



# zEnterprise Platform Performance Management: Overview

Hiren Shah  
[hiren@us.ibm.com](mailto:hiren@us.ibm.com)



8/4/2010

© 2010 IBM Corporation

# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	HiperSockets	POWER7	System z10	zSeries*
BladeCenter*	IBM*	PowerVM	WebSphere*	z/VM*
DataPower*	IBM eServer	RP/SM	z9*	z/VSE
DB2*	IBM (logo)*	RACF*	z10 BC	
FICON*	InfiniBand*	System x*	z10 EC	
GDPS*	Parallel Sysplex*	System z*	zEnterprise	
Geographically Dispersed Parallel Sysplex	POWER*	System z9*	z/OS*	

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

InfiniBand is a trademark and service mark of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* All other products may be trademarks or registered trademarks of their respective companies.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

## Agenda

- **Platform Management Objectives**
- **Workload based monitoring**
- **Management functions**

# Platform Performance Management Objectives

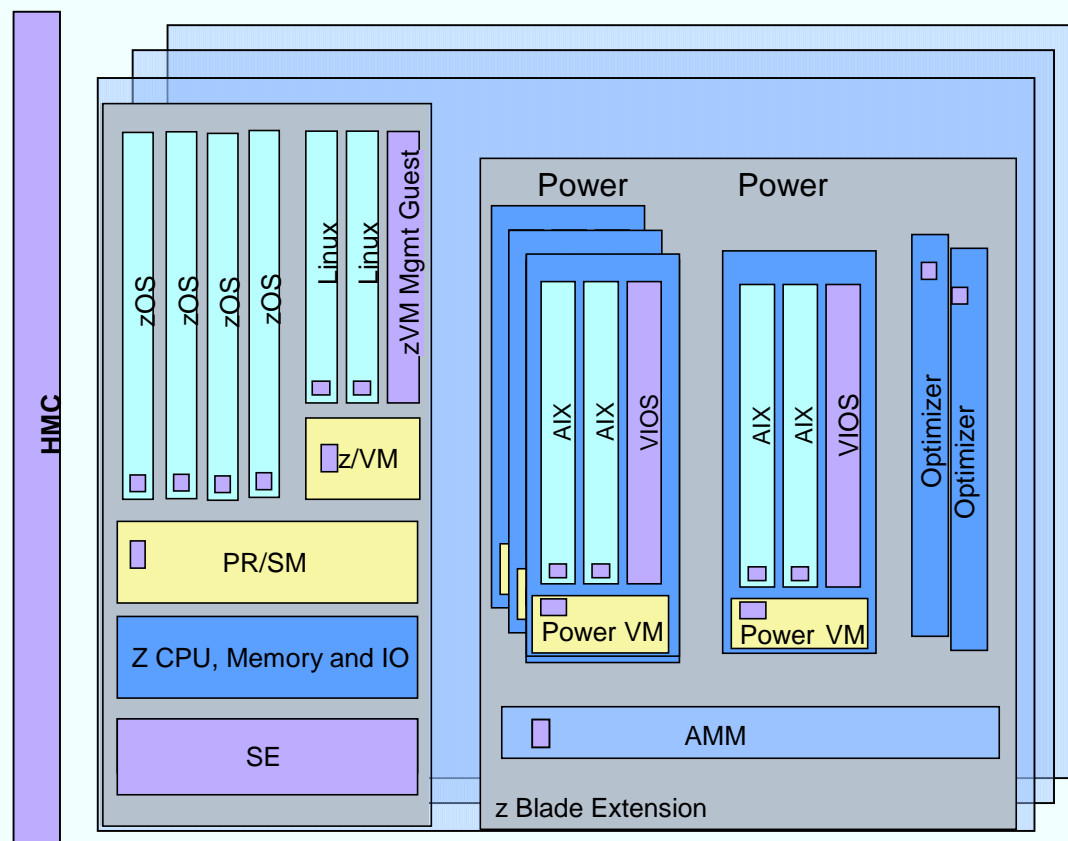
# zEnterprise Unified Resource Manager

## Ensemble:

- A zEnterprise Ensemble is a collection of zEnterprise Nodes managed as a single virtualized pool of server resources
  - Native LPAR and z/VM Virtual Images
  - Power VM Virtual images
  - IBM Smart Analytics Optimizer for DB2
- A zEnterprise Node can be a member of at most one Ensemble

