

# Hybrid Management

Dr. Stefan Wirag

z/OS Workload Management Development  
IBM Deutschland Research & Development GmbH

October 14, 2010

[stefan.wirag@de.ibm.com](mailto:stefan.wirag@de.ibm.com)



---

## Agenda

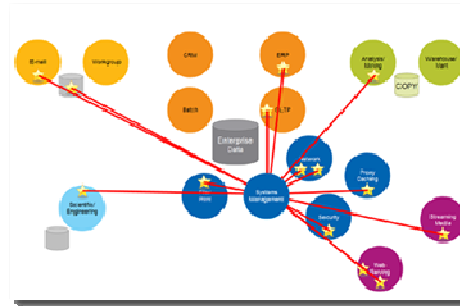
- System z Hybrid
- zEnterprise Unified Resource Manager
- z/OS Workload Management
- zEnterprise Platform Performance Management
  - Workload based Monitoring
  - Management Functions

# System z Hybrid

## System z Mainframe



## Integrated Systems Management firmware



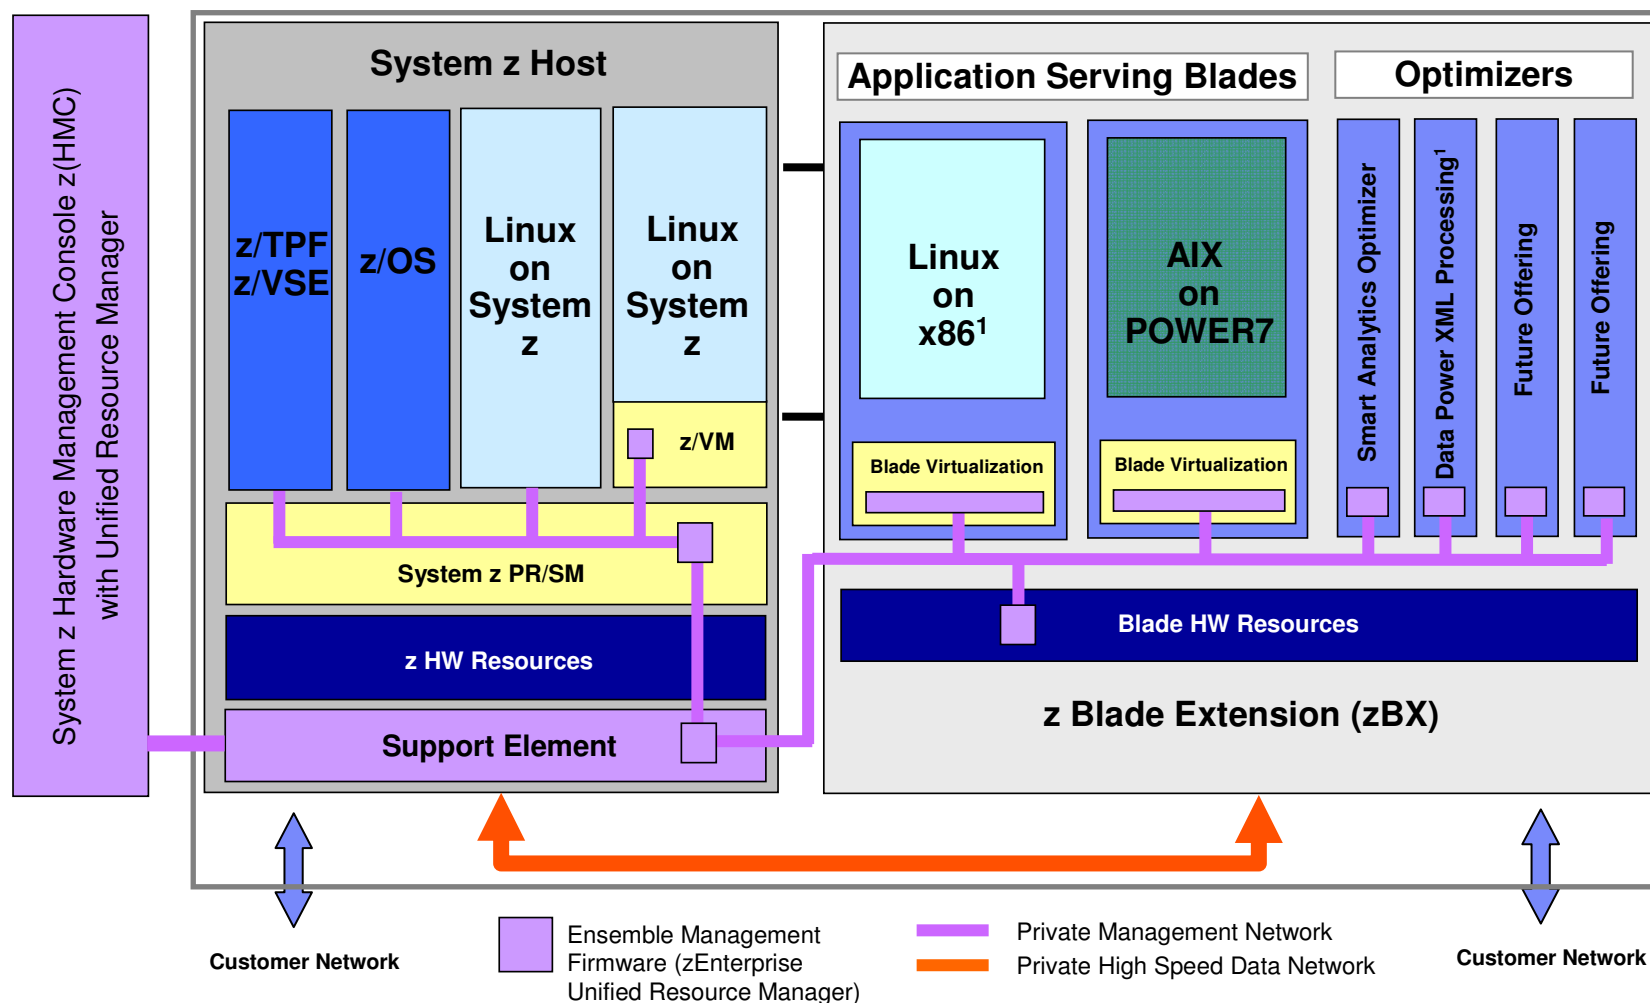
## Optimizers

- Extend and accelerate System z workloads
- Lower cost per transaction while improving application response time for CPU intensive applications

