

SHARE Winter 2013 San Francisco



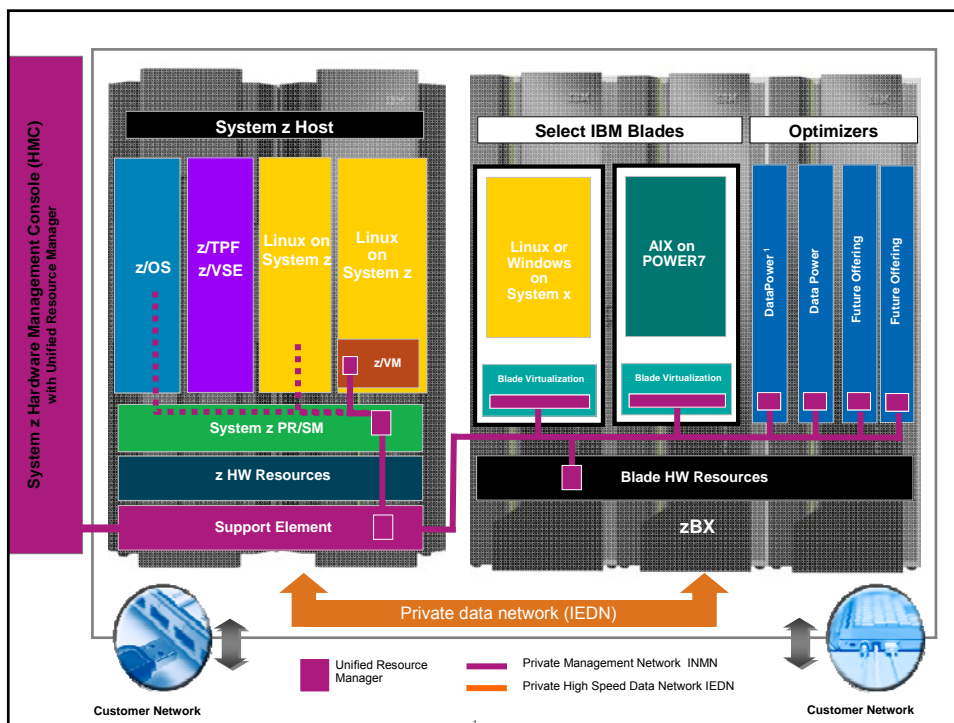
Application Performance Management and Capacity Planning for IBM zEnterprise Hybrid Workloads

Session 12946

Glenn Anderson, IBM Technical Training



© 2013 IBM Corporation



IBM

Agenda

- zEnterprise Workload Management
 - z/OS Virtual Servers
 - WLM and IRD
 - z/VM Guests and zBX Blade Virtual Servers
 - PPM

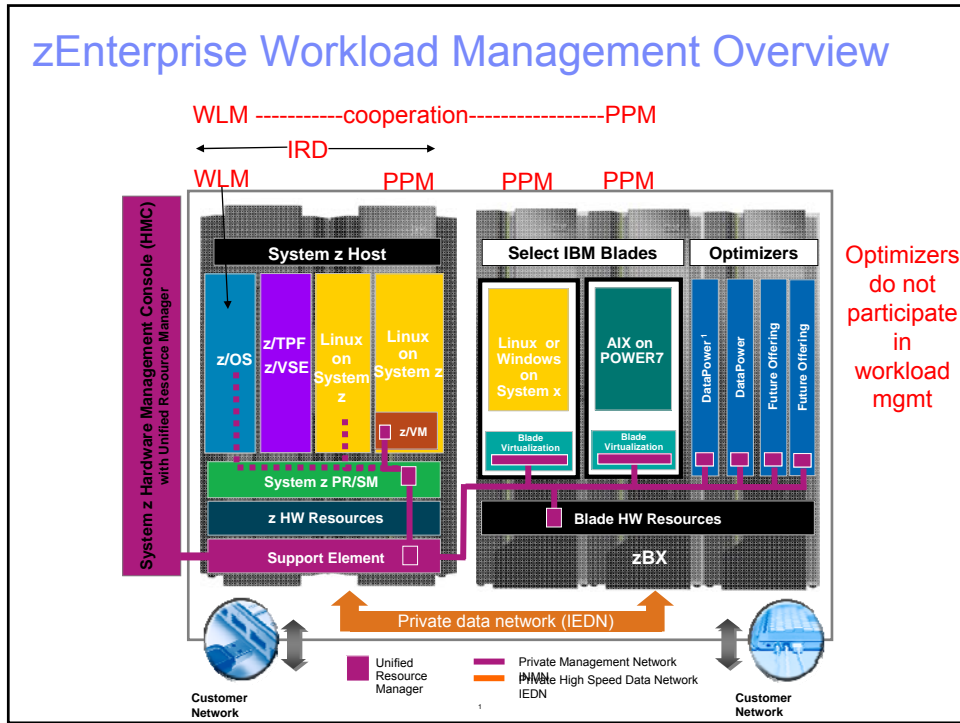
- zEnterprise Resource Monitoring
 - z/OS Virtual Servers
 - RMF
 - z/VM Guests and zBX Blade Virtual Servers
 - PPM
 - RMF XP

© Copyright IBM Corporation 2012

IBM

zEnterprise Workload Management

© Copyright IBM Corporation 2012



Platform Performance Manager

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
 - Extend goal oriented approach of z/OS WLM to platform managed resources
 - Common approach to monitoring / management of platform resources across zEnterprise
 - Orchestration of autonomic management of resources across virtual servers
 - Provide Intelligent Resource Director like function across the zEnterprise
 - 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

© Copyright IBM Corporation 2012

PPM Components

- **HMC**
 - HMC is management server and console
 - Provides ensemble wide aggregation of performance data
 - UI for defining workloads, performance policy and reporting data
 - Pushes management directives to all the nodes of ensemble
- **Support Element (SE)**
 - Provides node (or CPC) level aggregation of performance data
 - Pushes management directives to all the hypervisors in the node.

The diagram illustrates the PPM components. On the left, a vertical purple bar represents the HMC. To its right, a yellow box represents the PR/SM layer, which sits on top of a blue box representing Z CPU, Memory and IO. Below that is a purple box for the SE. The main part of the diagram shows a stack of components: z/OS, z/OS, z/OS, Linux, Linux, z/VM Mgmt Guest, z/VM, Power7, System x, AIX, AIX, VIOS, Linux, Windows, Power VM, xHyp, Optimizer, and Optimizer. At the bottom, there is a blue box for AMM and a grey box for z Blade Extension.