## zEnterprise Unified Resource Manager

- HMC is management console
- Ensemble-Wide scope of responsibility
- Hardware configuration and operational control
- Virtual server life cycle management
- Virtual network and storage provisioning
- Energy Management
- Goal-oriented performance management

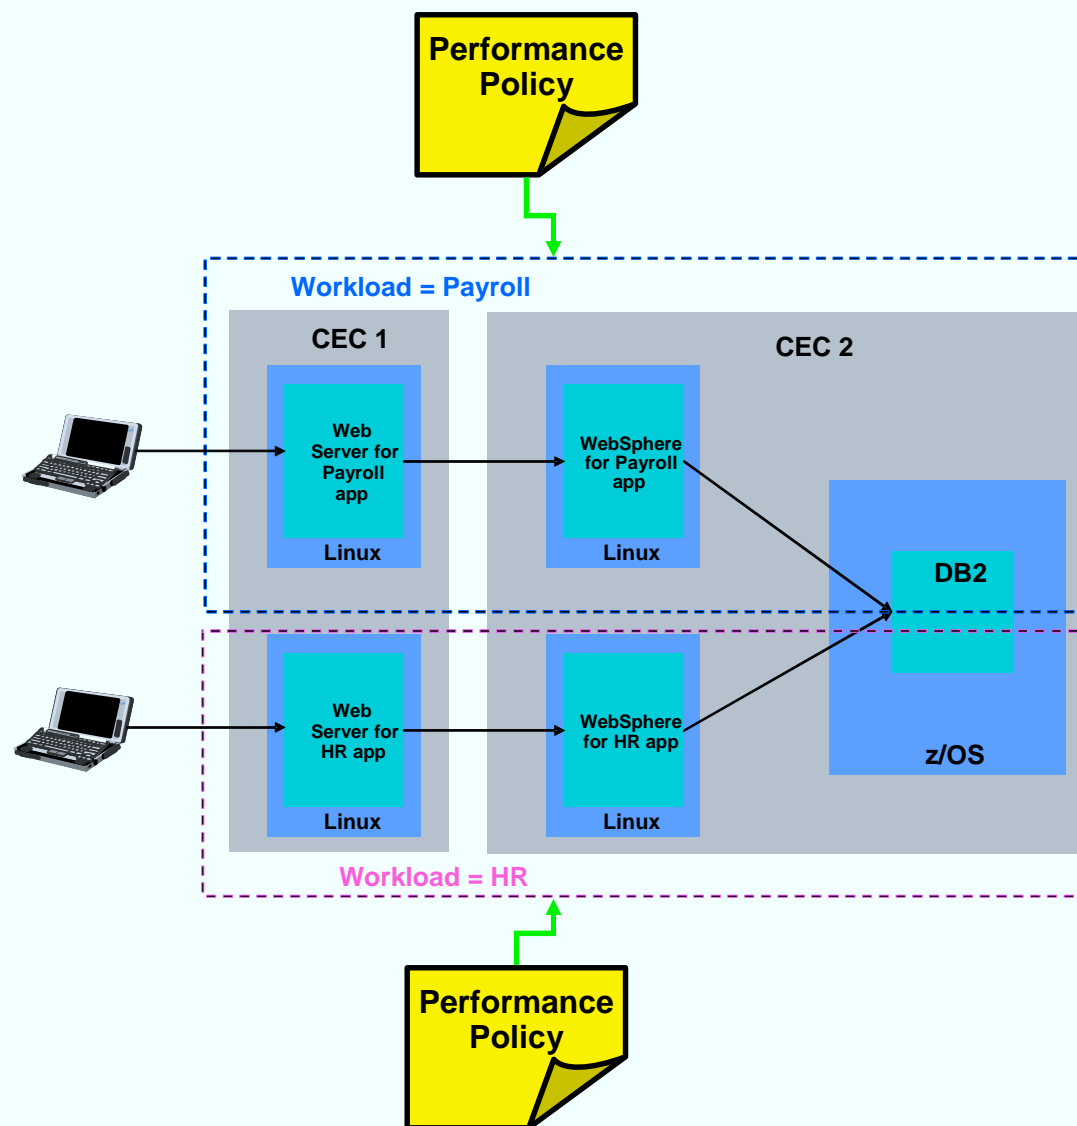


# zEnterprise Platform Performance Manager

- **Platform management component responsible for goal-oriented resource monitoring, management, and reporting across the zEnterprise Ensemble**
  - Core component responsible for definition and implementation of goal-oriented management policy
  - Workload monitoring and reporting based on management policy
  - Common approach to monitoring / management of platform resources across zEnterprise
  - Extend goal oriented approach of z/OS WLM to platform managed resources
  - Orchestration of autonomic management of resources across virtual servers
    - Provide Intelligent Resource Director like function across the zEnterprise
    - Management function will evolve over time
  - Pushes management directives to the SE, Hypervisors, and OS agents as required across the zEnterprise
- **Integration of HMC console support**
  - Integrated UI for monitoring, display of workload topology relationships, status alerts, etc
  - Definition of Performance Management Goals and Policy Administration
- **Functionality integrated into the Unified Resource Manager**
  - Code structured and packaged as System Z firmware
  - Inter-Component communication over trusted internal platform management network

## Platform Workload

- A Platform Workload is a grouping mechanism and “management view” of virtual servers supporting a business application
- Provides the context within which associated platform resources are presented, monitored, reported, and managed