## Application Serving Blades

- Logical device integration between System z resources and application serving commodity devices
- Providing competitive price-performance and improved QoS for applications with a close affinity to mainframe data

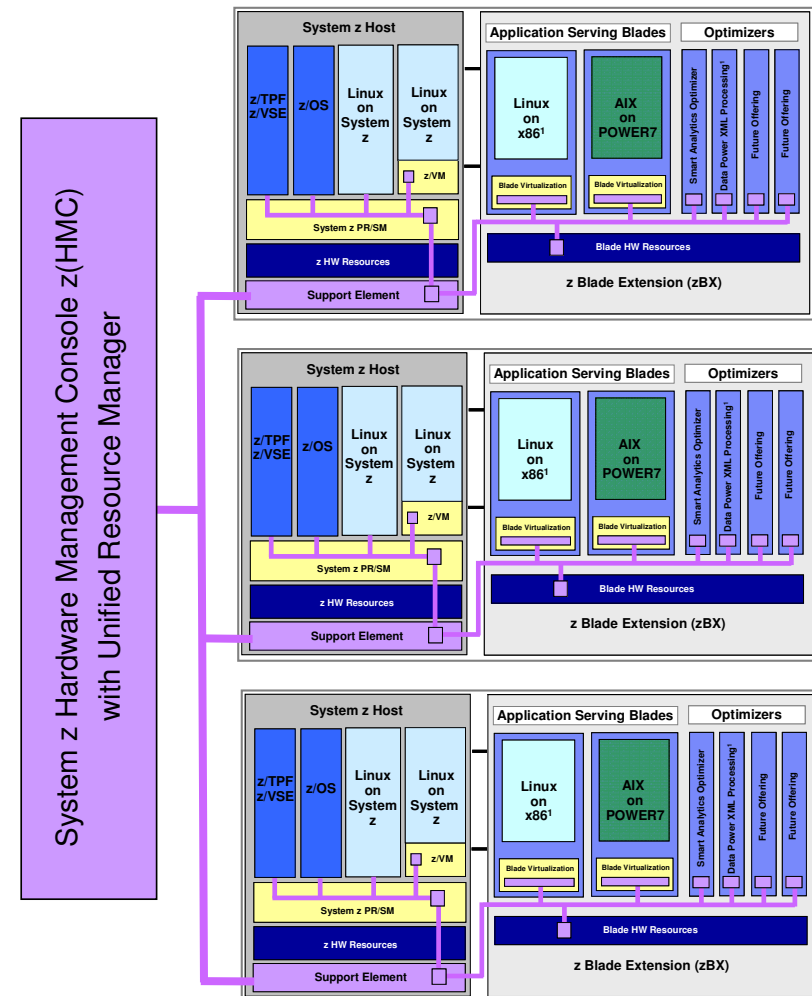
# z196 with zBX (z Blade Extension)



<sup>1</sup> All statements regarding IBM future directions and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

# zEnterprise Ensembles

- 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
  - allows for the management and optimization of a zEnterprise Ensemble as a single resource pool
  - System z Hardware Management Console (HMC) is management console
  - Ensemble-wide scope of responsibility



