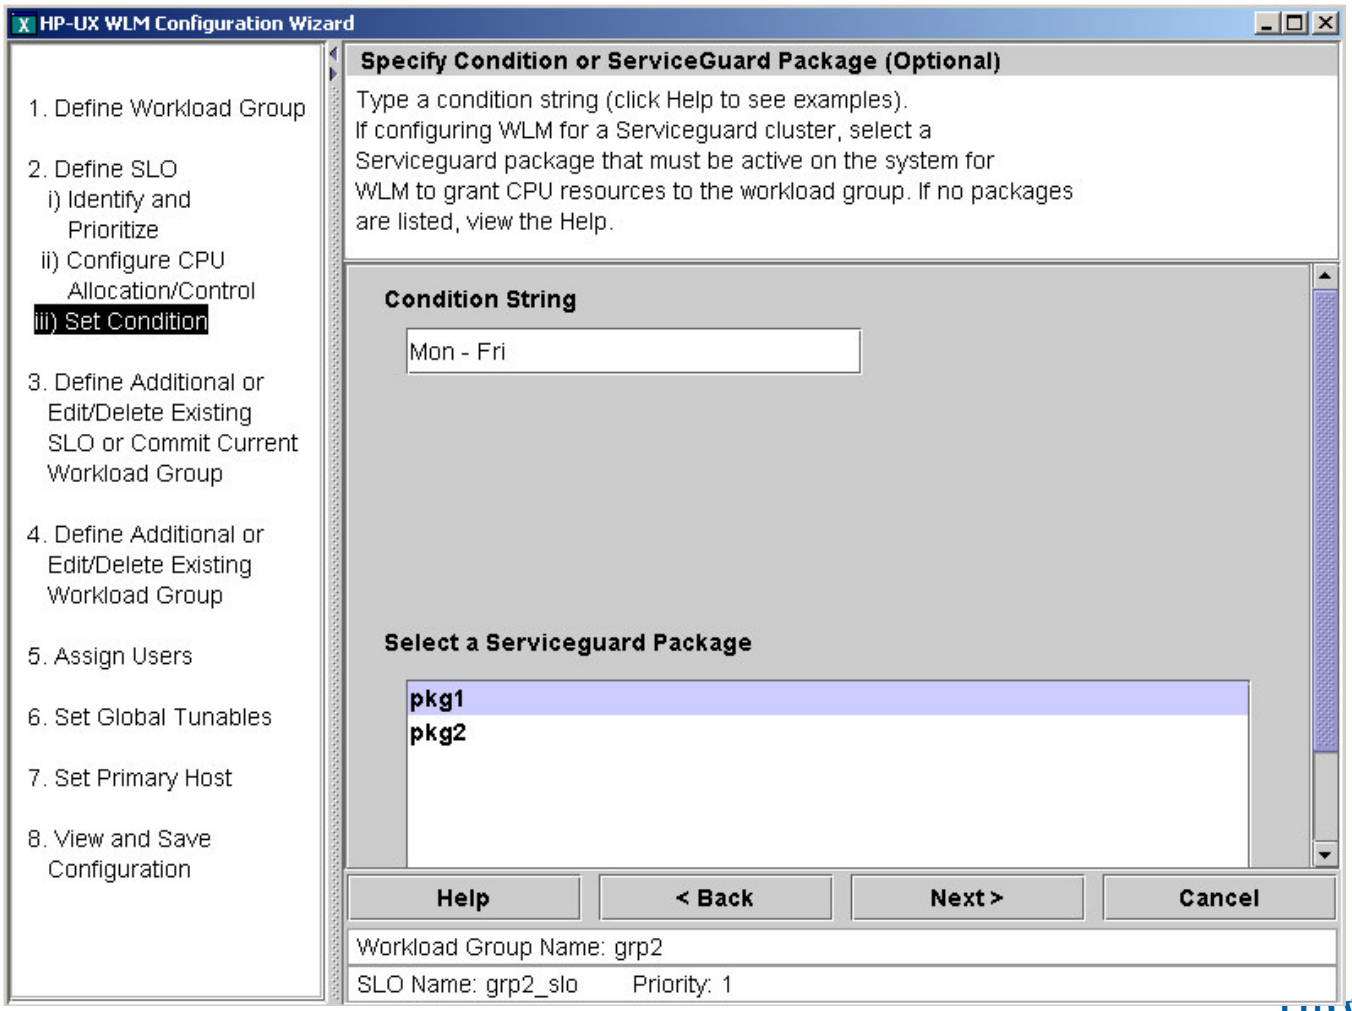

Add / Edit / Delete Workloads

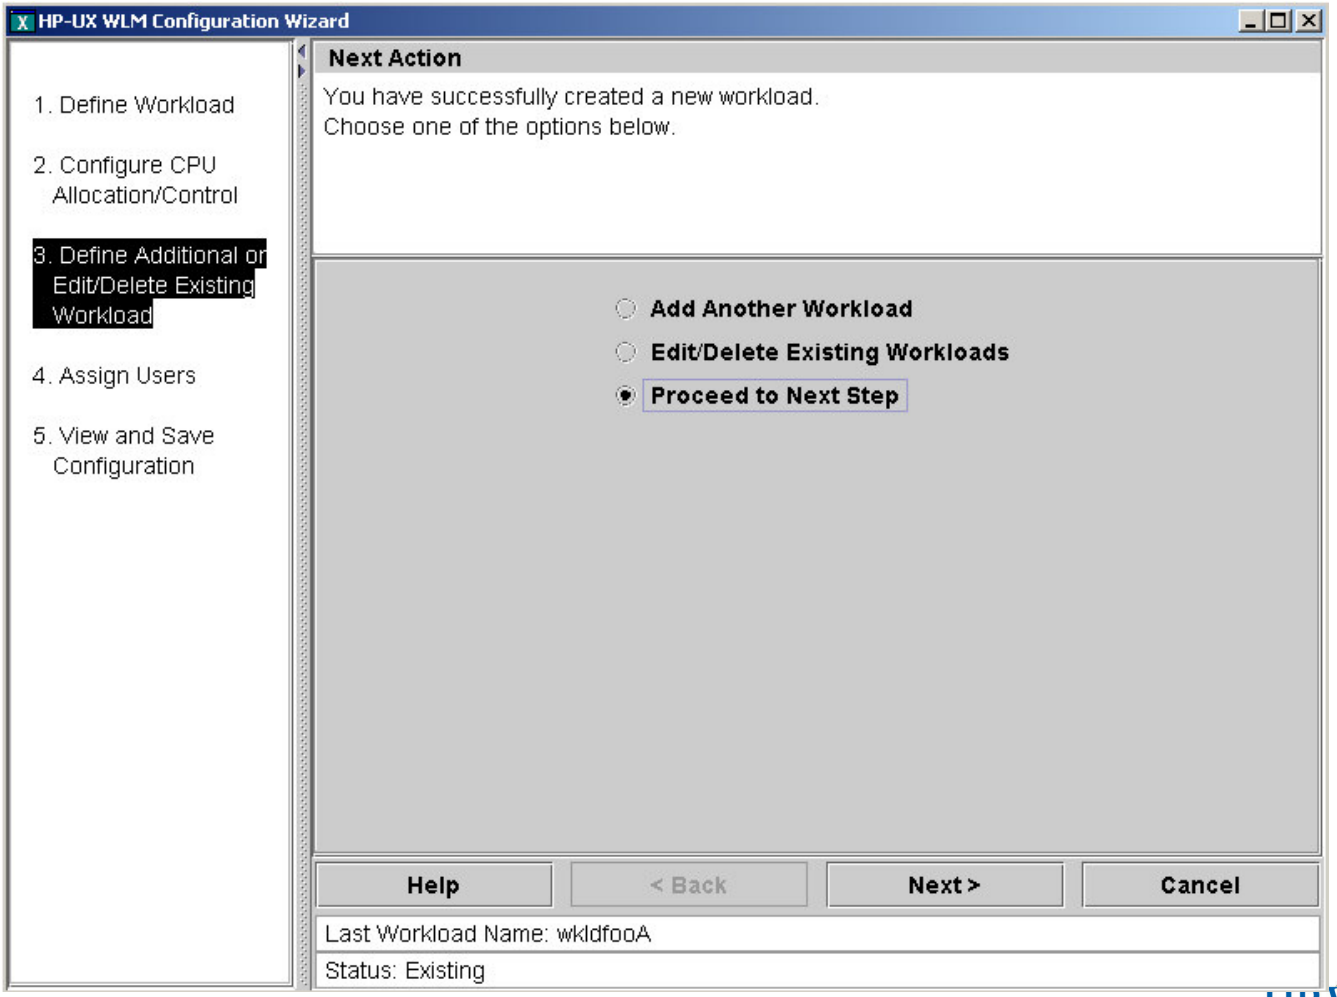

Assigning users to workloads

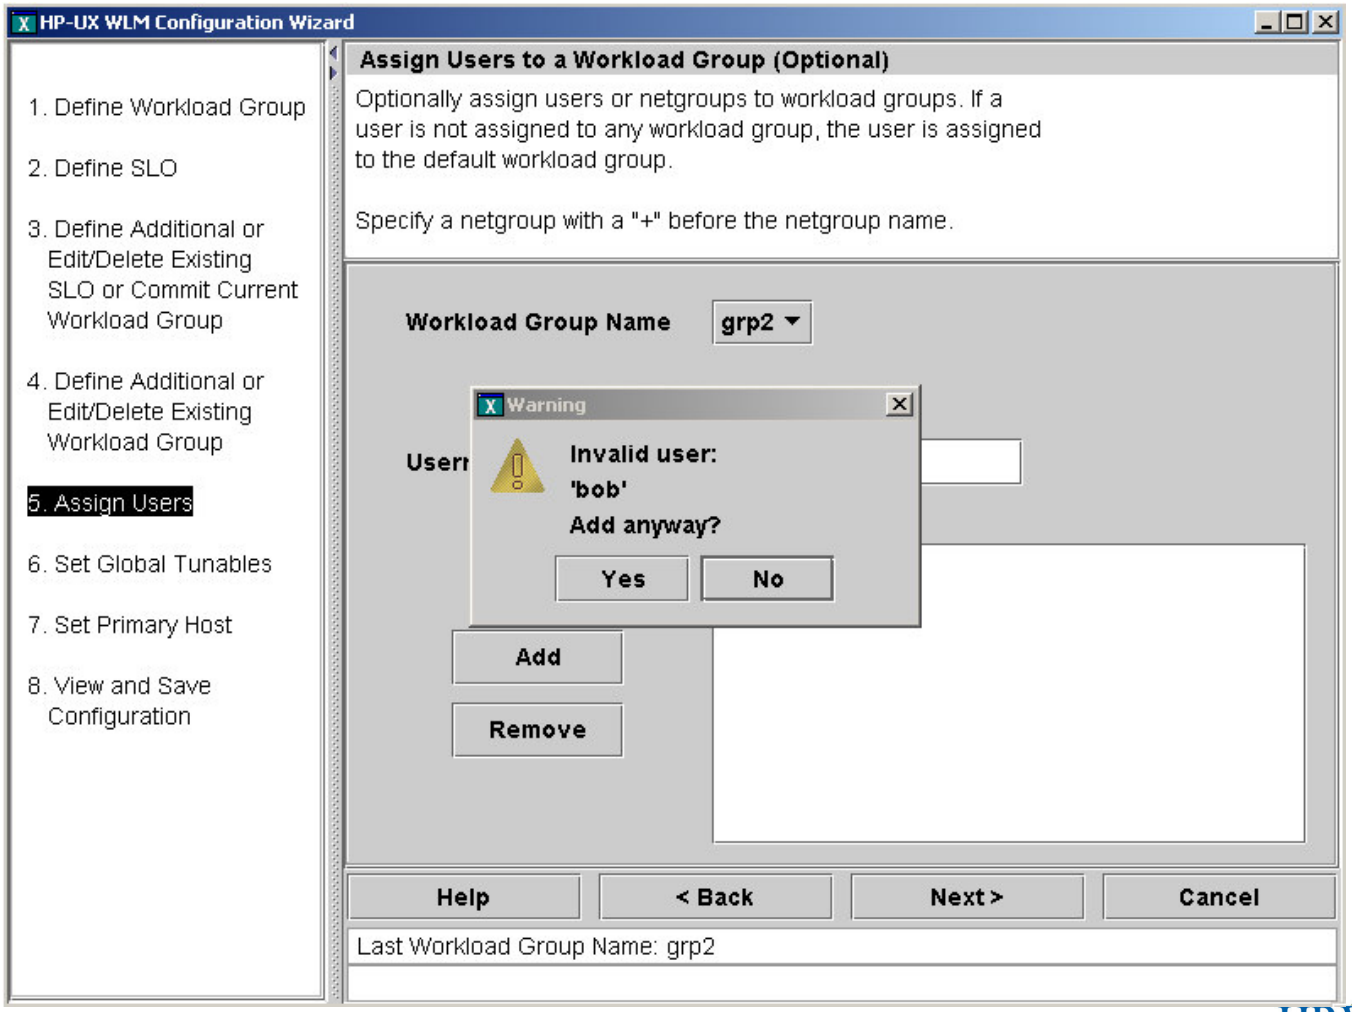

Adding global tunables

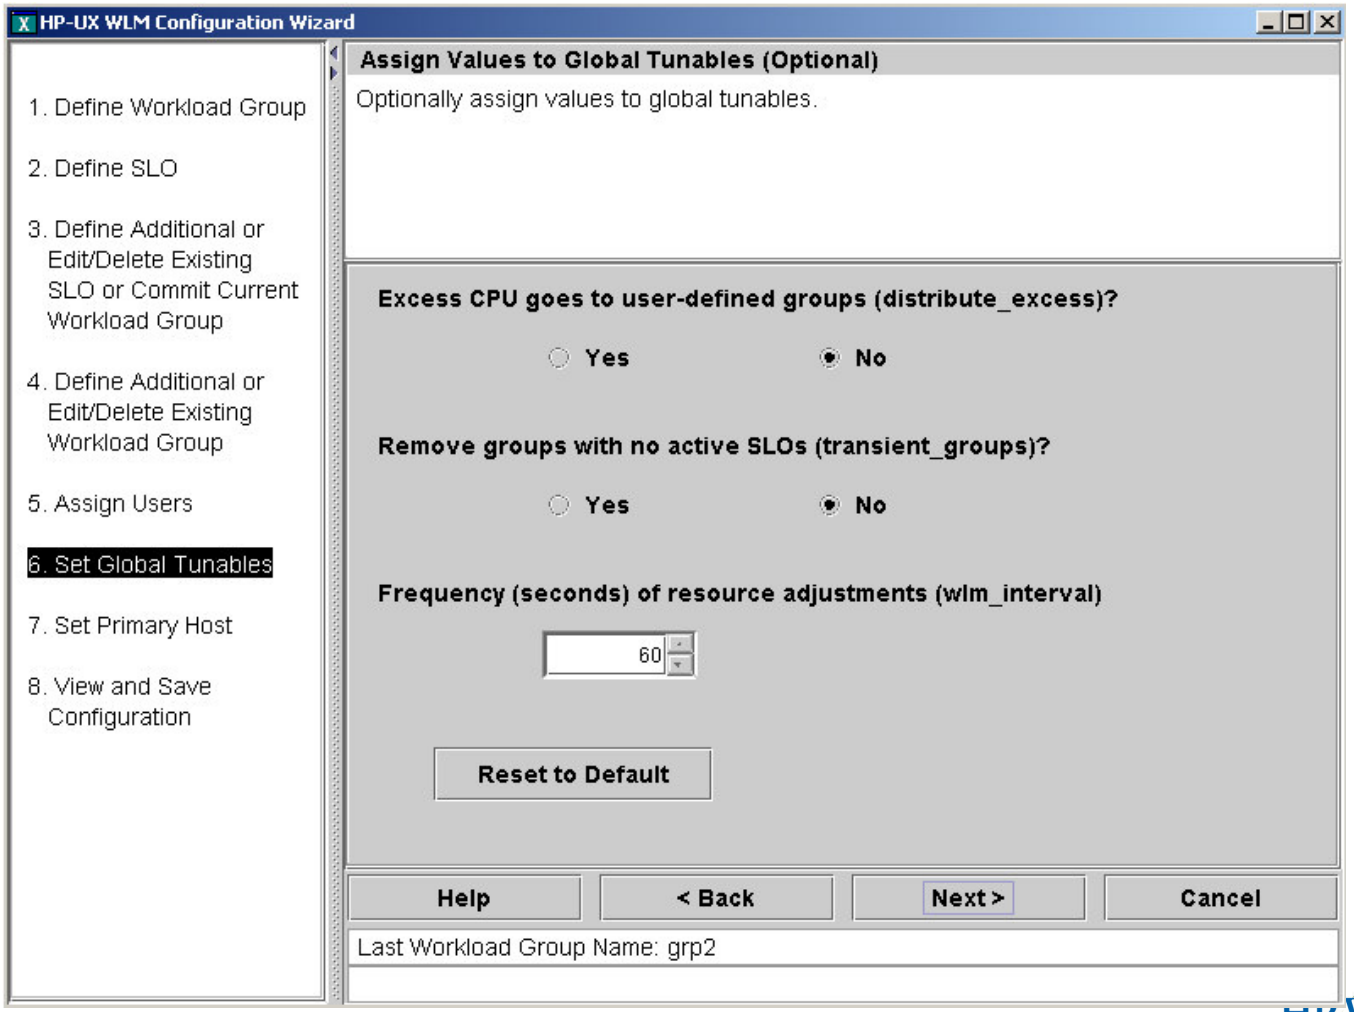

Specifying the primary host

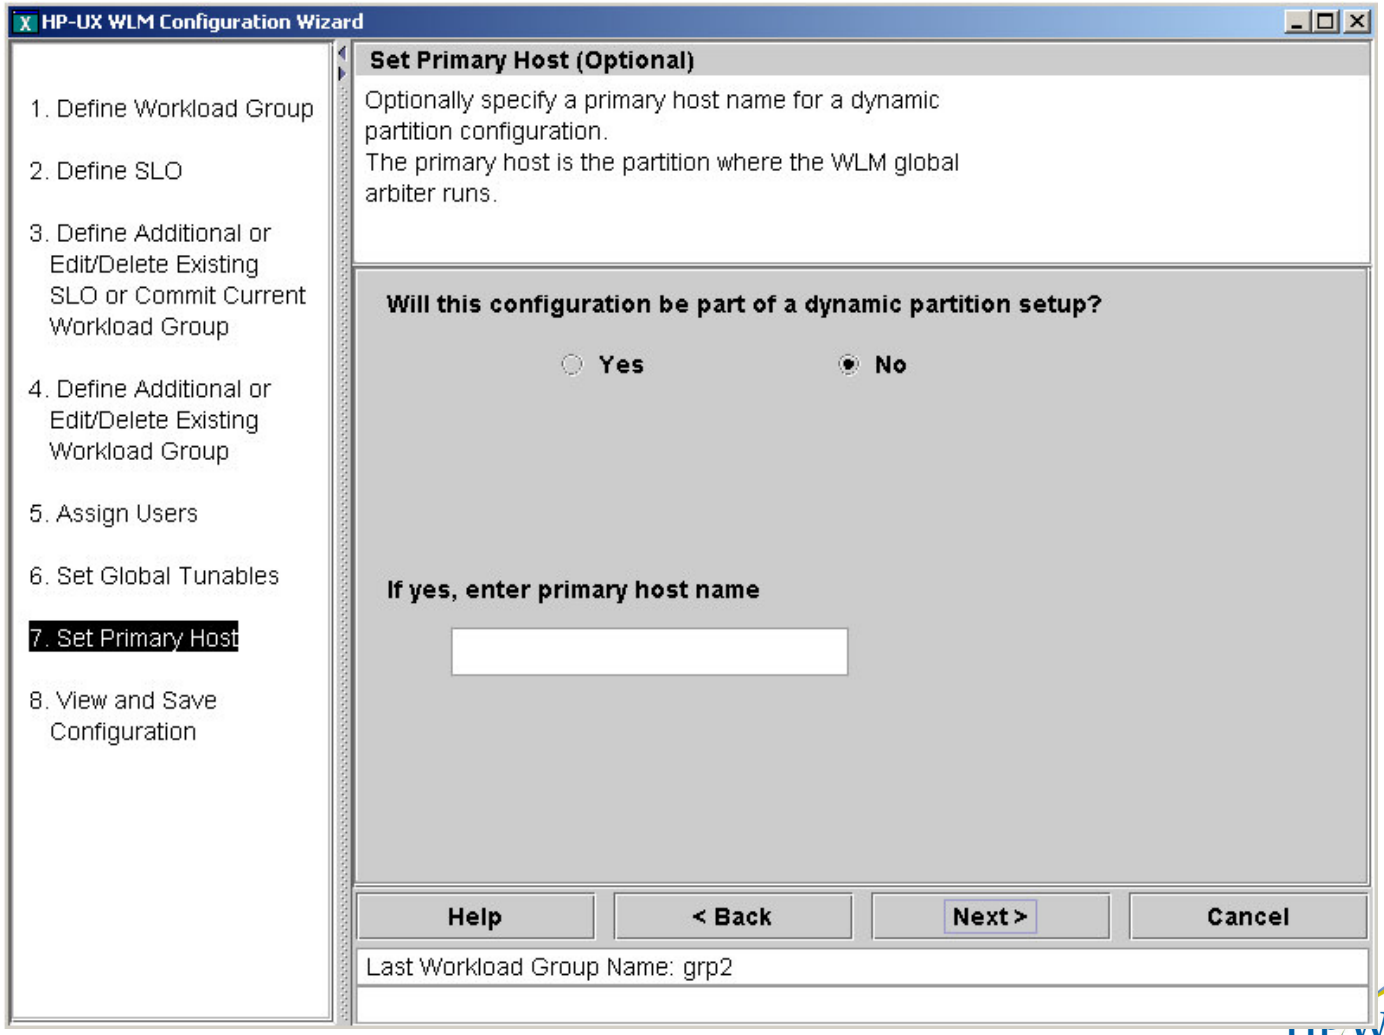

Viewing/Saving the configuration

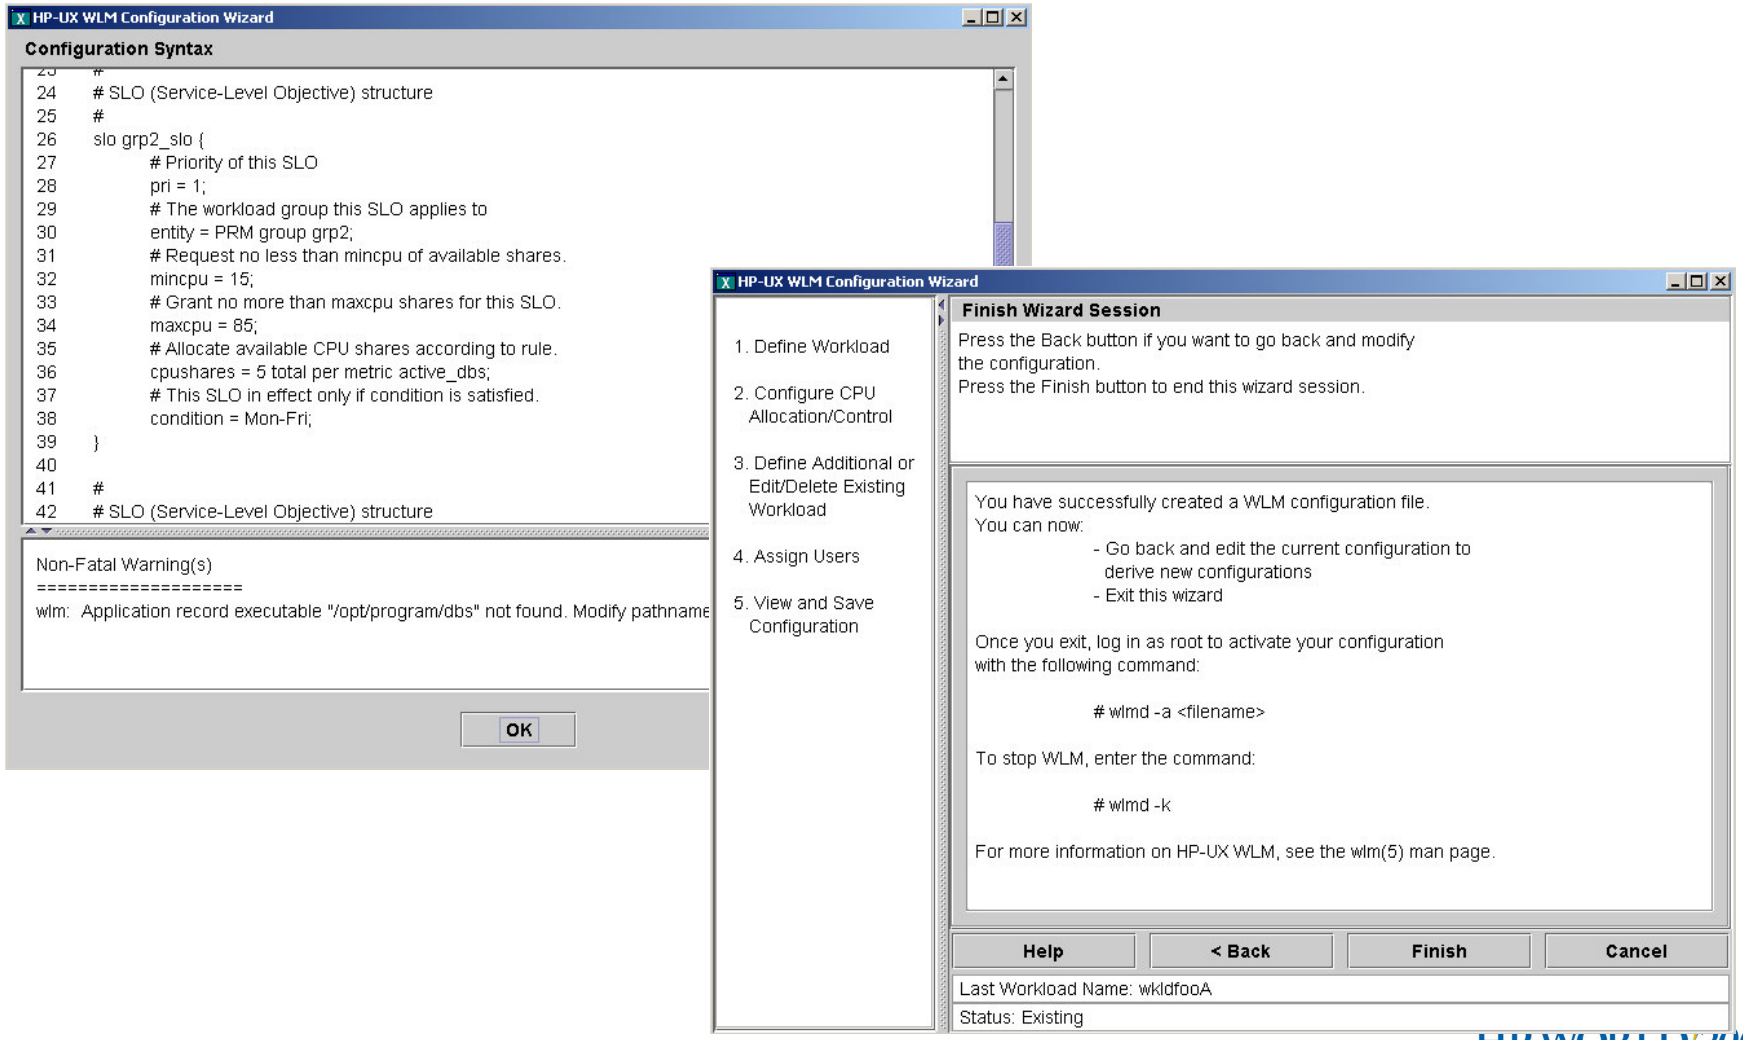

Help!

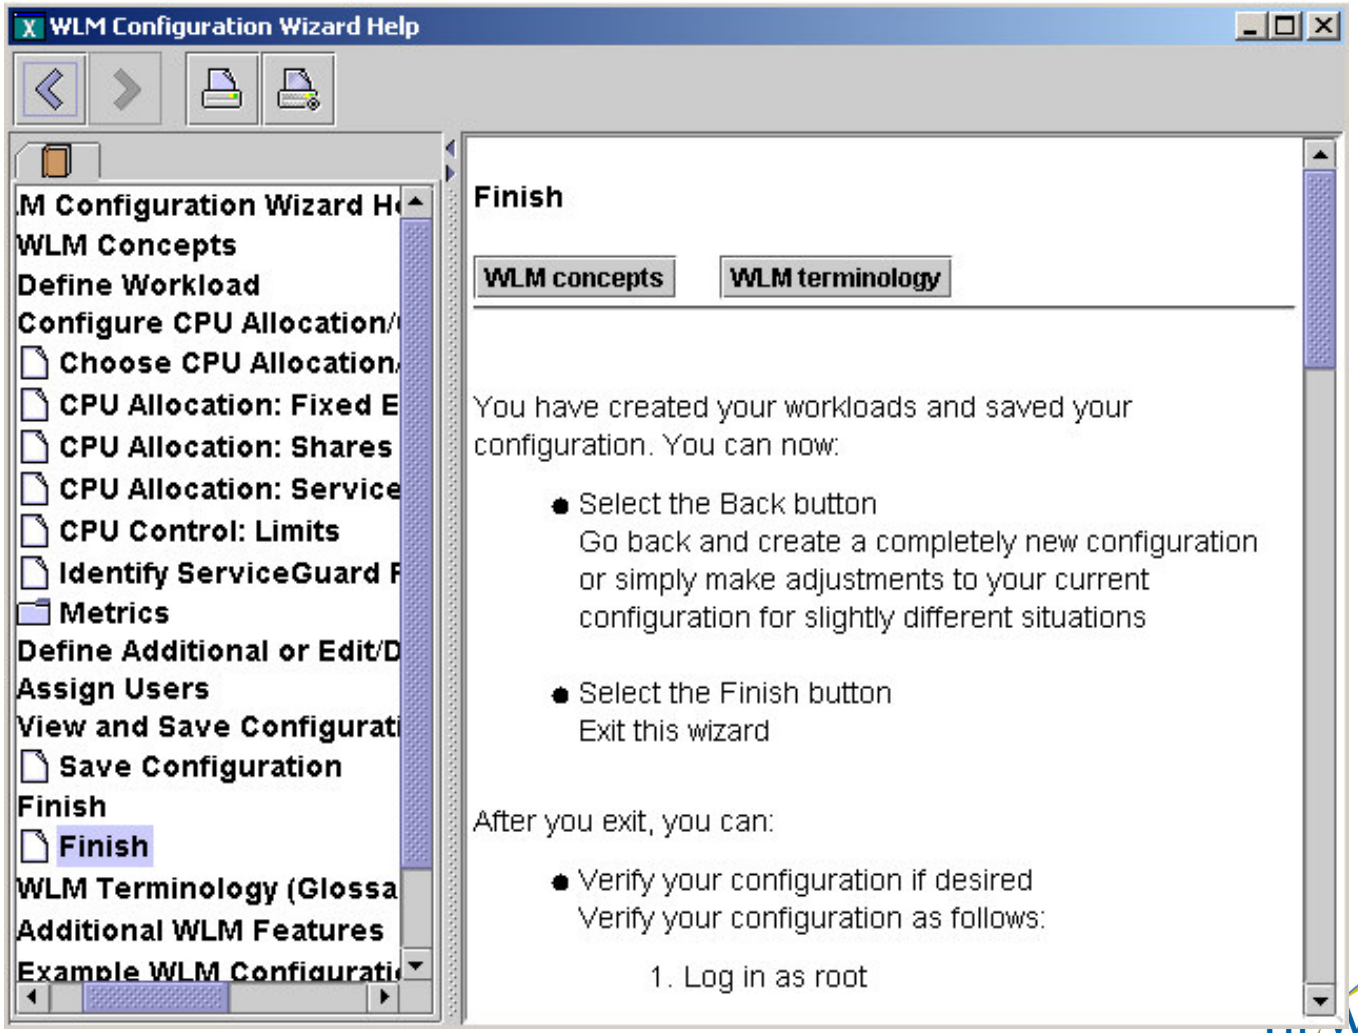
WLM Configuration Wizard Help
M Configuration Wizard H
WLM Concepts
Define Workload
Configure CPU Allocation/
Choose CPU Allocation
CPU Allocation: Fixed E
CPU Allocation: Shares
CPU Allocation: Service
CPU Control: Limits
Identify ServiceGuard F
Metrics
Define Additional or Edit/D
Assign Users
View and Save Configurati
Save Configuration
Finish
Finish
WLM Terminology (Glossa
Additional WLM Features
Example WLM Configurati
Finish
WLM concepts
WLM terminology
You have created your workloads and saved your configuration. You can now:
Select the Back button
Go back and create a completely new configuration or simply make adjustments to your current configuration for slightly different situations
Select the Finish button
Exit this wizard
After you exit, you can:
Verify your configuration if desired
Verify your configuration as follows:
1. Log in as root