- Management policies are associated to Platform Workload
  - Currently supports Performance Policy

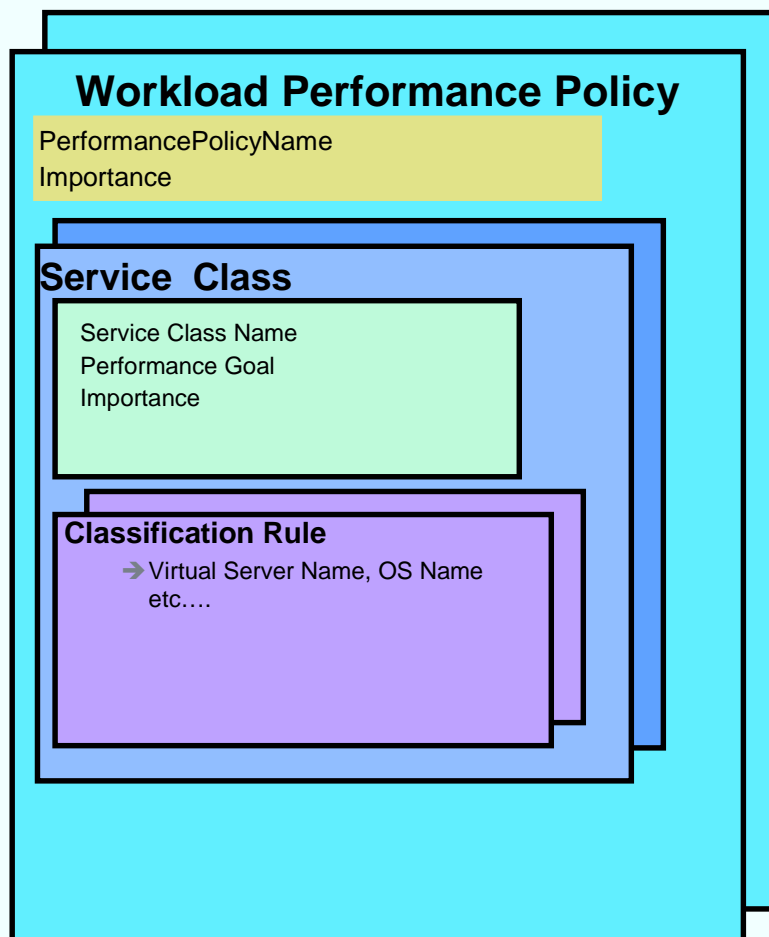


## Workload Performance Policy

- **Defines performance goals for virtual servers in a workload**
  - Conceptually similar to simplified z/OS WLM Policy
- **Provides basis for monitoring and management of platform resources used by virtual servers in a Workload**
- **Workload to performance policy relationship:**
  - A Workload can have multiple performance policies associated with it
  - Single policy is active at a given time
  - Can dynamically change the policy that is active
    - Through the UI
    - Through a timed based schedule
      - Example: Day shift policy / night shift policy



## Workload Performance Policy...



- **Policy structure:**
  - Policy contains a set of service classes
  - Classification rules map each virtual server within the workload to a service class
  - A service class assigns a performance goal and importance
- **HMC as console for policy creation and editing**
  - Wizard for policy creation
  - Repository for policies under development and saved policies
  - Links to Workload based performance reporting

# Workload Monitoring

## Workload Based Monitoring and Reporting

- **Provide reporting capability that shows usage of platform resources in a Workload context within a zEnterprise Ensemble scope**
  - Across virtual servers / partitions supporting the Workload
- **Workload goal vs actual reporting**
- **Drill down from overall Workload “performance health” view to contributions of individual virtual server**
- **Graphical views**
  - Topology, trending graphs, etc
- **Links to system activity displays to show hardware utilization views**
- **Reporting is limited to platform level resources, not trying to replicate tools that report on intra-OS resources and performance**