# zEnterprise Unified Resource Manager (zManager)

- What is it?
  - Unified Resource Manager provides **workload awareness** to optimize the system resources in accordance with understanding the policies assigned to that particular workload
  - Functions are grouped into two suites of tiered functionality that enable different levels of capability – Manage suite and Automate suite
- How is it different?
  - **Heterogeneous management:** Total systems management across heterogeneous resources
  - **Integration:** Single point of control, common skills for resources, reduced complexity of day to day operations
  - **Monitoring:** New dashboard for CPU resources and energy management
  - **Simplified installation:** Auto discovery and configuration of resources and workloads with single interface
  - **Secure:** Improved network security with lower latency, less hops and less complexity. Improved control of access due to management of hypervisors as firmware
  - **Service and support management:** Virtual machines and blades able to perform hardware problem detection, reporting and call home


# zEnterprise Unified Resource Manager (zManager)

## Hypervisor Management

- Integrated deployment and configuration of hypervisors
- Hypervisors (except z/VM) shipped and serviced as firmware
- Management and control communication between virtual server operating systems and the hypervisor

## Operational Controls

- Auto-discovery and configuration support for new resources
- Cross platform hardware problem detection, reporting and call home
- Physical hardware configuration, backup and restore
- Business management functions including user management, auditing, device status, service network configuration
- Ensemble creation



## Network Management

- Management of virtual networks including access control

## Workload Awareness and Platform Performance Management

- Management of resources in accordance with specified business service level objectives
- HMC provides a single consolidated and consistent view of resources
- Monitor resource use within the context of a business workload
- Define workloads and associated performance policies

## Energy Management

- Monitoring and trend reporting of CPU energy efficiency
- Ability to query maximum potential power
- Static power savings

## Virtual Server Lifecycle Management

- Single view of virtualization across platforms
- Ability to deploy multiple, cross-platform virtual servers within minutes



# z/OS Workload Management

Goal oriented workload management since 1994


**Service Policy**  
**Service Classes**

- CICS, Imp=1, 90% < 0.25s
- DDF, Imp=2, 80% < 0.5s
- TSO, Per=1, Imp=3, 90% < 1s
- TSO, Per=2, Imp=4, 80% < 3s
- Batch, Imp=5, ExVel=10



- Goal oriented approach
- Different classes of service
- Classification rules
- Importance based goal attainment
- Subsystems participate



**z/OS Workload Manager (WLM)**



determine goal achievement



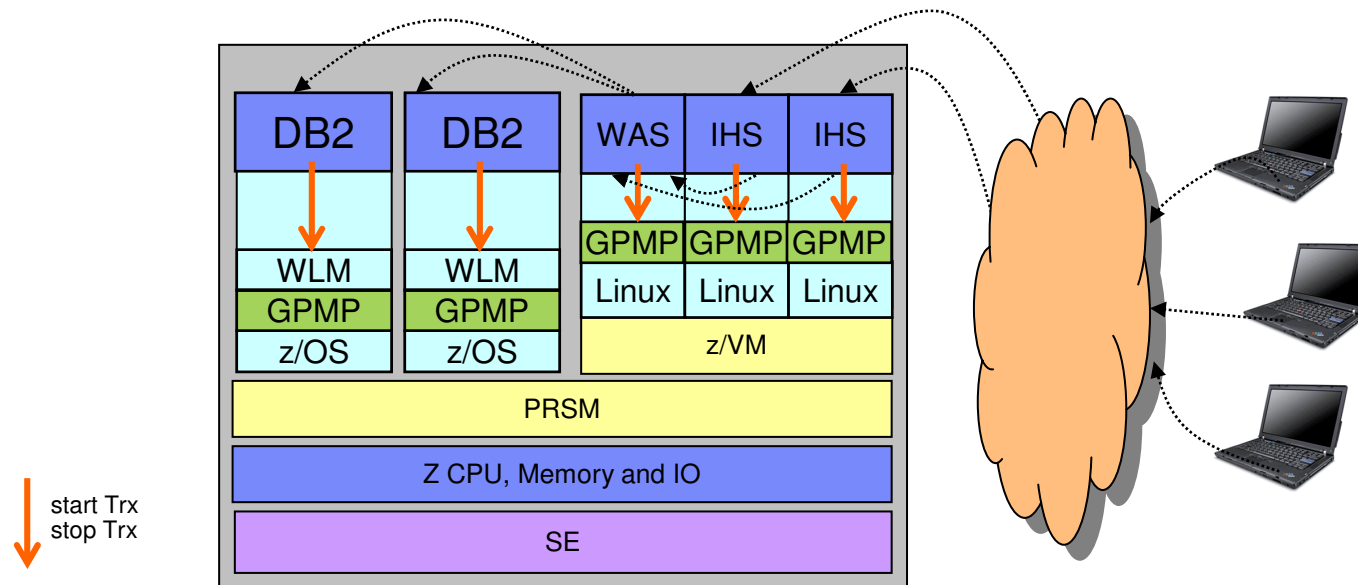


## zEnterprise Platform Performance Manager

- 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
  - 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 functions will evolve over time
  - Pushes management directives to the Support Element, 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 zEnterprise Unified Resource Manager
  - Code structured and packaged as System z firmware
  - Inter-Component communication over trusted internal platform management network