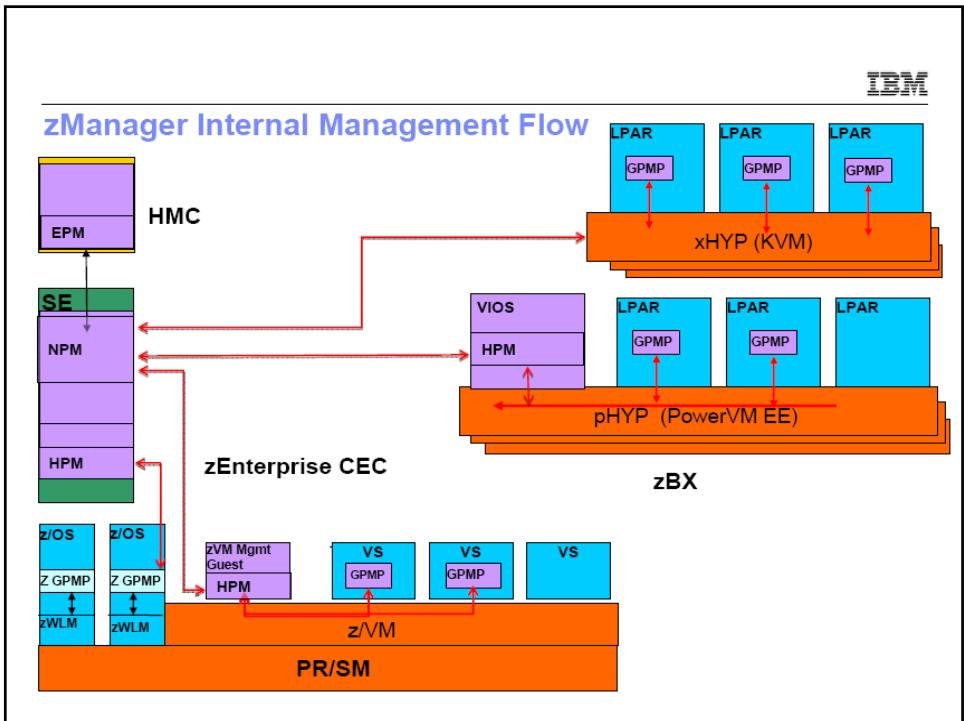
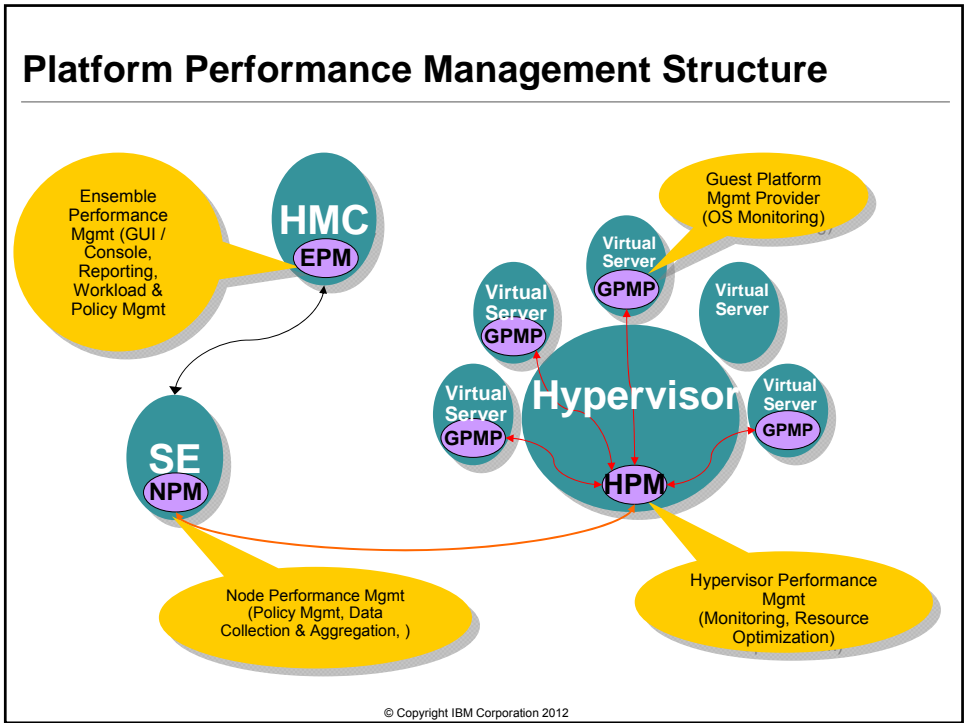
© Copyright IBM Corporation 2012

PPM Components

- **Hypervisors**
 - Monitors goal defined in performance policy and performs dynamic resource mgmt (z/VM and Power VM) to achieve performance goal where applicable
 - Collects virtual server statistics from hypervisor and guest platform management providers. Pushes aggregated metrics to SE
- **Virtual Servers**
 - **Optional** Guest Platform Management Provider software deployed in Virtual Server
 - Collects monitoring data from Operating system and ARM instrumented applications and pushes to hypervisors.

The diagram is identical to the one above, showing the HMC, SE, PR/SM, z/VM, z/OS, Linux, z/VM Mgmt Guest, Power7, System x, AIX, VIOS, Linux, Windows, Power VM, xHyp, Optimizer, and AMM components.

© Copyright IBM Corporation 2012

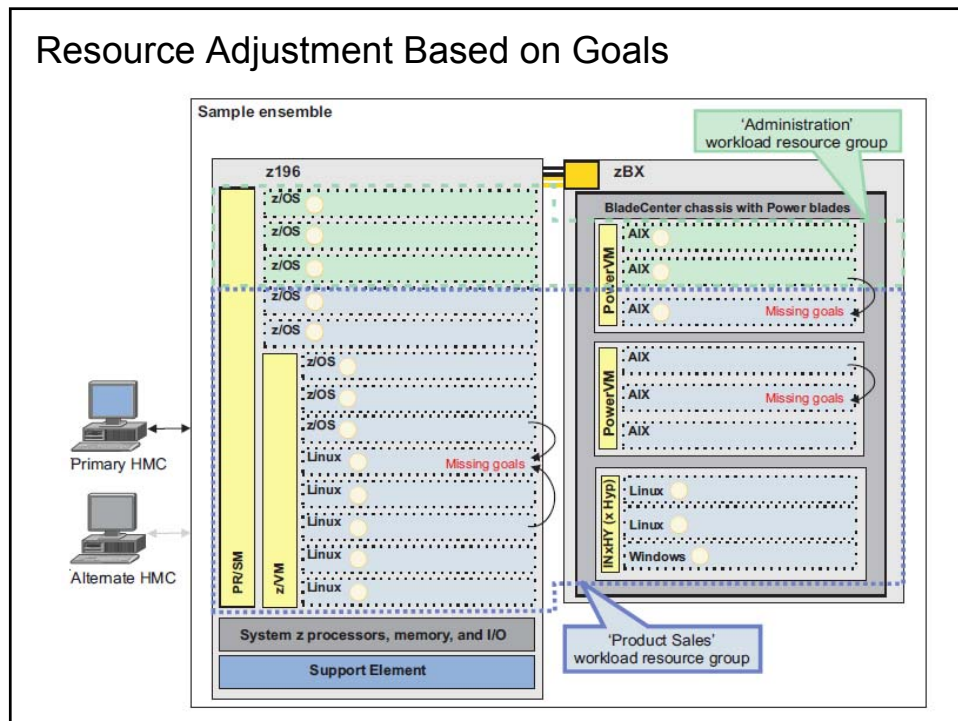


zManager CPU Resource Mgmt Function

- z/VM and PowerVM Hypervisors
 - Virtual Server CPU Management provides the ability to manage CPU resources across virtual servers based on a goal-oriented performance policy.
- System x (KVM based) Hypervisor
 - Does not currently participate in dynamic resource management
 - Statement of Direction (8/28/12): IBM intends to deliver workload-aware optimization for IBM System x blades in the zBX, allowing virtual CPU capacity to be adjusted automatically across virtual servers within a hypervisor
- PR/SM Hypervisor
 - Does not make resource management adjustments based on PPM Policy. Only IRD dynamically influences the PR/SM hipervisor

© Copyright IBM Corporation 2012

Resource Adjustment Based on Goals



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

The diagram illustrates a Platform Workload architecture. It features two Central End Cells (CEC 1 and CEC 2). Each CEC contains a Web Server for Payroll app (Linux) and a WebSphere for Payroll app (Linux). Similarly, each CEC contains a Web Server for HR app (Linux) and a WebSphere for HR app (Linux). Both workloads share a common DB2 database (z/OS). Performance Policies are associated with each workload, as indicated by yellow callout boxes labeled 'Performance Policy' pointing to the 'Workload = Payroll' and 'Workload = HR' groups.