## Workload Monitoring Overview

- **Provide monitoring on the HMC based on a Workload context**
- **Display of current data and fairly recent history**
  - Current stake in the ground is 36 hours of history
  - Interval of data displayed is user selectable
  - Granularity of data kept in repository changes over time
    - 1 minute granularity kept for most recent hour
    - 15 minute interval data kept after first hour
- **Workload Report**
  - Display high level view of “performance health” of each Workload
  - Indication if a Workload contains service class missing goals
  - Worst performing service class / performance index
  - Details of specific Workloads
    - Graph of PI of worst performing service class
      - Option to graph other service classes
    - Bar graph of virtual server utilization distribution
      - Visualize view of workload overall load
  - Drill down to Workload’s service class report

**HMC1: Workloads Report - Mozilla Firefox**

http://9.60.31.154:8080/hmc/wcl/T107e

Workloads Report - AgParts Ensemble

Report Interval: Last 5 minutes [Modify](#)

--- Select Action --- Filter

Select	Workload	Service Class With Largest PI (PI)	Performance Policy
<input type="radio"/>	Default	Default (1.51)	Default
<input type="radio"/>	Payroll	Default (1.81)	Default
<input checked="" type="radio"/>	Weinheimer Agriculture Parts	Buyers (3.35)	Peak Period

Total: 3 Filtered: 3 Selected: 1

Workload Charts

Charts: CPU Utilization | [Performance Index](#)

CPU Utilization Distribution for Workload Weinheimer Agriculture Parts

CPU Utilization Range	Number of Servers
60.0-70.0	1
70.0-80.0	2
80.0-90.0	1

Close Help

Waiting for 9.60.31.154...

### HMC1: Workloads Report - Mozilla Firefox

http://9.60.31.154:8080/hmc/wcl/T1069?timestamp=129c2b9e22f#W1058

Report Interval: Last 15 minutes [Modify](#)

Last: 15 minutes

Starting from: Date: 7/11/10 Time: 1:14:34 PM Duration: 15 minutes

OK Apply Cancel

--- Select Action --- Filter

Select	Workload	Service Class With Largest PI (PI)	Performance Policy
<input type="radio"/>	Default	Default (1.36)	Default
<input type="radio"/>	Payroll	Default (1.74)	Default
<input checked="" type="radio"/>	Weinheimer Agriculture Parts	Buyers (2.92)	Peak Period

Total: 3 Filtered: 3 Selected: 1

Workload Charts

Charts: [CPU Utilization](#) | Performance Index

Service Class Performance Index (PI) for Workload Weinheimer Agriculture Parts

Time	Buyers (PI)	Default (PI)	PI=1 (PI)
14:00:00	2.6	1.0	1.0
14:01:00	2.6	1.0	1.0
14:02:00	2.9	1.4	1.0
14:03:00	2.6	1.0	1.0
14:04:00	3.4	1.6	1.0
14:05:00	1.8	1.0	1.0
14:06:00	3.2	1.6	1.0
14:07:00	1.6	0.8	1.0
14:08:00	2.8	1.4	1.0
14:09:00	3.8	1.6	1.0
14:10:00	2.5	1.0	1.0
14:11:00	3.2	1.6	1.0
14:12:00	3.4	1.0	1.0
14:13:00	4.2	2.0	1.0
14:14:00	2.5	1.0	1.0