# zEnterprise Platform Performance Manager

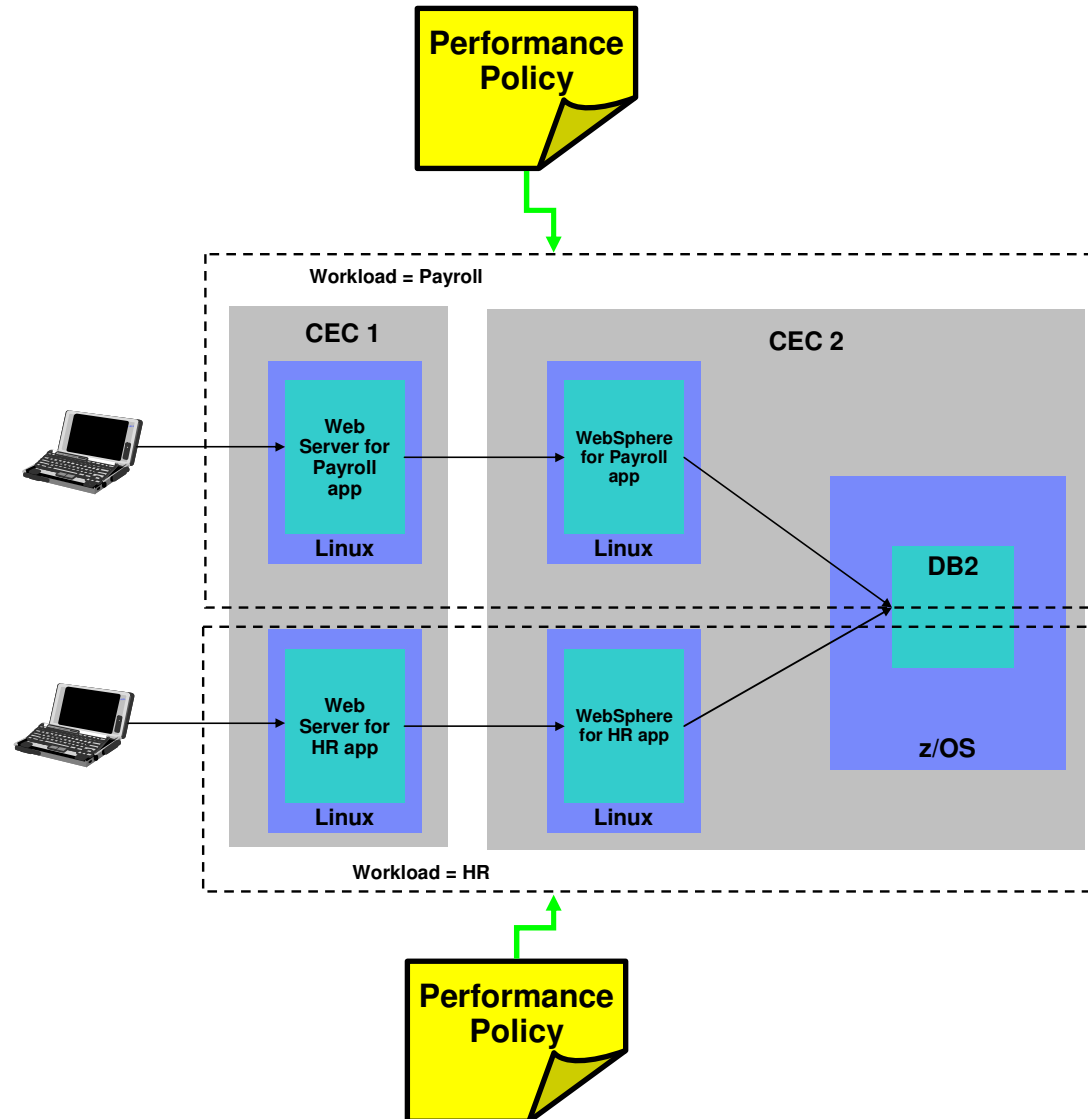
Resource management based on understanding of overall workload flow



- Applications / middleware has to be instrumented with ARM – Application Response Measurement (Open Group Standard) to collect transaction statistics
  - Enables to monitor the flow of transactions
  - Enables to monitor transaction response times and processing statistics
- OS Agent – guest platform management provider (GPMP)
  - is required to identify individual units of work
  - collects data about processes / address spaces and transactions
  - passes data to Platform Performance Manager
  - On z/OS the data is collected by WLM