© Copyright IBM Corporation 2012

Workload Performance Policy...

The diagram shows the structure of a Workload Performance Policy. It is a nested hierarchy:

- Workload Performance Policy** (outermost layer, light blue)
 - PerformancePolicyName
 - Importance
 - Service Class** (middle layer, light purple)
 - Service Class Name
 - Performance Goal
 - Importance
 - Classification Rule** (innermost layer, light pink)
 - Virtual Server Name, OS Name etc....

- 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

© Copyright IBM Corporation 2012

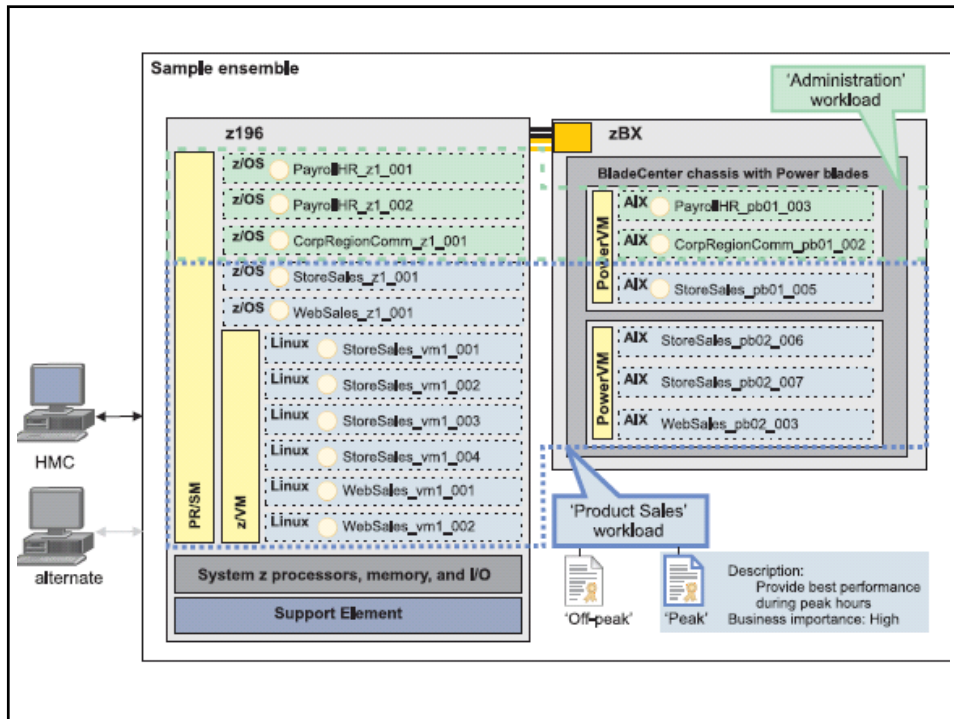
Elements of a Service Class

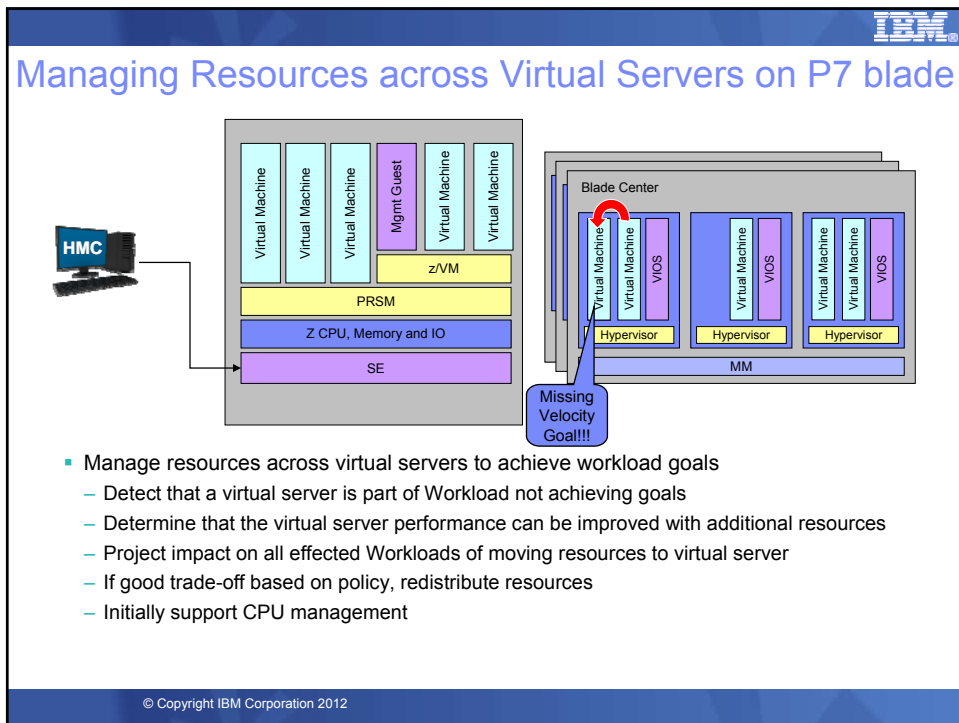
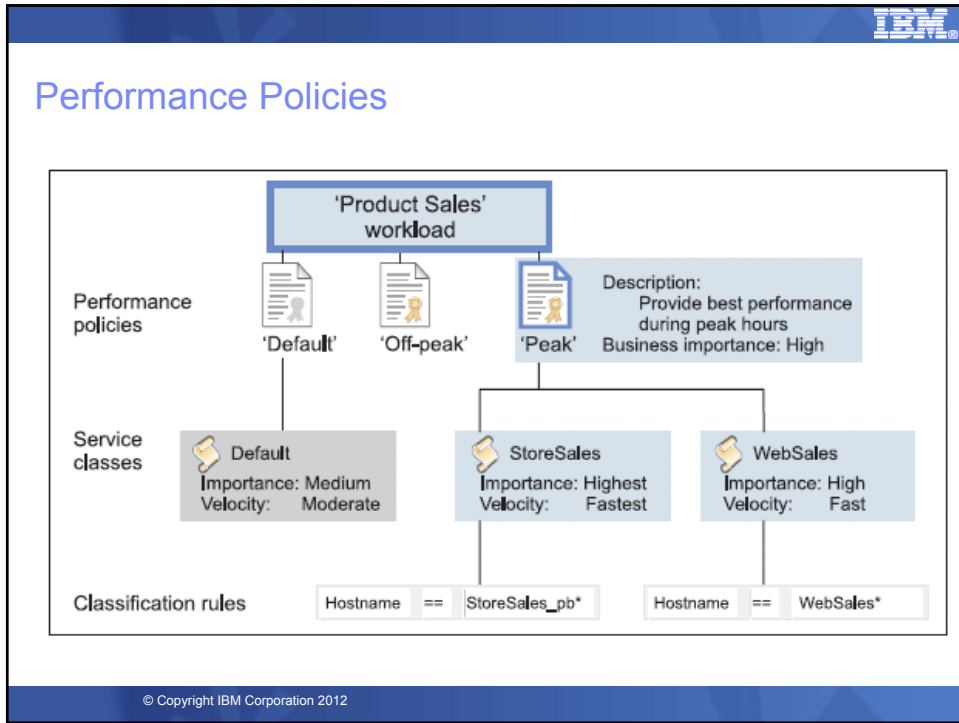
- **Performance Goal (managed at the virtual server level)**
 - Velocity: Fastest, Fast, Moderate, Slow, Slowest
 - Discretionary: No performance goal