Close Help

Done

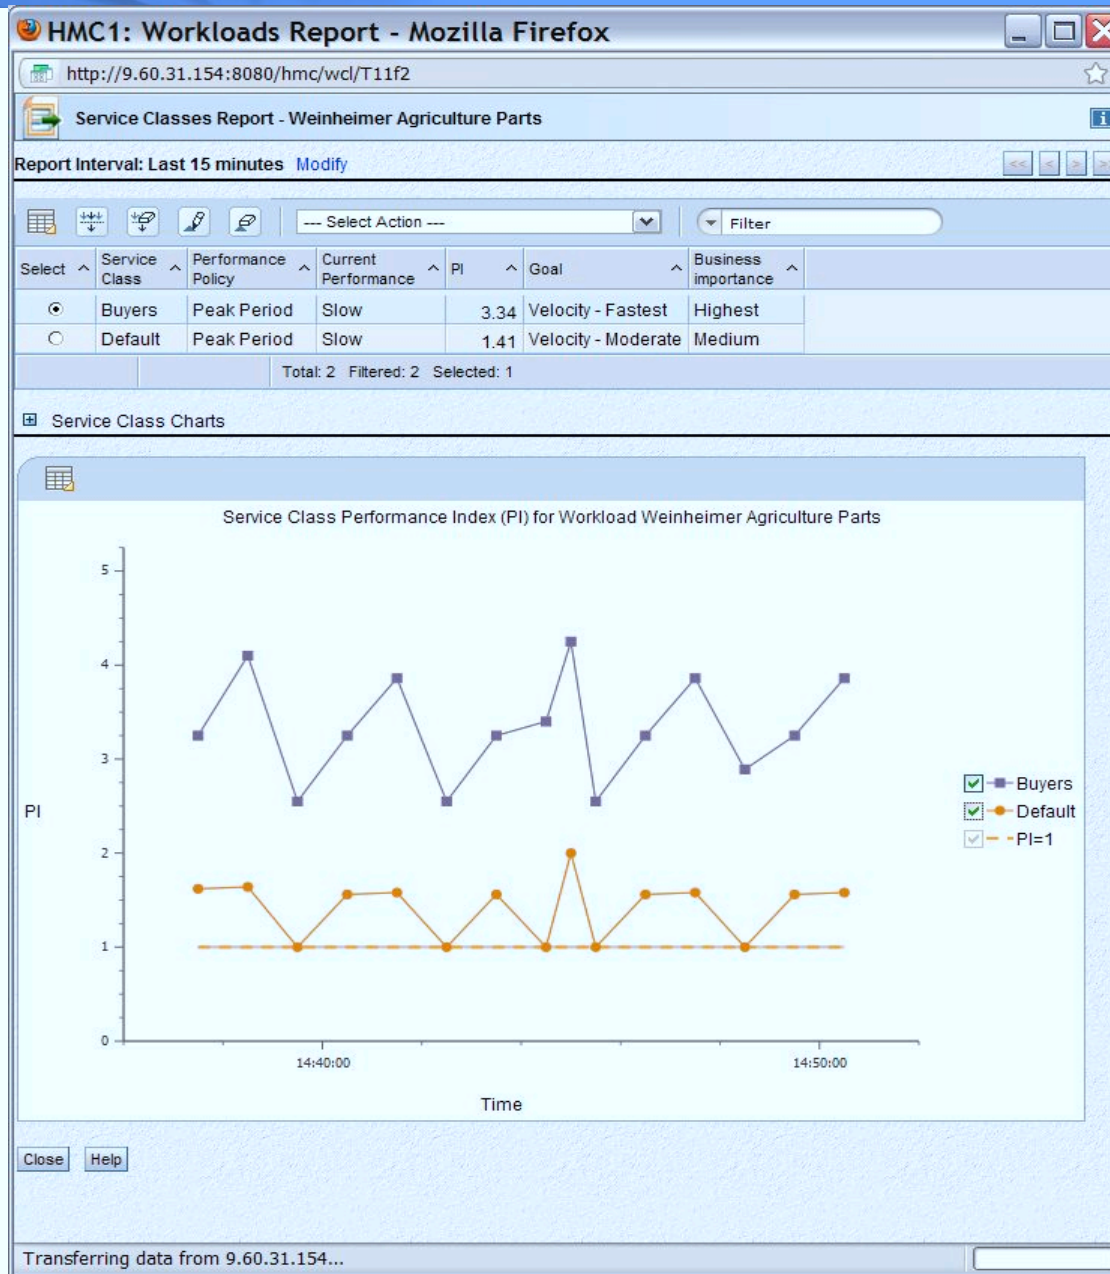
## Workload Monitoring Overview...