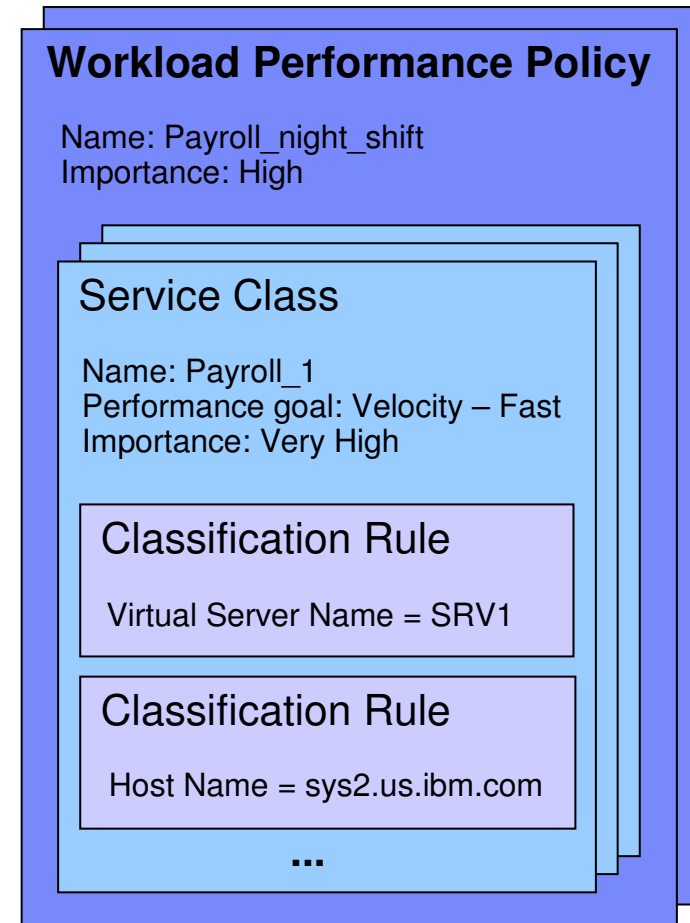
# zManager Platform Workload Definition

- A Platform Workload is a grouping mechanism and “management view” of virtual servers and optimizers 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




# zManager 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



## Goal Achievement is Determined via a Performance Index (PI)

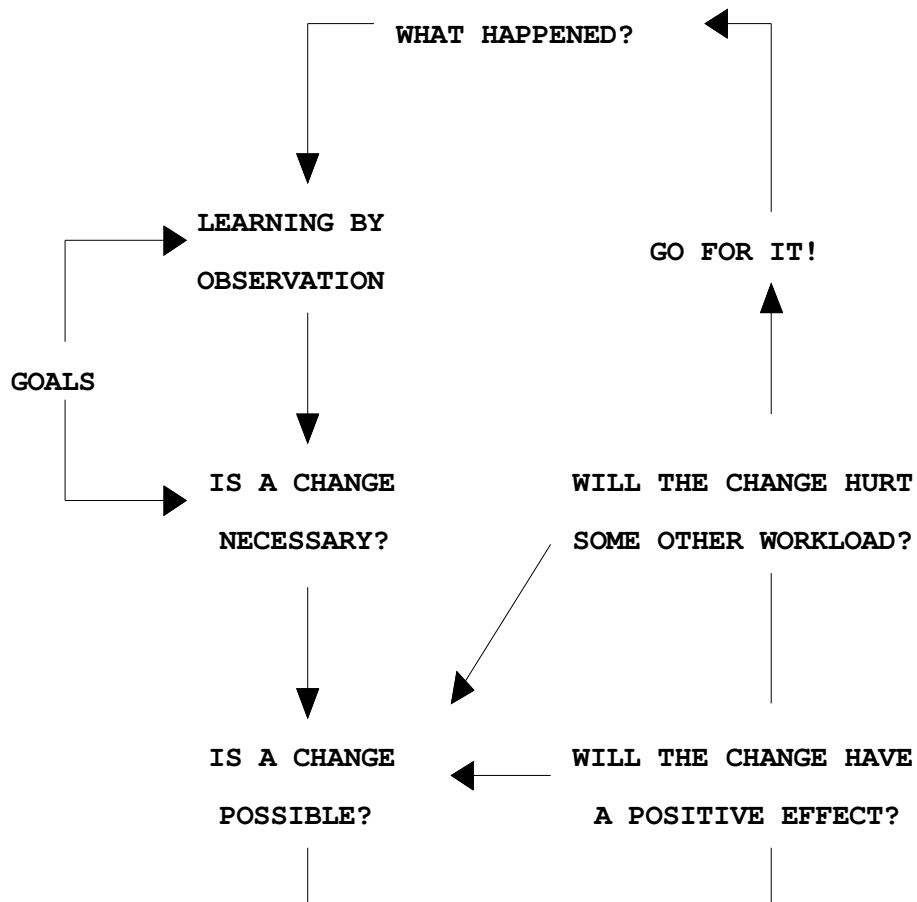
- Indicator how well work (service class) is doing
- Independent from business importance
- Basically
  - If  $PI = 1$ : service class meets goal
  - If  $PI > 1$ : service class is eligible for help
  - If  $PI < 1$ : service class is a potential donor



$$\text{Response Time Goal} : PI = \frac{\text{Actual Achieved Response Time}}{\text{Response Time Goal}}$$

$$\text{Execution Velocity Goal} : PI = \frac{\text{Execution Velocity Goal}}{\text{Actual Achieved Execution Velocity}}$$