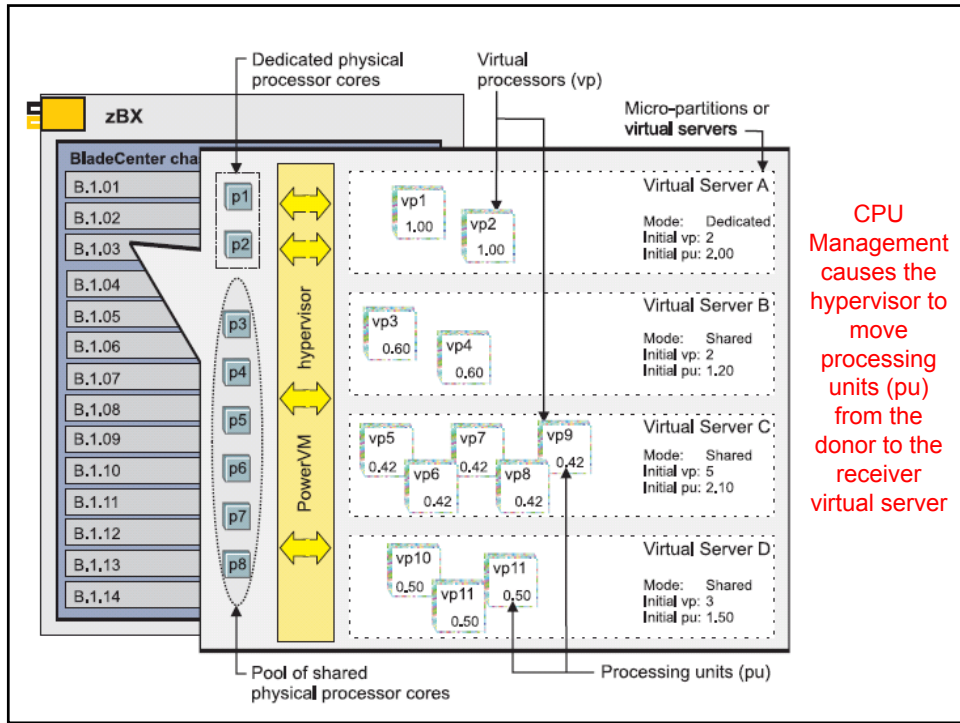
- **Business Importance: Highest, High, Medium, Low, Lowest**

- **Classification Rule**
 - Use Virtual Server Name as qualifier to assign Service Class
 - Virtual Servers under the PR/SM and System x hypervisors should be classified into a Service Class for resource monitoring purposes

© Copyright IBM Corporation 2012





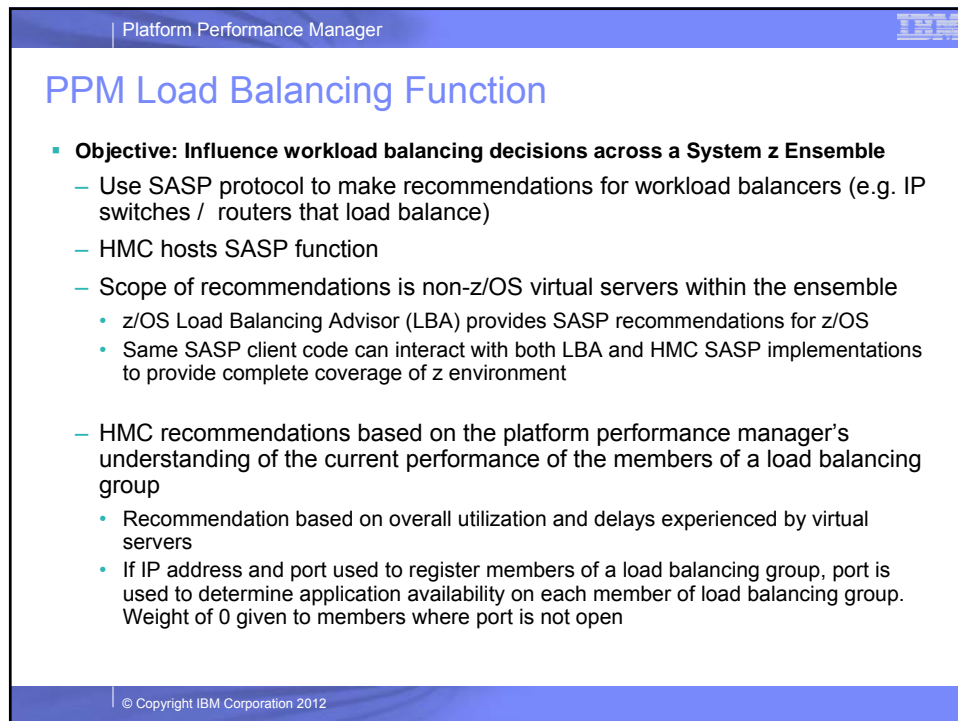
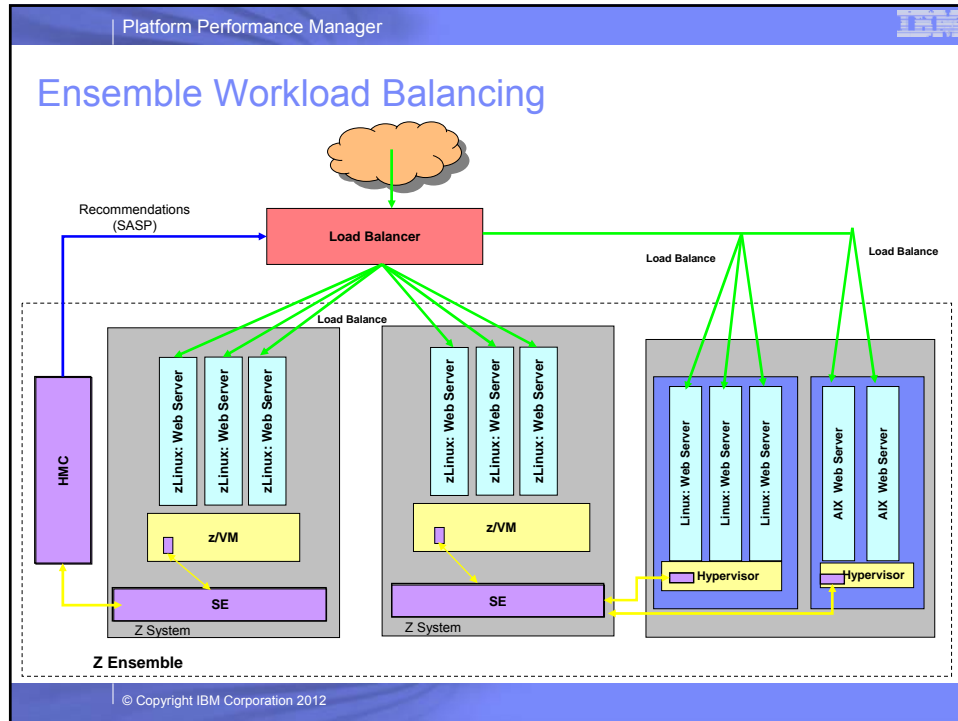