### ■ **Service Class Report**

- High-level view of each service class in Workload's performance policy
  - Goal and importance
  - Actual performance
  - Indication if monitoring event is established for service class and event is triggered
  - Service class details
    - Graph of service class performance index
  - Drill down to virtual server report for Workload

### ■ **Event Monitoring**

- Initial support:
  - Leverage HMC event monitoring
    - Send e-mail when selected metrics reach threshold
  - Service Class PI threshold
  - Virtual Server CPU Utilization threshold





**HMC1: Service Classes Report - Mozilla Firefox**

http://9.60.31.154:8080/hmc/wcl/T1615

Service Classes Report - Weinheimer Agriculture Parts

Report Interval: Last 15 minutes [Modify](#)

--- Select Action --- Filter

Select	Service Class	Performance Policy	Current Performance	PI	Goal	Business importance
<input checked="" type="radio"/>	Buyers	Peak Period	Fast	.65	Velocity - Fastest	Highest
<input type="radio"/>	Default	Peak Period	Fast	0.77	Velocity - Moderate	Medium

Total: 2 Filtered: 2 Selected: 1

**Service Class Charts**

Service Class Performance Index (PI) for Workload Weinheimer Agriculture Parts

PI

Time

Close Help

Done

## Workload Monitoring Overview...

- **Workload virtual server report**
  - List of virtual servers in a service class
    - Virtual server velocity
    - Resource usage
      - Physical CPU utilization
      - OS view of CPU utilization
      - Physical memory used
    - Hypervisor delay percentage
- **Resource adjustment report**
  - Resource adjustment actions taken over report interval
- **Transaction topology and hops report**
  - Topology of virtual servers
  - Transactions hops report

**HMC1: Workloads Report - Mozilla Firefox**

http://9.60.31.154:8080/hmc/wcl/T1325

Virtual Servers Report - Weinheimer Agriculture Parts

Report Interval: Last 15 minutes [Modify](#)

--- Select Action --- Filter

Select	Virtual Server	Virtual Processors	Allocated Memory (MB)	Physical CPU Utilization (%)	Hypervisor CPU Delay (%)	Idle Time (%)	Other Time (%)	Service Class With Largest PI (PI)	OS Processes Total CPU Using Samples (%)	OS Processes Total CPU Delay Samples (%)	OS Processes Total I/O Delay Samples (%)	OS Processes Total Page Delay Samples (%)
<input checked="" type="radio"/>	Buyer 1	1	1,024	88.7	94.0			Buyers (1.64)				
<input type="radio"/>	Buyer 2	1	1,024	83.0	66.0			Buyers (1.64)				
<input type="radio"/>	Vendor 1	1	1,024	83.0	66.0			Default (0.76)				
<input type="radio"/>	Vendor 2	1	1,024	66.6	66.1			Default (0.76)				

Total: 4 Filtered: 4 Selected: 1

**Virtual Server Charts**

CPU Utilization for Virtual Server Buyer 1

CPU Utilization (%)

Buyer 1

Done

**HMC1: Workloads Report - Mozilla Firefox**

http://9.60.31.154:8080/hmc/wcl/T14d0

**Virtual Server Resource Adjustments Report - Buyer 1**

Report Interval: Last 15 minutes [Modify](#)

Successful Adjustments

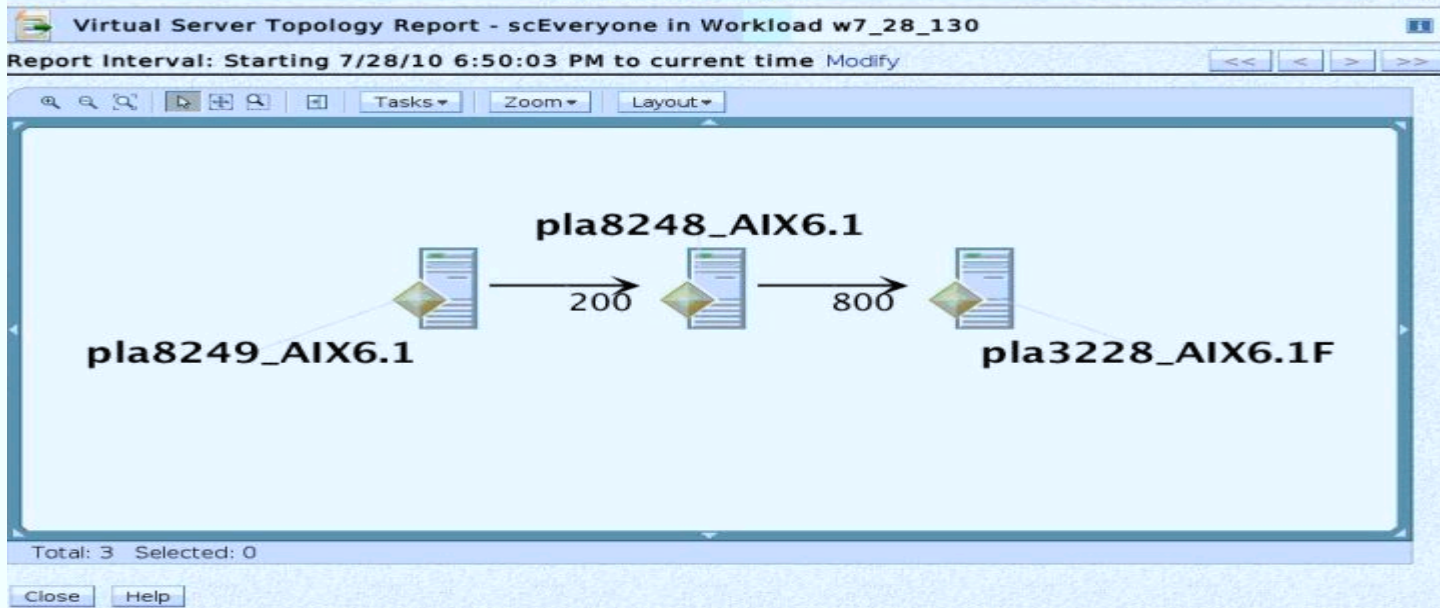
Receiver Virtual Servers ^	Receiver Workload ^	Receiver Service Class ^	Receiver Processing Units After (Before) ^	Donor Virtual Servers ^	Donor Workload ^	Donor Processing Units After (Before) ^	Time ^
Buyer 1	Weinheimer Agriculture Parts	Buyers	0.52 (0.50)	Payroll App	Payroll	0.49 (0.50)	Jul 11, 2010 4:13:18 PM
Buyer 1	Weinheimer Agriculture Parts	Buyers	0.52 (0.50)	Vendor 1	Weinheimer Agriculture Parts	0.49 (0.50)	Jul 11, 2010 4:13:18 PM
		Total: 2		Filtered: 2			