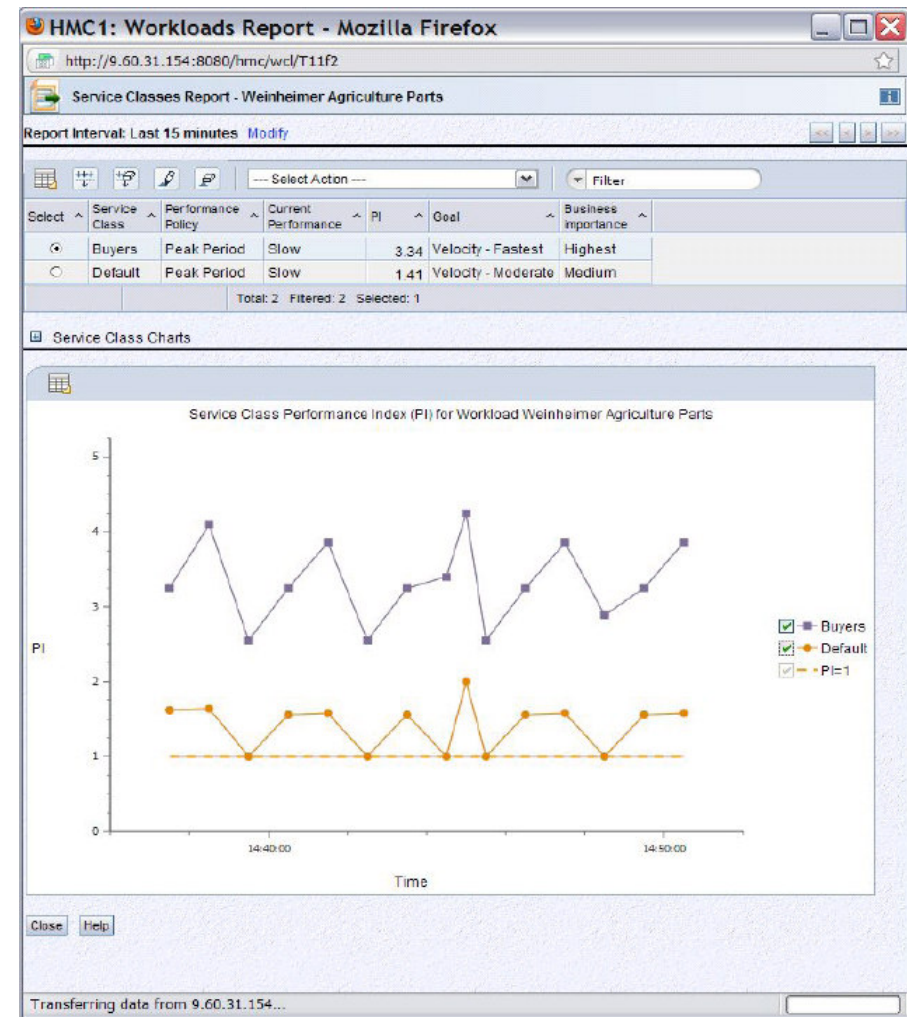
# Workload Management Control Loop



- The Problem
  - Resources must be dynamically adjusted to meet service goals
- Required
  - Seamless adjustment when load changes
  - End user expectations (goals) for important work must always be guaranteed
  - System must be highly utilized (up to 100% for z/OS)
- Solution
  - Periodically measure system state
  - Periodically assess workload state
  - Adjust access of Service Classes to resources (CPU, storage, IO) to help workloads by importance
- In each adjustment cycle resource access is moved from one donor service class to one receiver service class