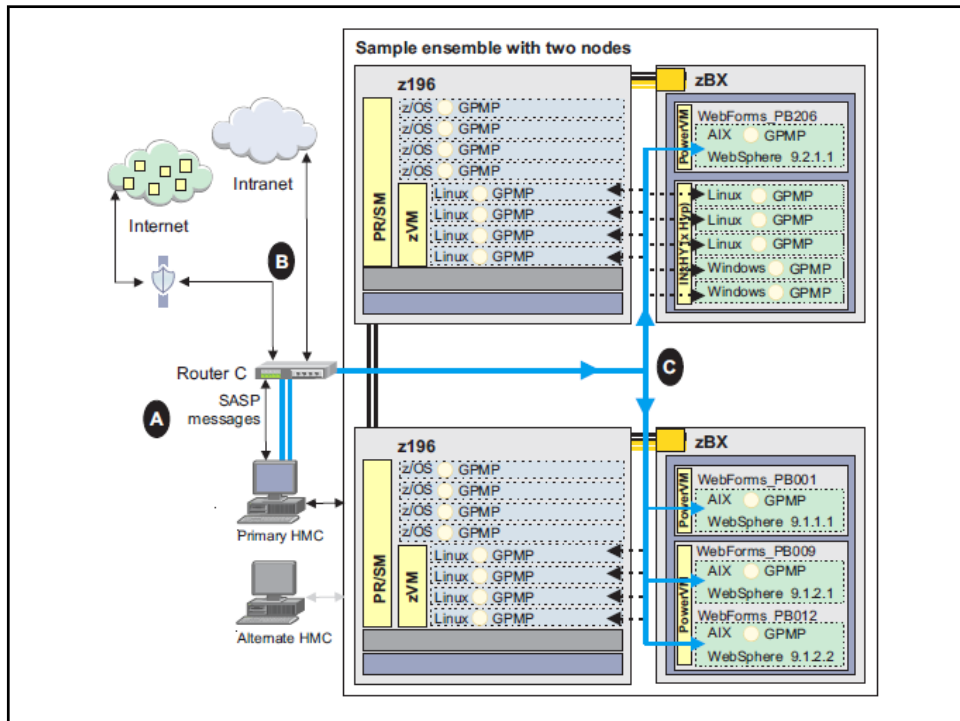
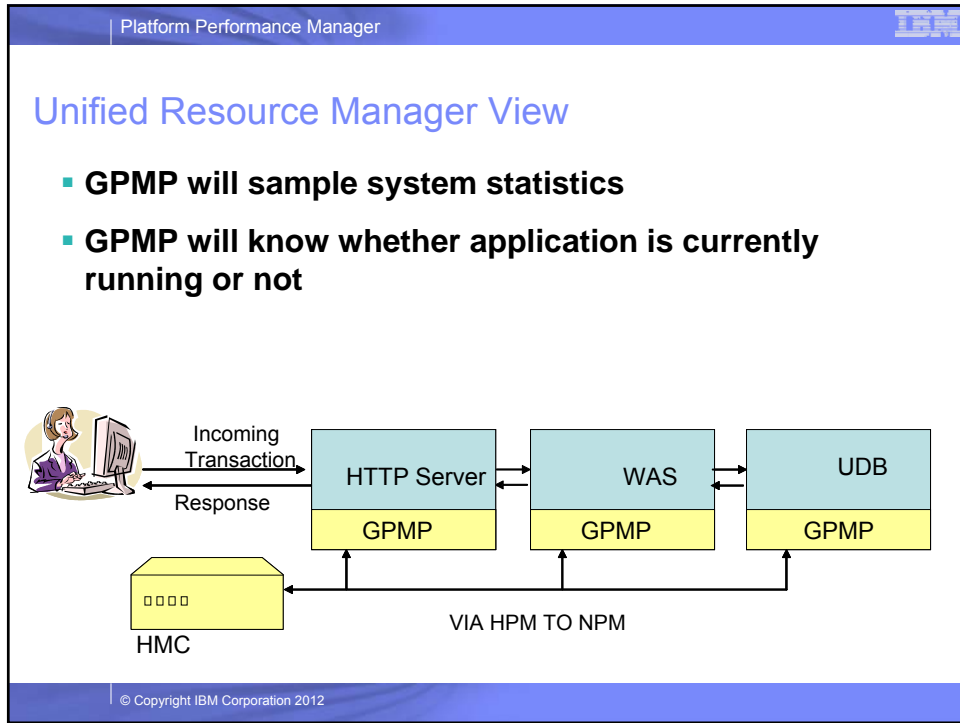
Managing Resources across z/VM Virtual Machines

The diagram shows the resource management stack. At the top is the HMC (Hardware Management Console). Below it is the z/VM layer, which includes Virtual Machines, a Mgmt Guest, and VIOS. This is managed by PRSM (Resource Management Subsystem). Below PRSM is the Z CPU, Memory and IO layer, and at the bottom is the SE (System Environment) layer. A callout bubble points to the z/VM layer with the text 'Missing Velocity Goal!!!'. To the right, a Blade Center is shown with its own set of Virtual Machines, VIOS, Hypervisor, and MM (Memory Manager) layers.