Failed Adjustments

Receiver Virtual Servers ^	Receiver Workload ^	Receiver Service Class ^	Failure Reason ^	Time ^
Total: 0 Filtered: 0				

Close Help

Done



Hops Report - scEveryone in Workload w7\_28\_130  
 Report Interval: Starting 7/28/10 6:50:03 PM to current time Modify


Details for scEveryone

Workload: w7\_28\_130 Performance policy: p1  
 Performance goal: Velocity - Moderate Business importance: Medium  
 PI: 0.40 Performance: Fastest

Name	Hop Number	Group Name	Successful Transactions	Failed Transactions	Stopped Transactions	Inflight Transactions	Queue Time (s)	Execution Time (s)	Successful Average Response Time (s)	
▣ Hop 0		0	200	0	0	0	2	0.000	0.000	0.014
▣ IBM DB2 Universal Database		0 db2inst1	0	0	0	0	0	0.000	0.000	0.000
▣ IBM Webserving Plugin		0 IBM_HTTP_Server	200	0	0	0	0	0.000	0.000	0.014
pla8249_AIX6.1		0	200	0	0	0	0	0.000	0.000	0.014
▣ WebSphere:APPLICATION_SERVER		0 server1	0	0	0	0	1	0.000	0.000	0.000
pla8248_AIX6.1		0	0	0	0	0	1	0.000	0.000	0.000
▣ HelloWorld		0 Examples	0	0	0	0	1	0.000	0.000	0.000
▣ Hop 1		1	200	0	0	0	0	0.000	0.000	0.006
▣ WebSphere:APPLICATION_SERVER		1 server1	200	0	0	0	0	0.000	0.000	0.006
pla8248_AIX6.1		1	200	0	0	0	0	0.000	0.000	0.006
▣ Hop 2		2	800	0	0	0	0	0.000	0.000	0.000
▣ IBM DB2 Universal Database		2 db2inst1	800	0	0	0	0	0.000	0.000	0.000