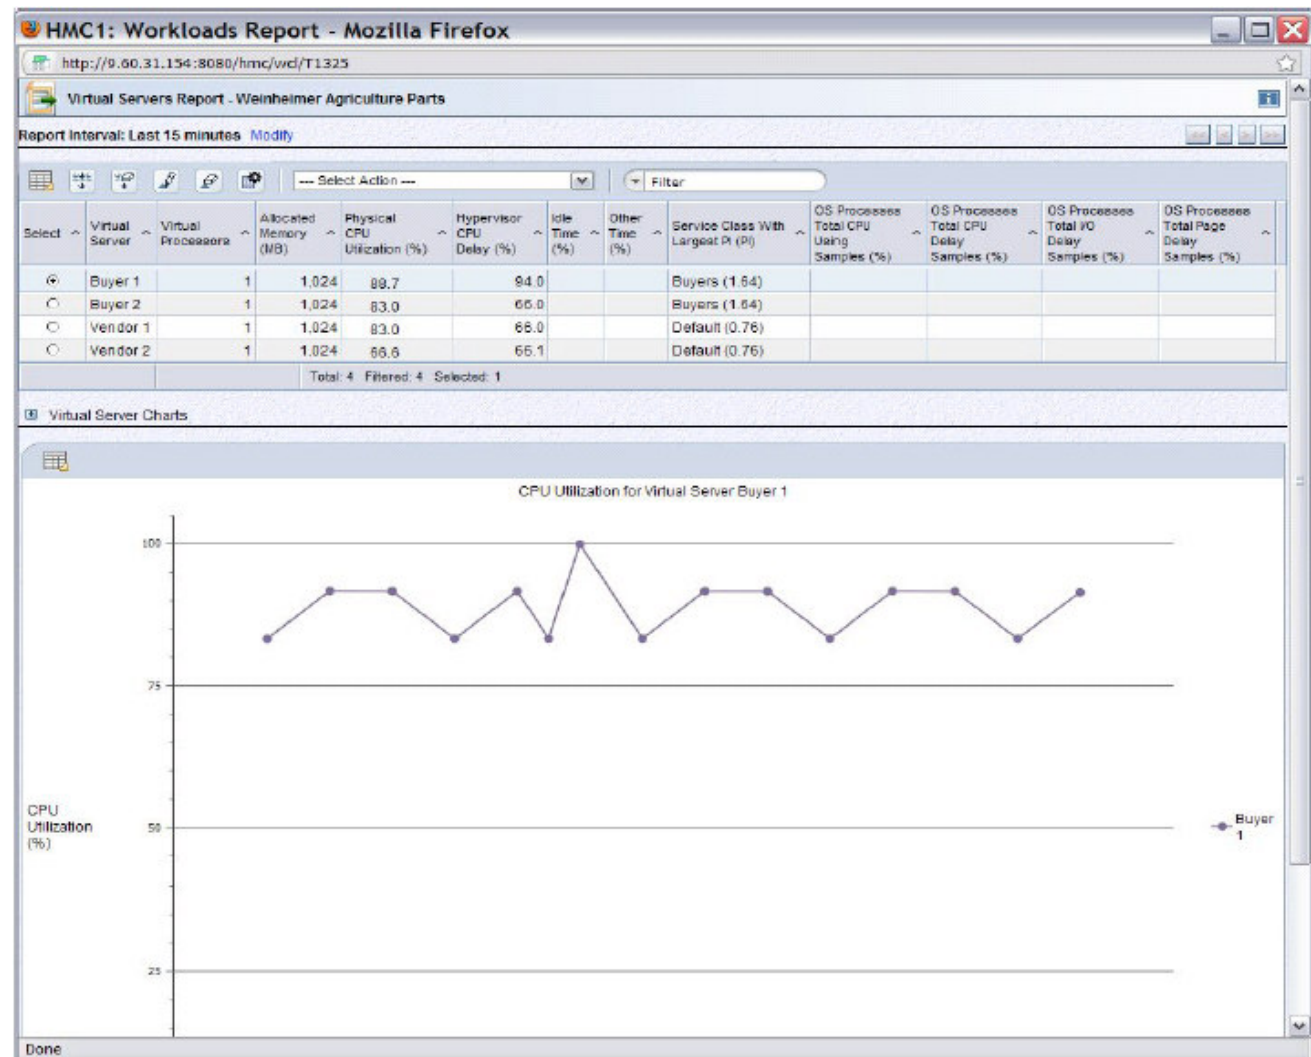
# zManager Workload Based Monitoring and Reporting