- 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

© Copyright IBM Corporation 2012



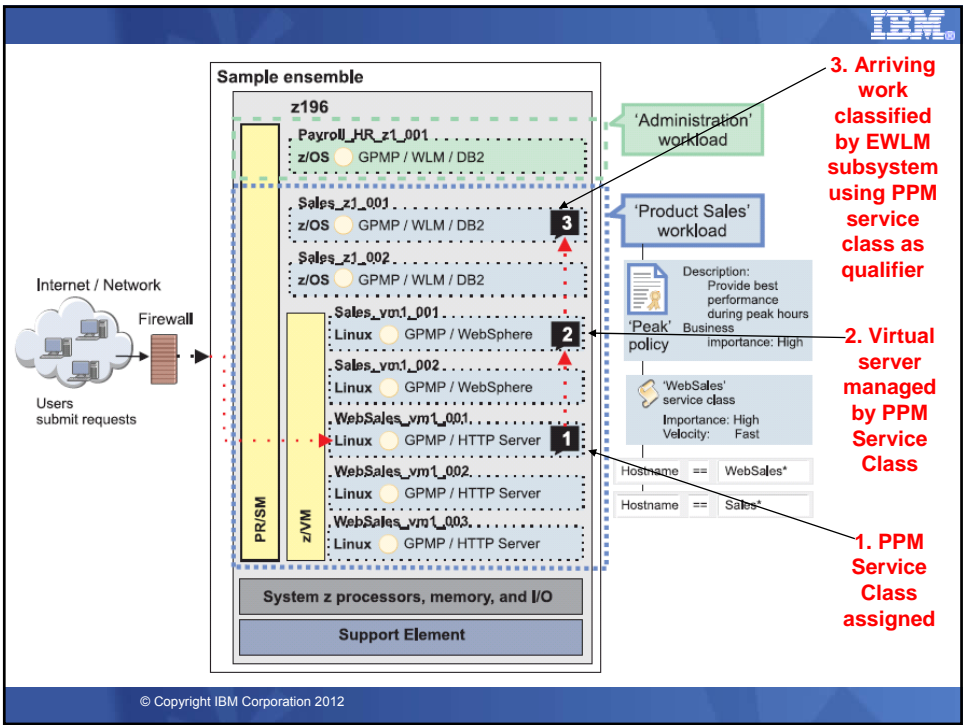



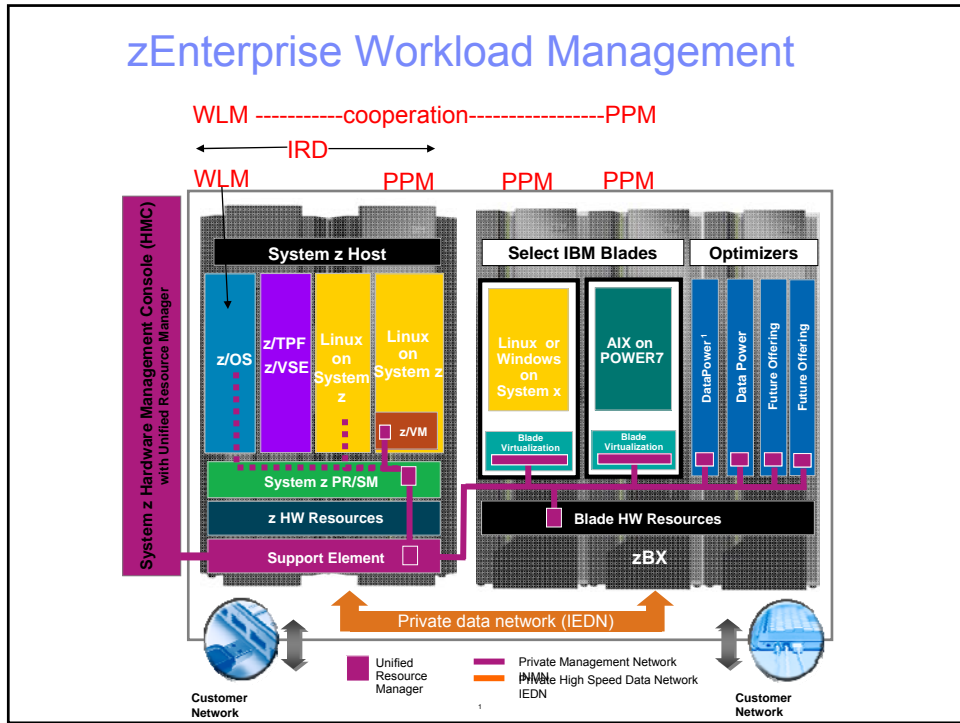
IBM

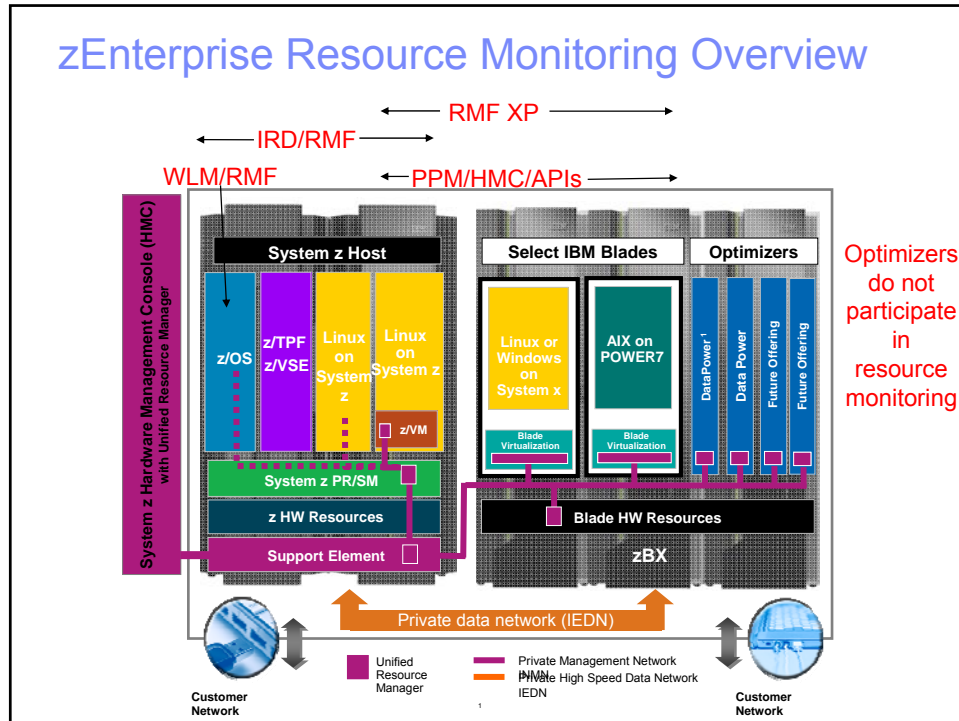
Co-operative management with z/OS WLM

- z/OS provides differentiated service to PPM classified work
- Transaction coming to z/OS needs to be ARM instrumented via Guest Platform Management Provider (GPMP) implementation
- WLM service definition needs to map PPM service classes to z/OS WLM service classes via EWLM classification rules
- PPM service class associated with transaction is used by WLM to classify work unit to a different WLM service class.
- WLM manages the resources based on the goal assigned to this specific service class.