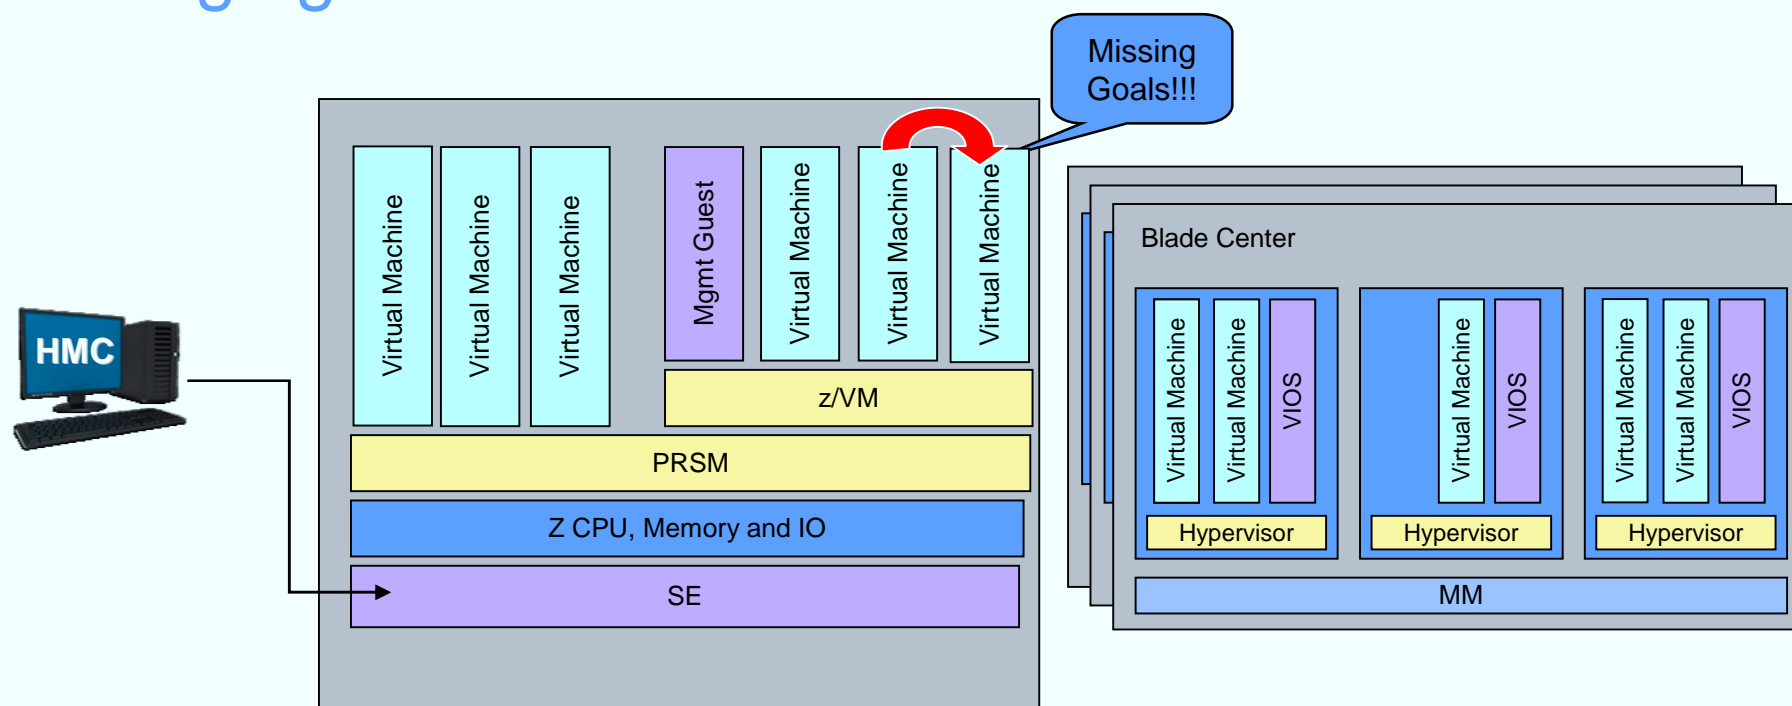
# Management Functions

# Managing Resources across Virtual Servers on P7 blade



- **Manage resources across virtual servers to achieve workload goals**
  - Detect that a virtual server is part of Workload not achieving goals
  - Determine that the virtual server performance can be improved with additional resources
  - Project impact on all effected Workloads of moving resources to virtual server
  - If good trade-off based on policy, redistribute resources
  - Initially support CPU management

# Managing Resources across z/VM Virtual Machines



- **Manage resources across z/VM virtual machines**

- Detect that a virtual machine that is part of Workload is not achieving goals
- Determine that the virtual machine performance can be improved with additional resources
- Project impact on all effected Workloads of moving resources to virtual machine
- If good trade-off based on policy, redistribute resources
- Initially support CPU management



## Platform Performance Manager Summary

- **Extend z/OS goal oriented workload management concepts across zEnterprise mixed processors environment**
- **Integrated function of zEnterprise Unified Resource Manager firmware**
- **Workload based goal oriented policy definition**
- **Monitoring and reporting in context of Workload and associated performance policy**
- **Goal oriented resource management**