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



# zManager – Workload Virtual Server Report

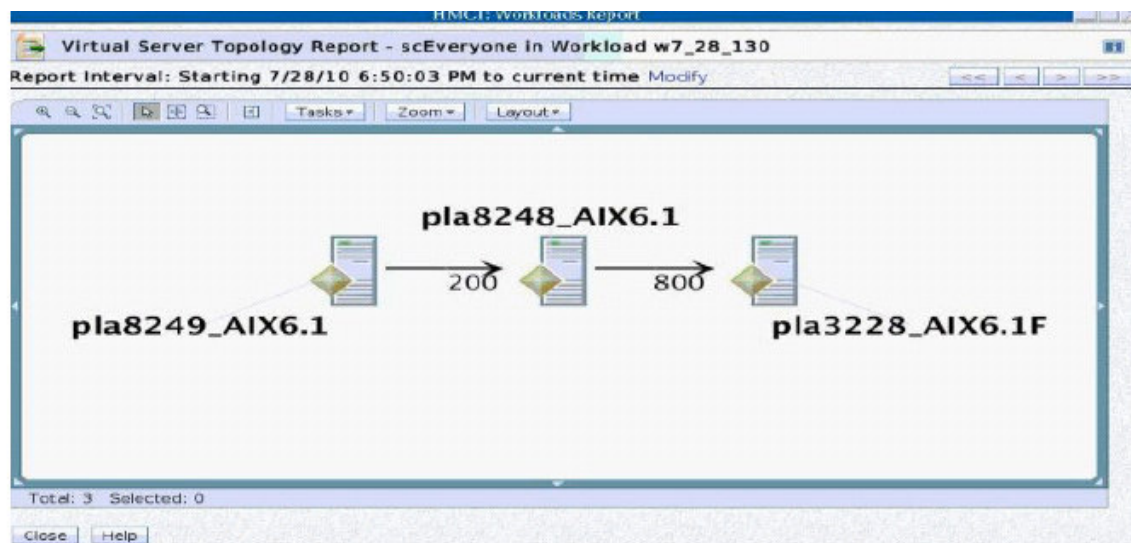
- Workload Virtual Server Report lists virtual server in a service class
  - Virtual server velocity
  - Physical resource utilization
  - OS view of resource utilization
  - Physical memory used
  - Hypervisor delay percentage





# zManager – Transaction Topology and Hops Report

- Topology of virtual servers
- Transaction statistics

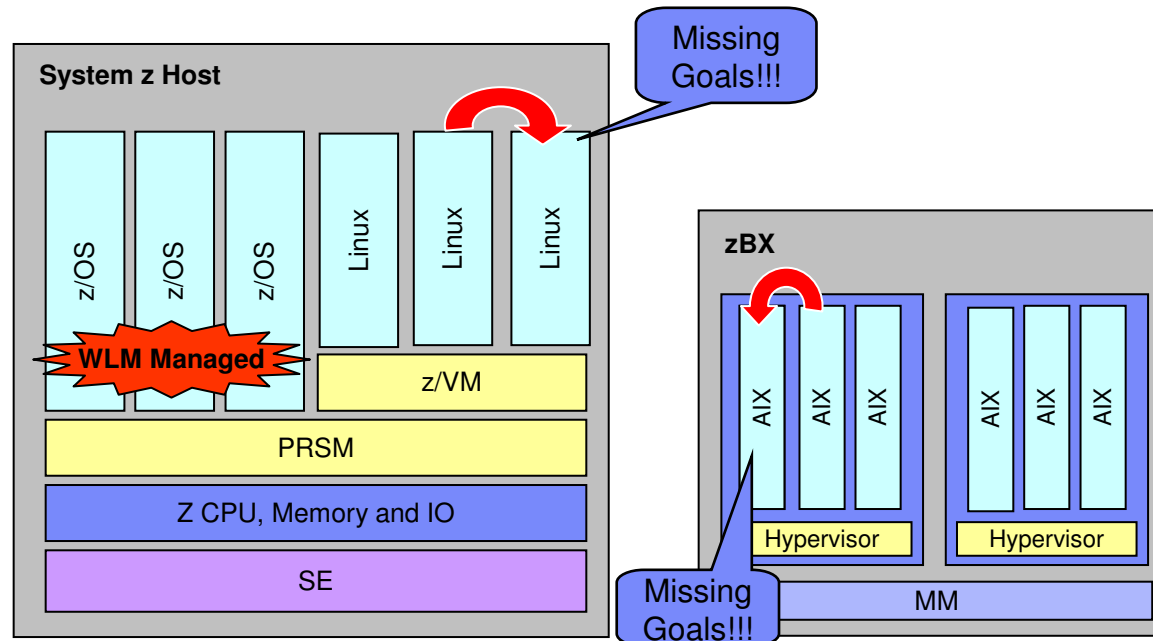


**Hops Report - scEveryone in Workload w7\_28\_130**  
 Report Interval: Starting 7/28/10 6:50:03 PM to current time

**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	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	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

# zManager – Managing Resources across Virtual Servers on P7 Blade and z/VM Virtual Machines



- 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

# zManager – Reporting of Resource Adjustment Actions

**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

- Resource Adjustment Report displays movement of processing capacity between service classes

---

## Summary

- zEnterprise Unified Resource Manager provides workload awareness to optimize the system resources in accordance with understanding the policies assigned to a particular workload
- zEnterprise Platform Performance Manager extends z/OS goal oriented workload management concepts across the platform / ensemble
- Workload based goal oriented policy definition
- Monitoring and reporting in context of Workload Performance Policy
- Goal oriented resource management
  - Extending Intelligent Resource Director management across environments of the platform
  - Management functions will evolve over time