© Copyright IBM Corporation 2012







PPM 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

IBM

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

© Copyright IBM Corporation 2012

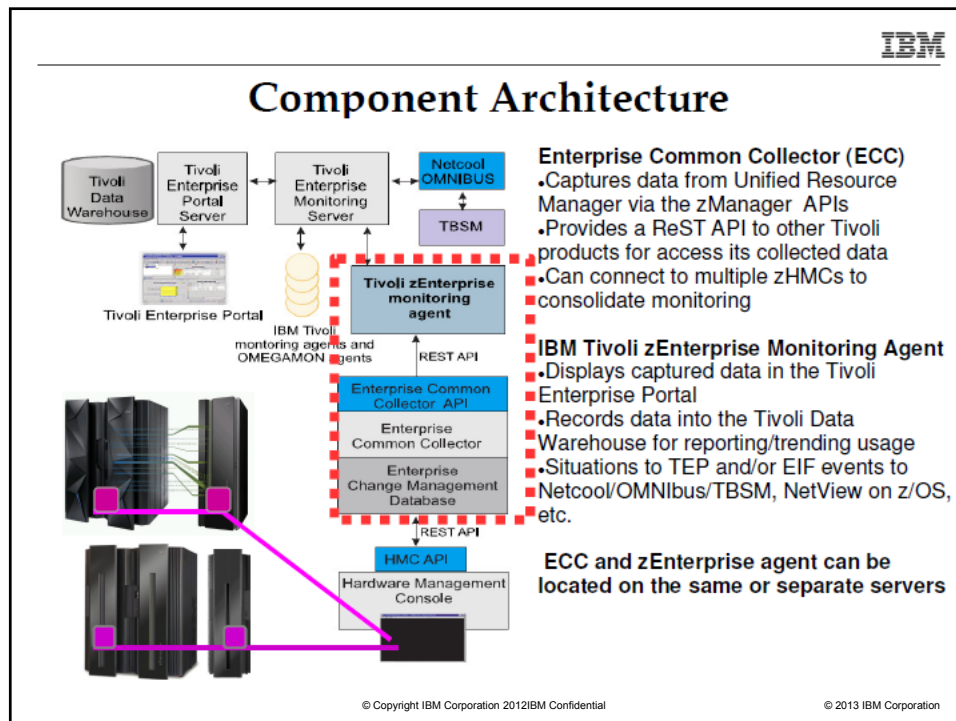
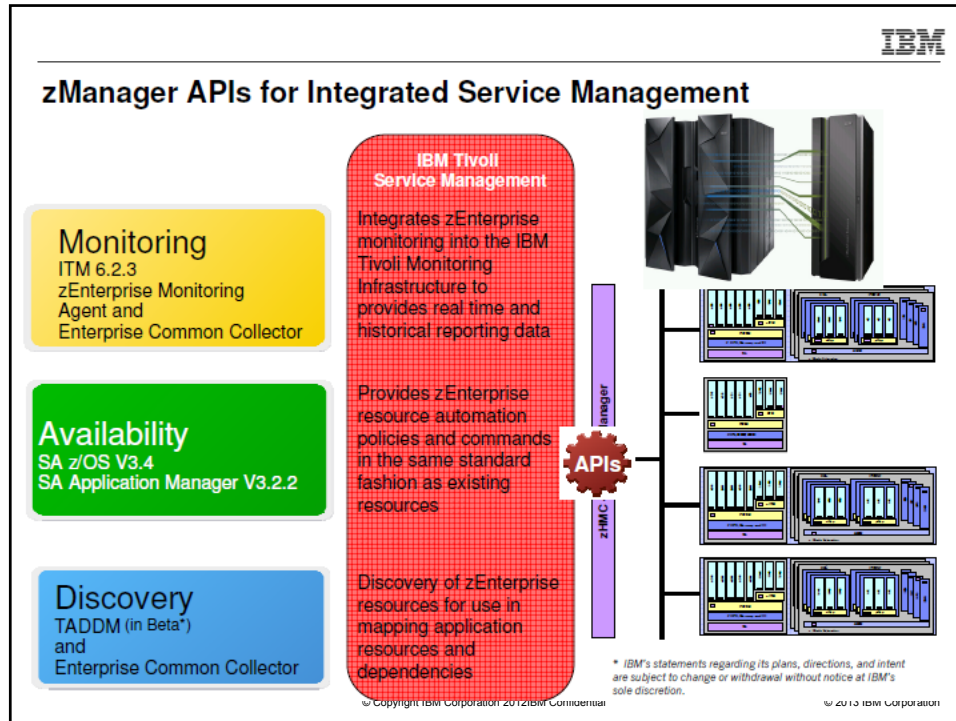
IBM


Unified Resource Manager APIs *Enabling External Management Tools*

- New API support allows programmatic access to the same underlying functions exploited by the HMC user interface (UI)
 - ▶ Same resource types, instances and policies
 - ▶ API functions corresponding to views and tasks in the UI
 - Listing resource instances
 - Creating, changing, deleting resource instances
 - Operational control of resource instances
- Access to functions will enable management of Unified Resource Manager from external (to HMC) tools
- Initially the priority scenarios will be the discovery, monitoring, and provisioning use cases

zEnterprise System

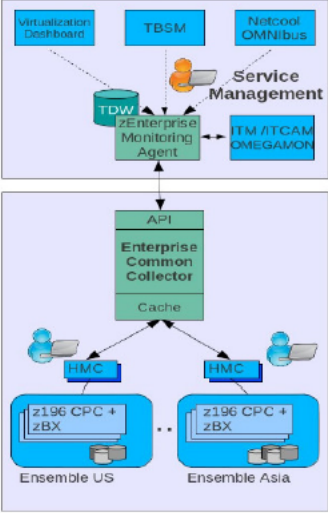
© Copyright IBM Corporation 2012 IBM Confidential © 2012 IBM Corporation






zEnterprise Monitoring Agent

- Provides visibility into the IBM zEnterprise hybrid infrastructure, including hardware resources, hypervisors, virtual servers, and workload resource groups
- Automated discovery of resources within the monitoring environment
- Integrates the information into IBM Tivoli Monitoring infrastructure to provide:
 - ✓ Information highlighting and alerting capabilities
 - ✓ Integration with data from other agents in the ITM infrastructure
 - ✓ Optional long term history collection and reporting/trending with the Tivoli Data Warehouse and Tivoli Common Reporting
 - ✓ Optional event integration and Business Service Management with Netcool/OMNIBus and Tivoli Business Service Manager (TBSM)



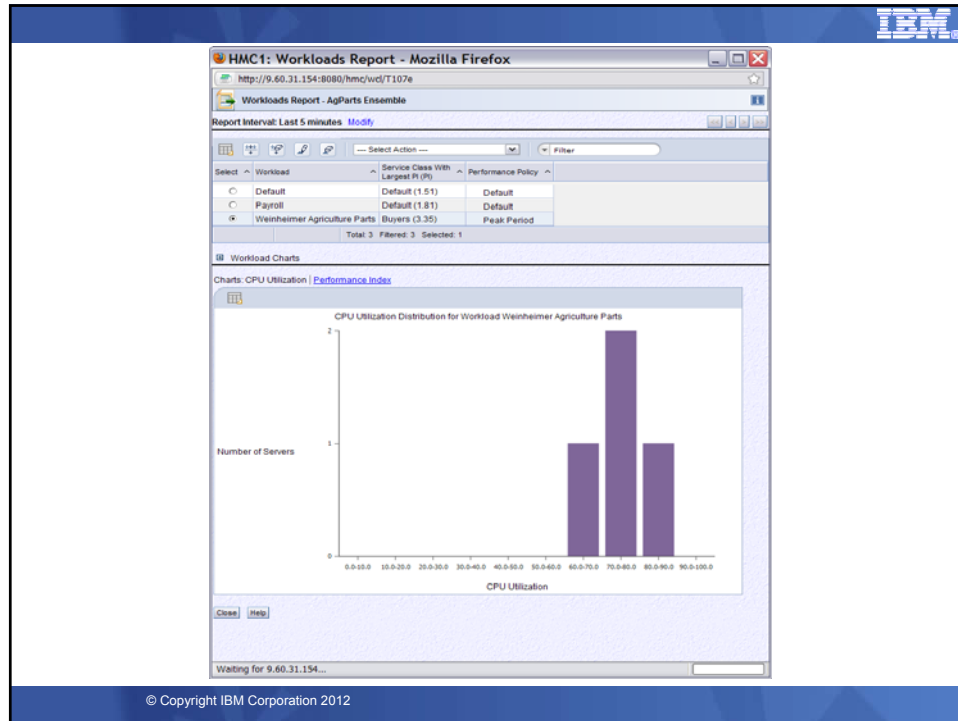
© Copyright IBM Corporation 2012 IBM Confidential © 2013 IBM Corporation



Workload Monitoring Overview...