Completing the Wizard Session

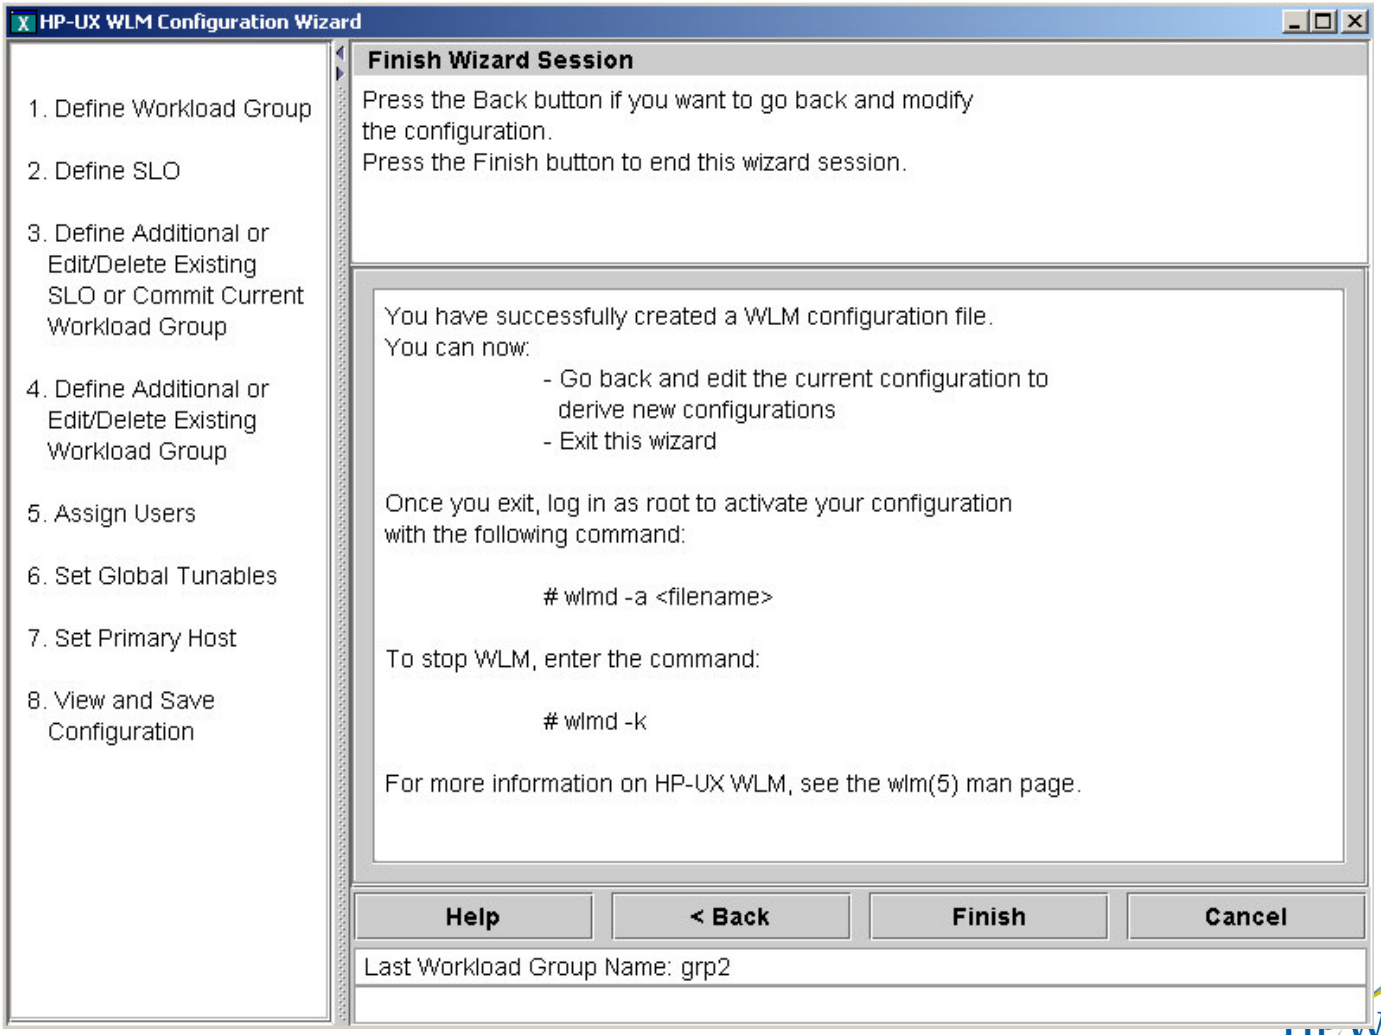

`wlmd` command

The **`wlmd`** command allows you to:

- Activate a configuration
- Validate a configuration file without altering the current configuration
- Specify **`wlmd`** logging options

wlmd syntax

wlmd has the following syntax and valid option combinations:

```
wlmd -h

wlmd -V

wlmd [-p] [-t] [-i] [-W] -A [-l
logoption[=n][,…]]

wlmd [-p] [-t] [-i] [-W] -a configfile [-l
logoption[=n][,…]]

wlmd [-W] -c configfile

wlmd -C

wlmd -k
```

Example:

WLM in passive mode

- WLM provides a passive mode that allows you to see how WLM will approximately respond to a given configuration without putting WLM in charge of your system's resources.
  - **`wlmd -p -a`** *configfile*
- Check that you configuration will behave as expected with minimal effect on your system.
- Possible uses for passive mode:
  - How does a **`condition`** statement work?
  - How does a **`cpushares`** statement work?
  - How do my goal work? Is my goal set up correctly?
  - How does a usage goal work?
  - Is my global configuration file set up as I wanted? If I used global arbitration on my production system, what might happen to the CPU layout?

Monitoring WLM with the **`wlminfo`** command

- **`wlminfo`** provides various WLM data, with reports focusing on SLOs, metrics, or workload groups.
- The command has both a command-line and graphical interface.
- The command syntax is below:

```
wlminfo -h [cmd]
wlminfo -V
wlminfo -i
wlminfo slo [-s slo] [-g grp] [-m met] [-l] [-o] [-v] [-b {0 | 1}]
wlminfo metric [-m met] [-l] [-o] [-b {0 | 1}]
wlminfo group [-g grp] [-l] [-o] [-b {0 | 1}]
wlminfo host [-l] [-o] [-b {0 | 1}]
wlminfo par [-h host] [-l] [-o] [-b {0 | 1}]
```

`wlminfo` command-line interface

```
# /opt/wlm/bin/wlminfo slo

Wed Apr 16 10:53:14 2003

SLO Name              Group         Pri     Req Shares   State Concern
slo2                  g2              1      30     30   FAIL  Max
slo3                  g3              1    20.5     21   PASS
```

```
# /opt/wlm/bin/wlminfo group

Wed Apr 16 10:56:14 2003

Workload Group    PRMID  CPU Shares   CPU Util   Mem Shares   State
OTHERS                1       50.00       0.00         0.00   ON
wkldfooA              2       30.00      15.55         0.00   ON
wkldfooB              3       20.00       5.00         0.00   ON
```

```
# /opt/wlm/bin/wlminfo metric

Wed Apr 16 10:56:14 2003

Metric Name                 PID State Value
count_time2               11337 NEW   4.000000
count_time3               11338 NEW   4.000000
```

wlmgui and **wlmcomd**

- The WLM graphical user interface, **wlmgui**, can be used on any system that has WLM installed or has the WLMUtilities bundle installed.
- In addition to monitoring, **wlmgui** can also be used to create, edit and deploy configurations.
- The **wlmcomd** daemon provides support for **wlmgui** allowing local and remote access to the system.
- As a security measure, **wlmcomd** must be explicitly started:

```
# /opt/wlm/bin/wlmcomd
```

- You can also start **wlmcomd** at boot time by editing the /etc/rc.config.d/wlm file.

Remote GUI Architecture

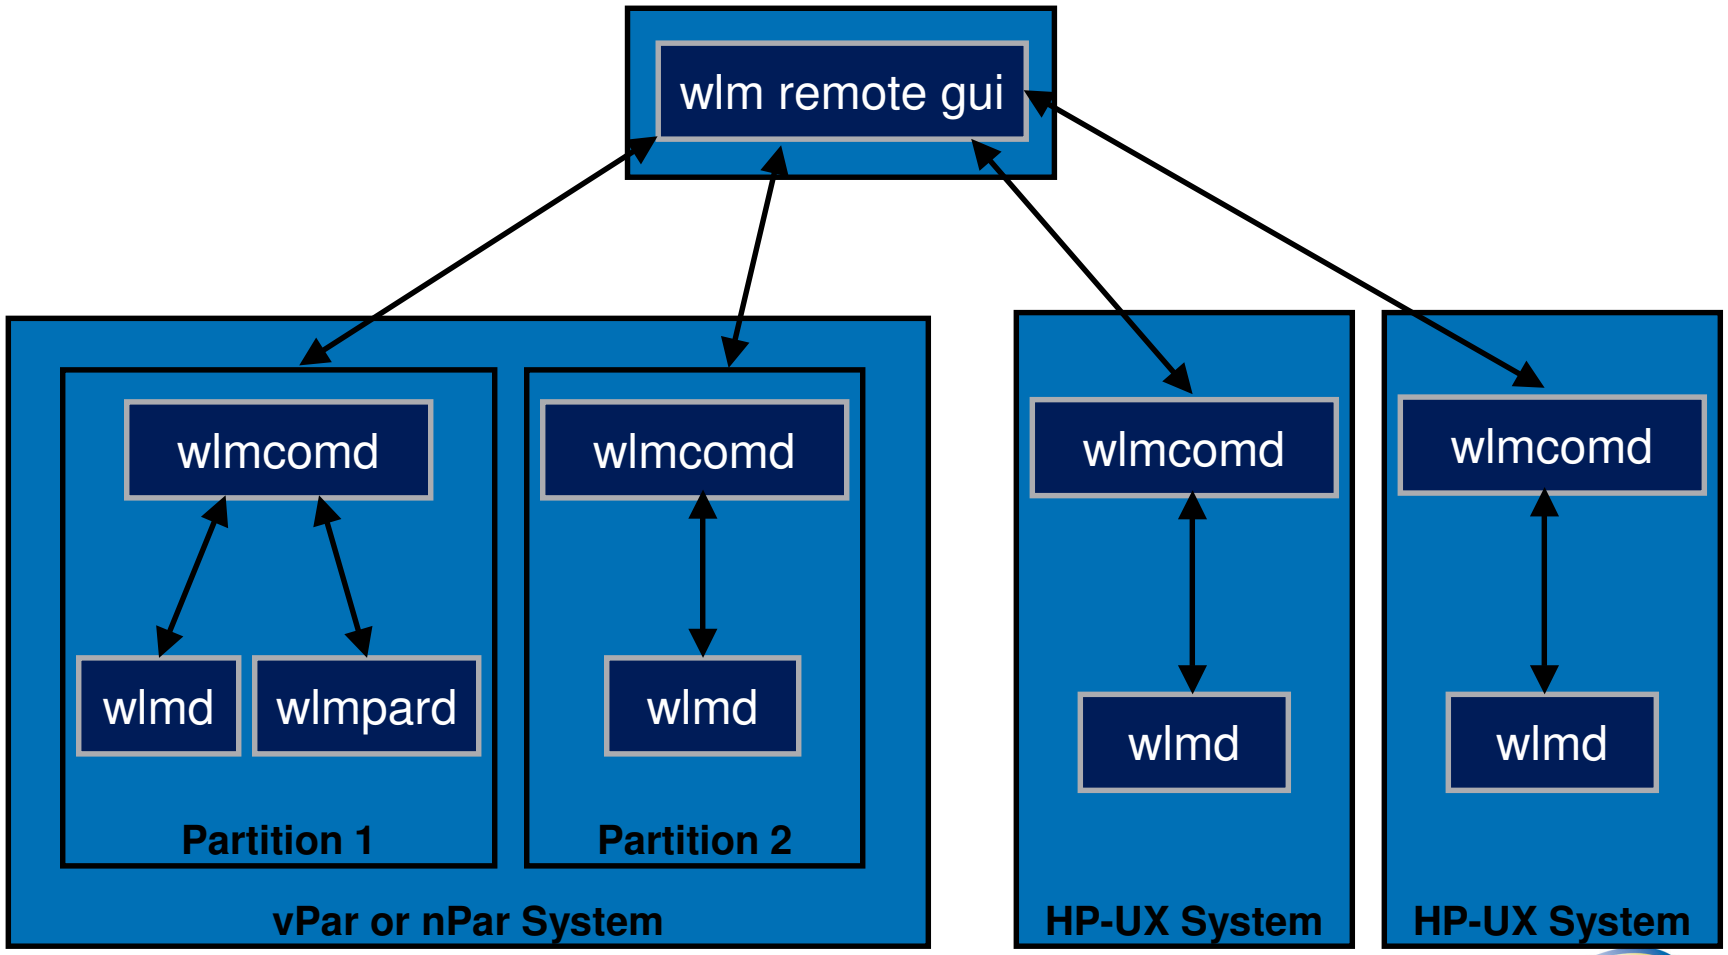

Monitoring WLM with **`wlmgui`**

- **`wlmgui`** has five views for monitoring:
  - Configuration, CPUs, Workload Groups, Service Levels, Custom
- To invoke the GUI, use the **`-i`** option to **`wlminfo`** or the **`wlmgui`** command:

```
# /opt/wlm/bin/wlminfo -i
# /opt/wlm/bin/wlmgui
```

- To perform remote management of a WLM-controlled system:
  - HP-UX workstation or server
    - Install the WLMUtilities bundle
    - Invoke **`wlmgui`**
  - Microsoft Windows
    - Install the utility using the setup.exe file
    - Invoke through the menu:

Starting `wlmgui` monitoring

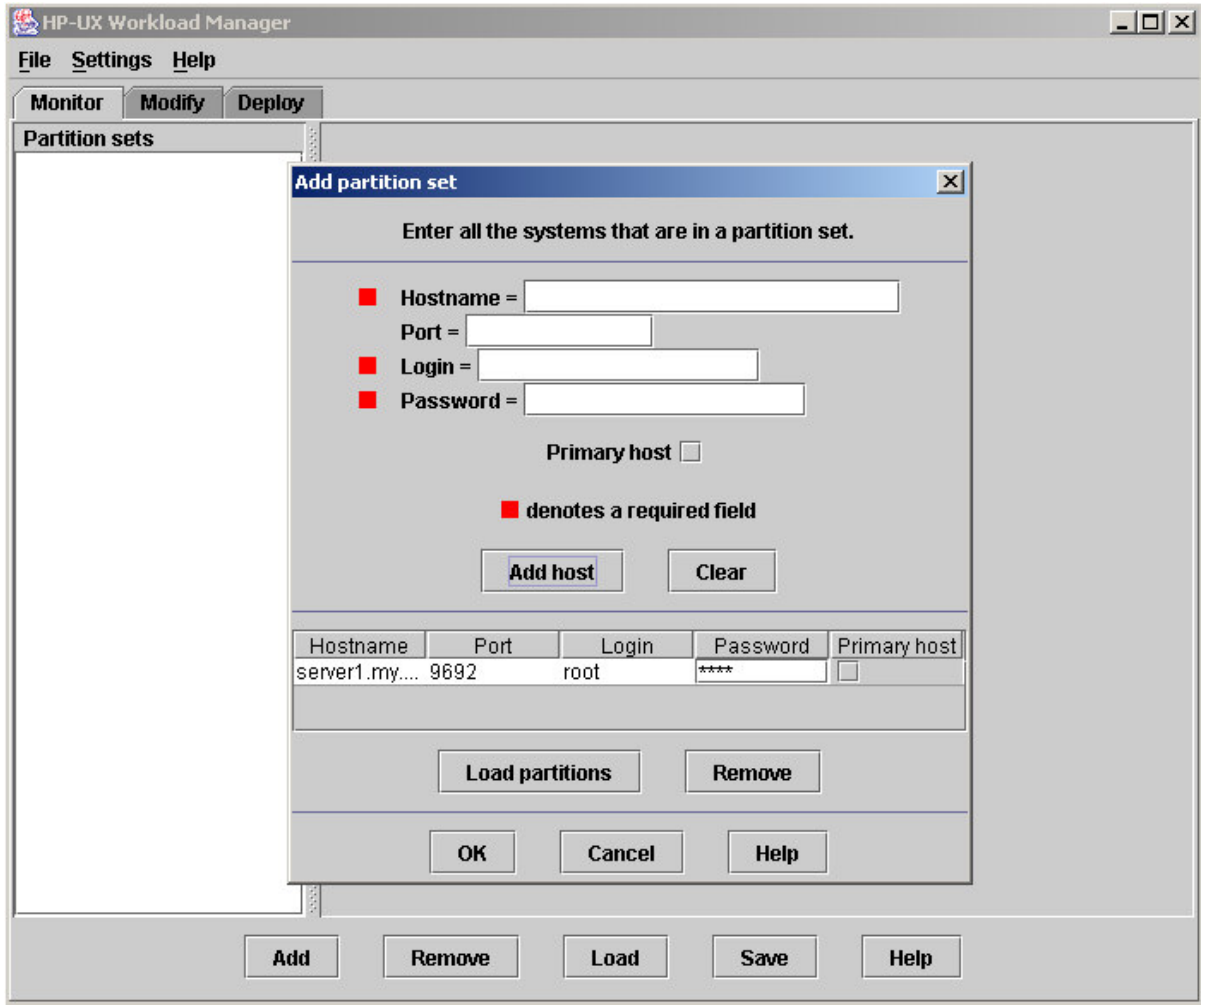

`wlmgui` monitoring views: Configuration

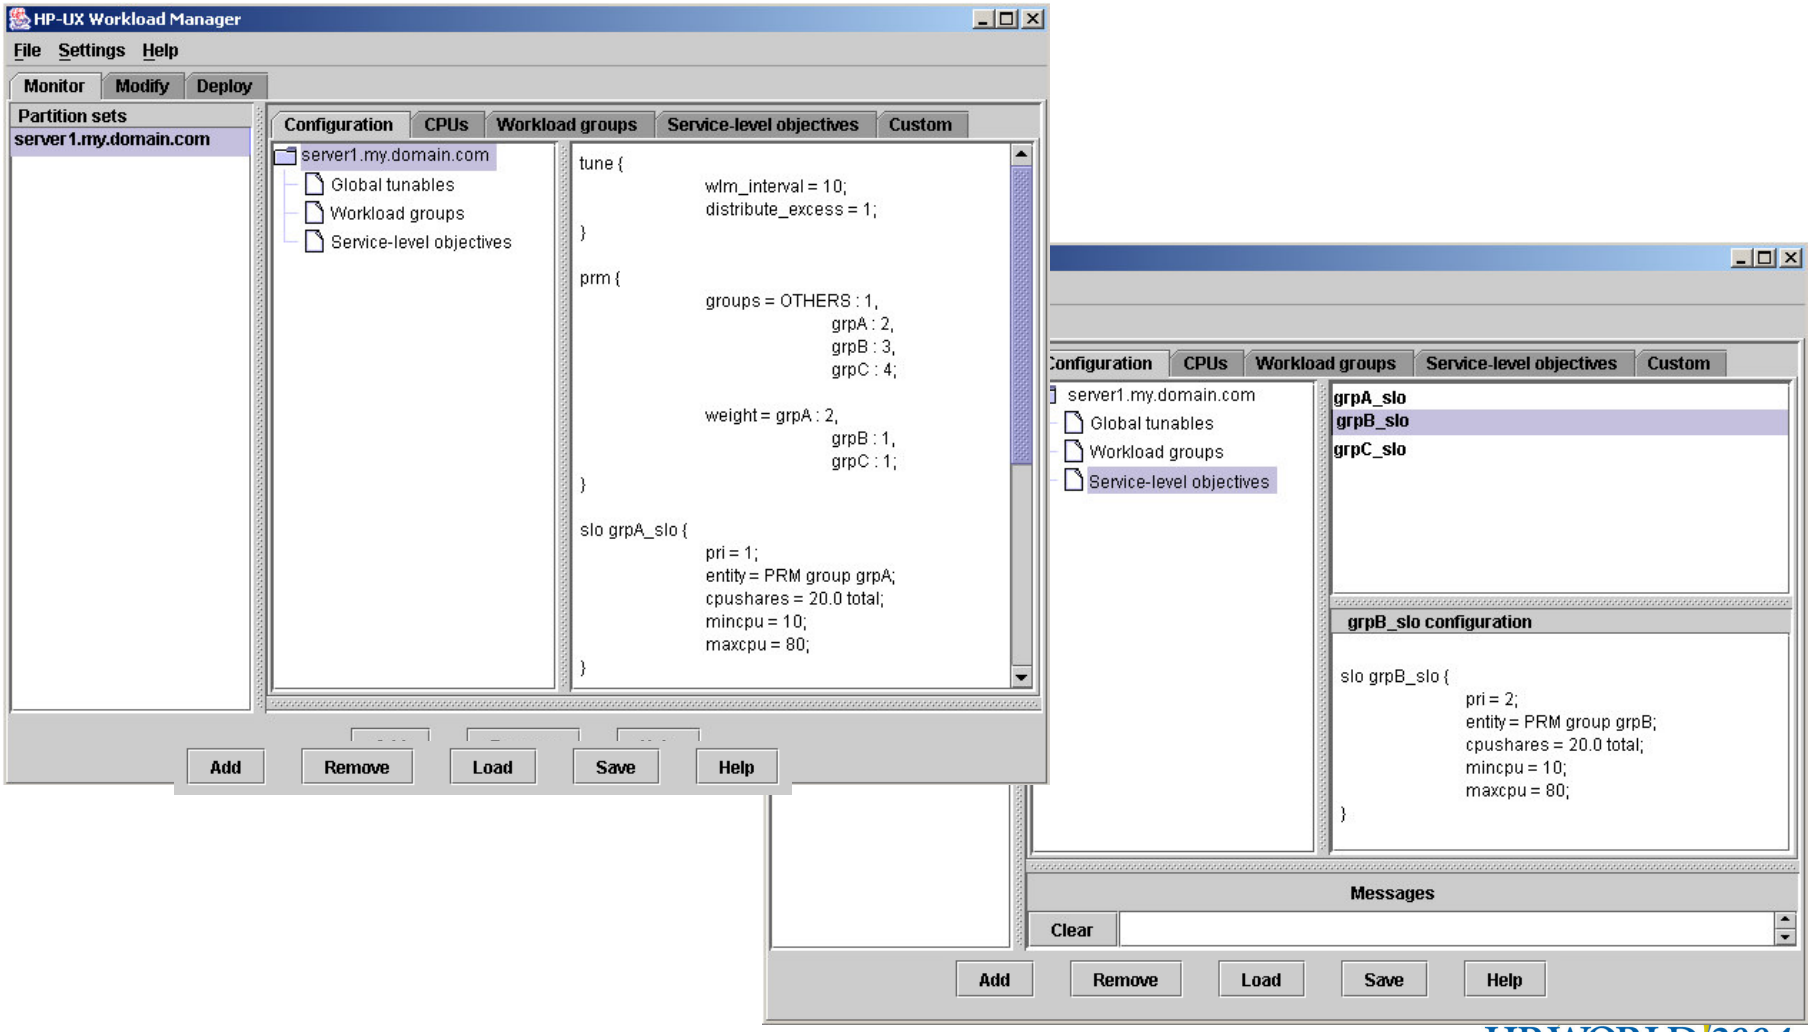

`wlmgui` monitoring views: Workload groups

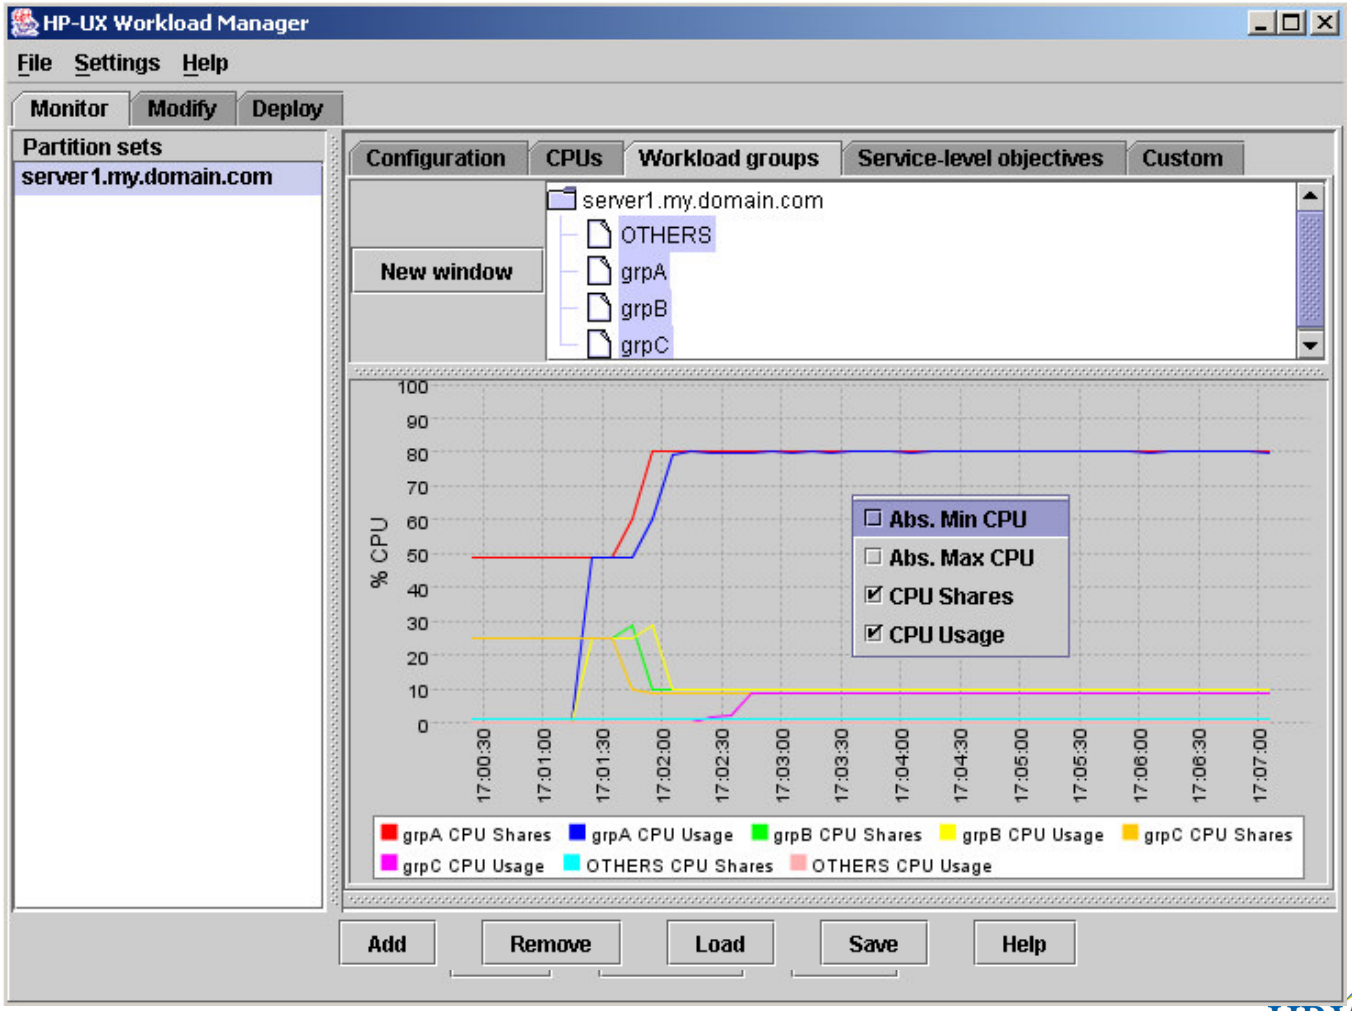

`wlmgui` monitoring views: Service-level objectives

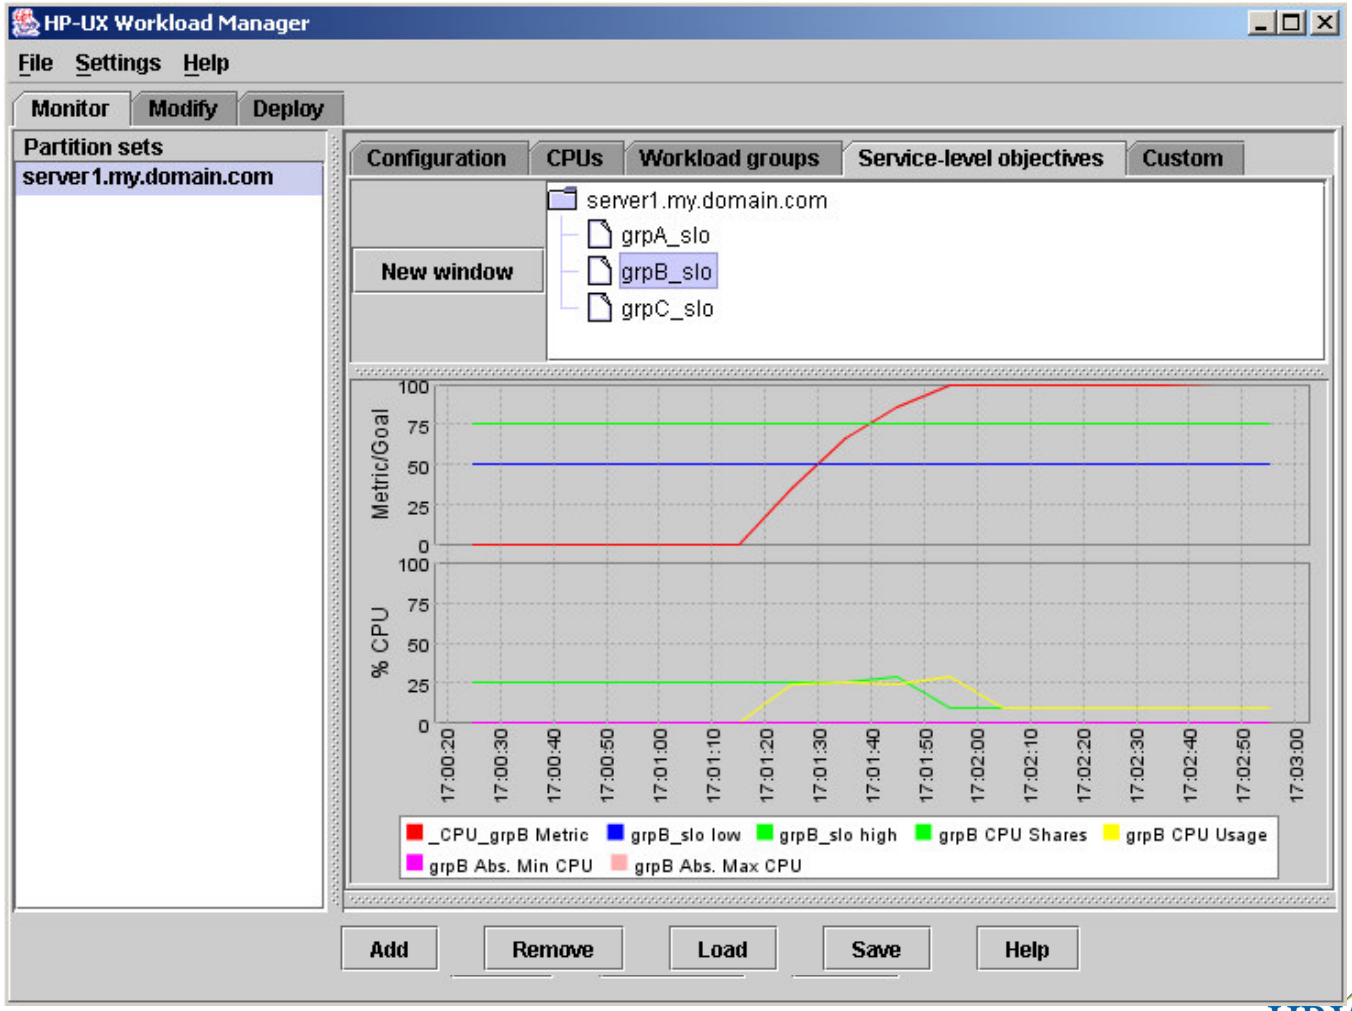

`wlmgui` monitoring views: Custom

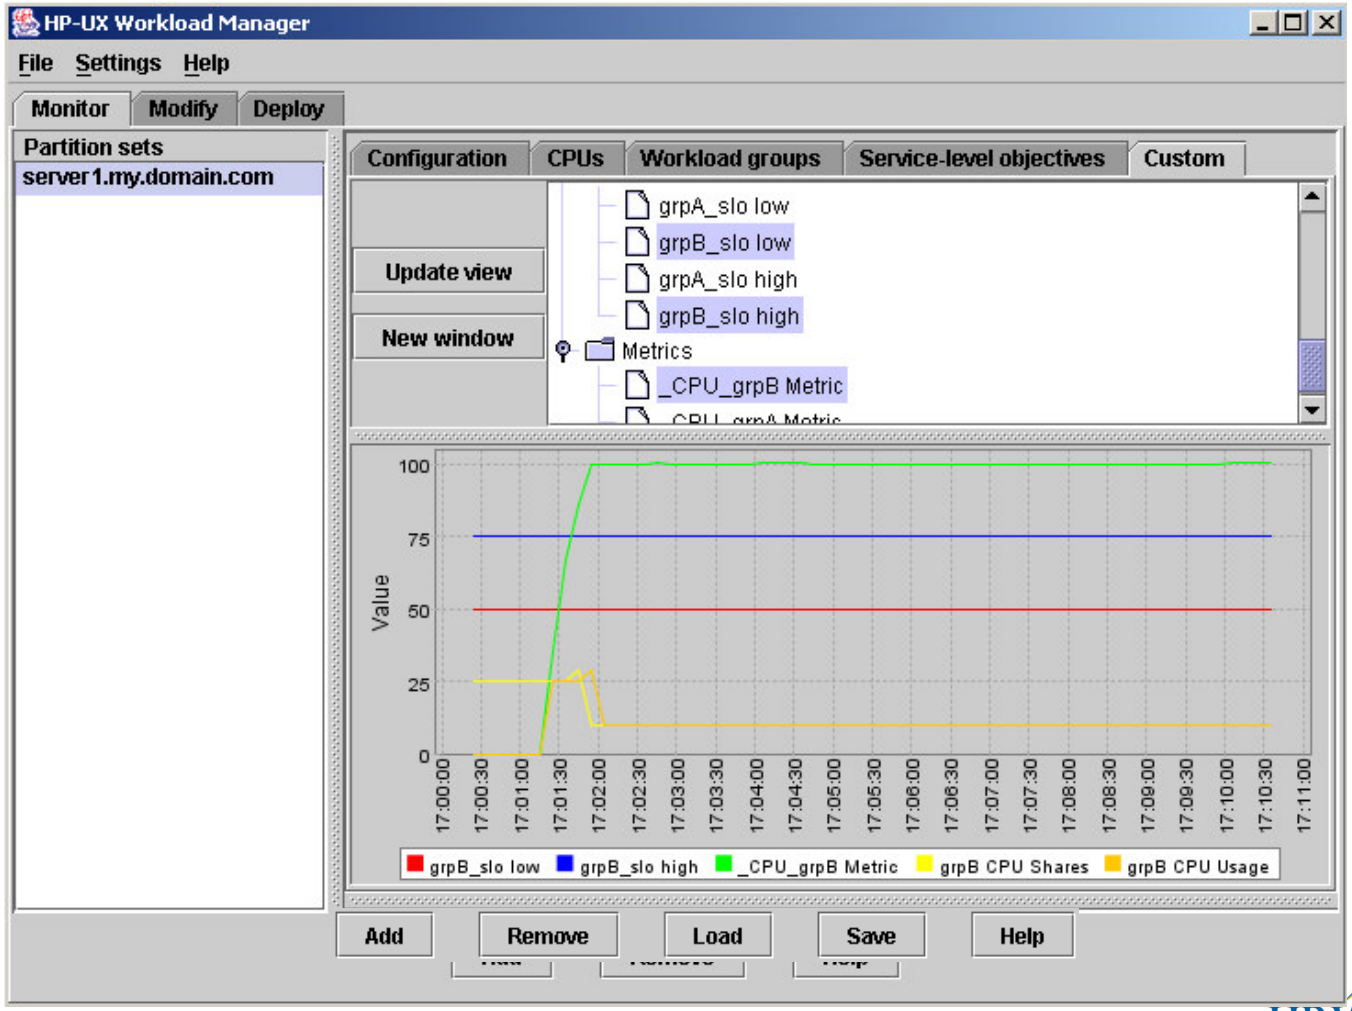

Changing graph view preferences

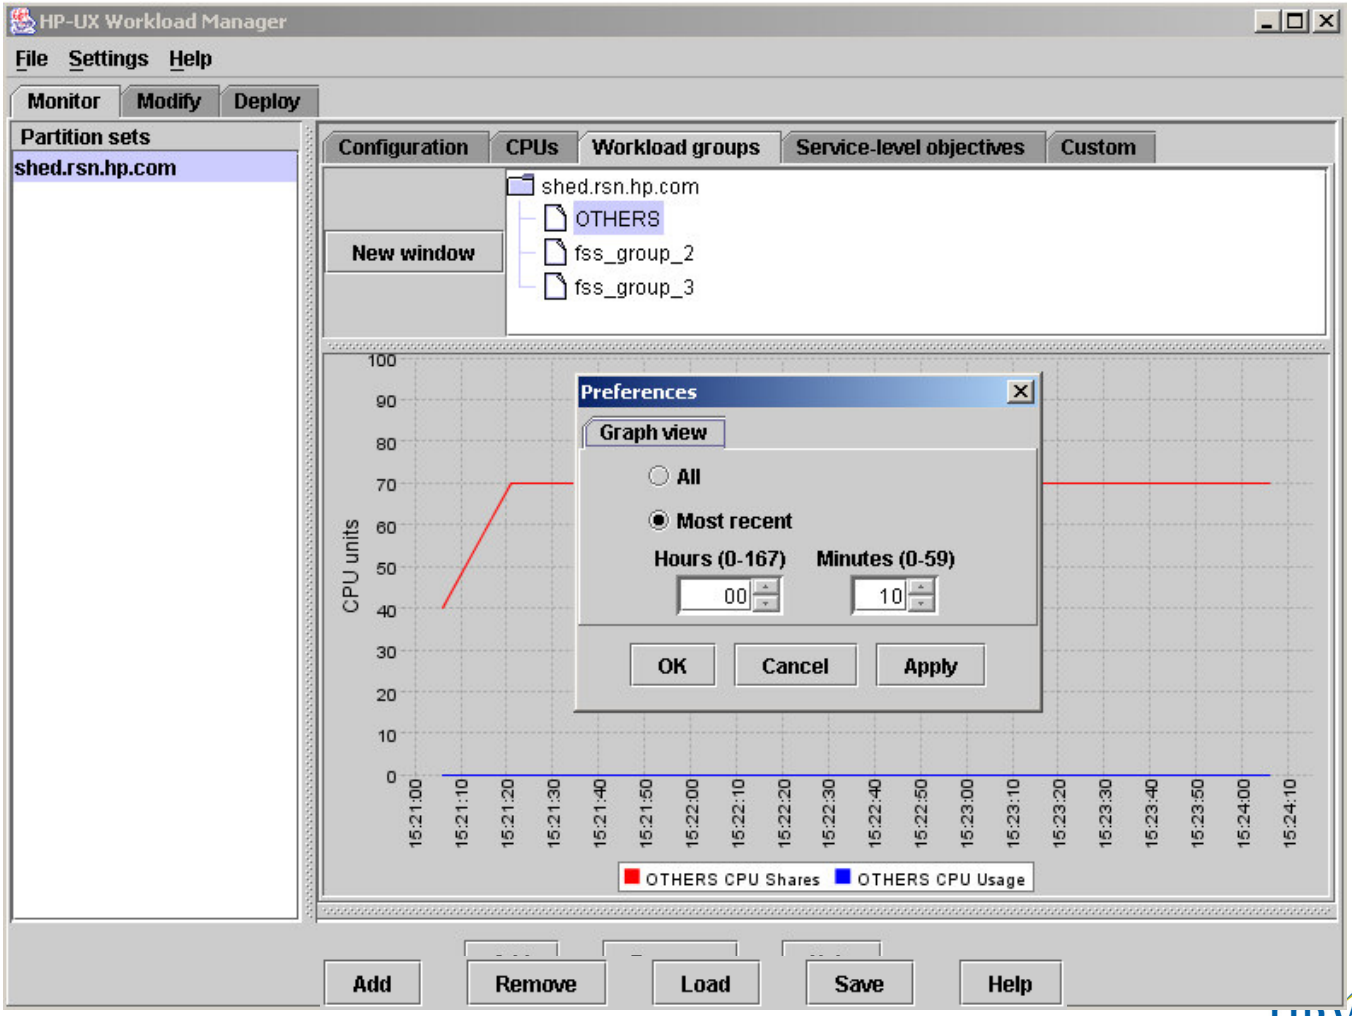

Saving partition sets

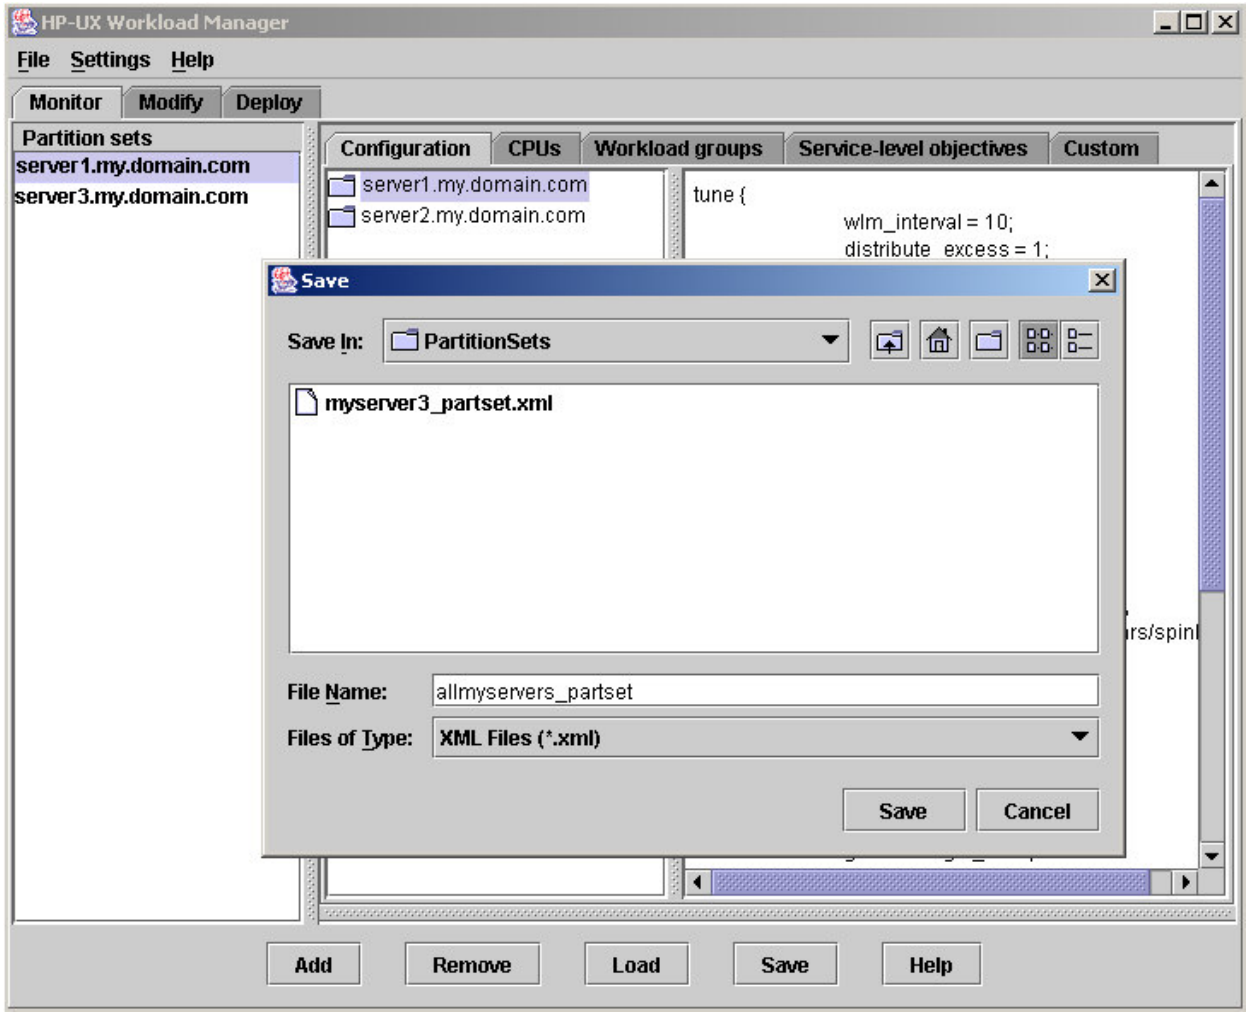

Loading a saved partition set definition

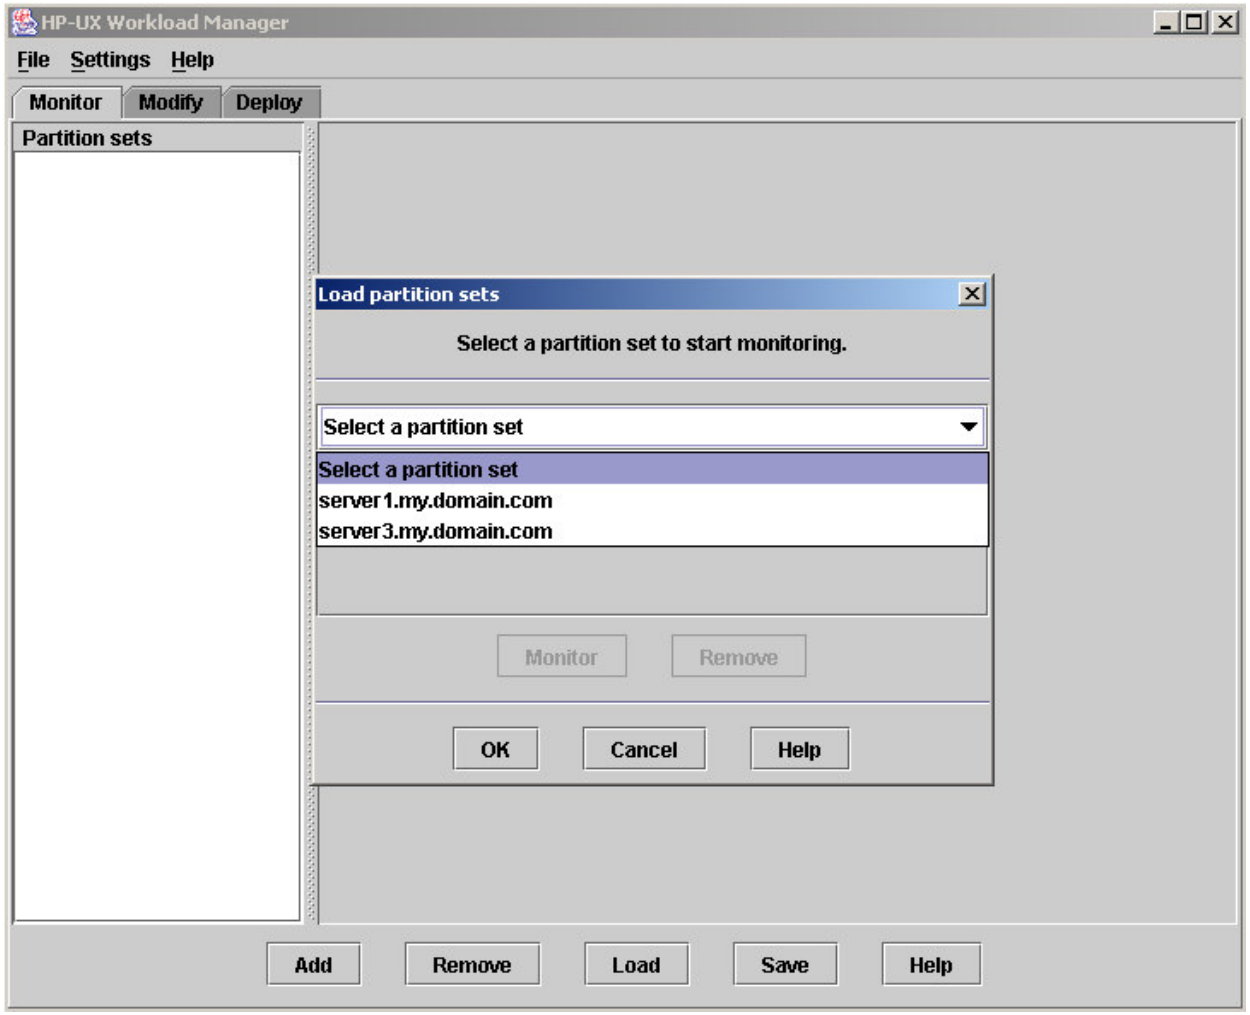

Reconfiguring WLM

To fine-tune an existing configuration, follow these steps:

1. Edit the WLM configuration file
2. (Optional) Activate the configuration in passive mode
3. Activate the configuration

It is not necessary to shut WLM down before activating a new configuration. Each of these steps can be done via the command line or **`wlmgui`**.

Setting WLM to start automatically at reboot

- For WLM to start automatically at reboot, set the **WLM_ENABLE** variable in the file /etc/rc.config.d/wlm to **1**:

  WLM_ENABLE=1

- When started at reboot, WLM by default uses the most recently activated configuration file
- To active a specific configuration at reboot, edit the following line in the /etc/rc.config.d/wlm file:

  WLM_STARTUP_SLOFILE="*configfile*"

  where ***configfile*** is the configuration file to use at reboot

Example configuration

- Look at the example configuration in your notes. It consists of three workload groups and three SLOs.

Module 2 Lab Exercise


hp
invent

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

HP Certified Professional