- 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

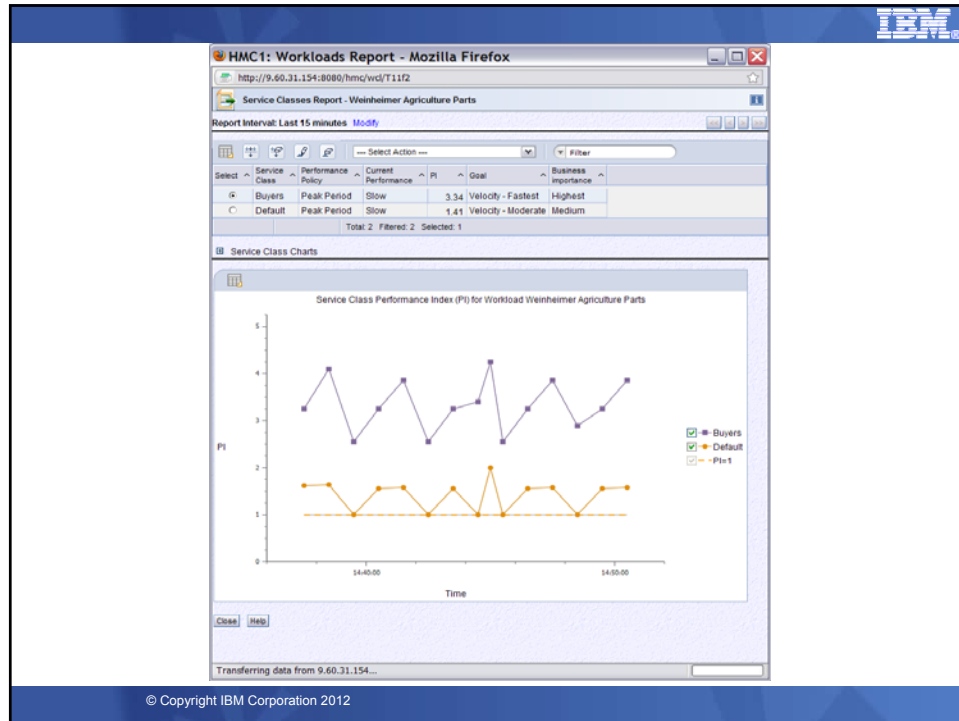
© Copyright IBM Corporation 2012



Workload Monitoring Overview...

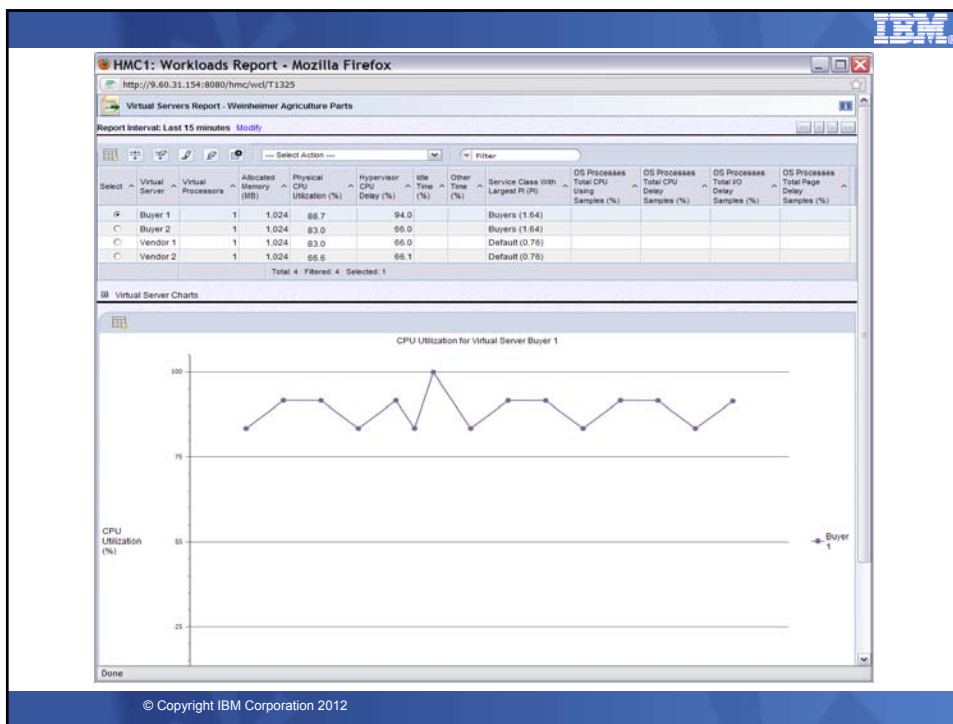
- Service Class Report
 - High-level view of each service class in Workload's performance policy
 - Goal and importance
 - Actual performance
 - Indication if monitoring event is established for service class and event is triggered
 - Service class details
 - Graph of service class performance index
 - 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

© Copyright IBM Corporation 2012



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



HMC1: Workloads Report - Mozilla Firefox
 http://9.60.31.154:8080/hmc/wcd/T1400
 Virtual Server Resource Adjustments Report - Buyer 1
 Report Interval: Last 15 minutes

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				

© Copyright IBM Corporation 2012

PPM Hypervisor Report

R90HMC1: Virtual Servers Report - Mozilla Firefox
 https://9.12.16.241/hmc/inv/743de4tableTop_6b66c6d6

Processor count: 8 Total CPU consumption: 4.6%
 Total memory allocated for LPARs: 32,768 MB Total memory: 65,536 MB
 Total processor entitlement: 6.84

Virtual Servers

Virtual Server	Processor Management Status	Processor Management Reason	Virtual Processor Count	Consumed Processors	Hypervisor Processing Unit Delay (%)	Allocated Memory (MB)	LPAR Capped	Uncapped Weight	Current Entitled Capacity	Defined Entitled Capacity	Min Entitled Capacity
r90f1b207v2	Active	None	2	0.01	0.0	4,096	--	128	180	180	10
r90f1b207v3	Active	None	2	0.02	0.0	4,096	--	128	20	20	20
r90f1b207v4	Active	None	2	0.01	0.0	4,096	--	128	139	139	10
r90f1b207v5	Active	None	2	0.01	0.0	4,096	--	128	140	140	10
Total: 8 Filtered: 8											

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
r90f1b207v7	WkldForModerateMedium	SrvClsForModerateMedium	0.45 (0.25)	r90f1b207v2	Default	1.72 (1.80)	Sep 29, 2010 11:33:19 AM
r90f1b207v7	WkldForModerateMedium	SrvClsForModerateMedium	0.45 (0.25)	r90f1b207v4	Default	1.35 (1.39)	Sep 29, 2010 11:33:19 AM
r90f1b207v7	WkldForModerateMedium	SrvClsForModerateMedium	0.45 (0.25)	r90f1b207v5	Default	1.36 (1.40)	Sep 29, 2010 11:33:19 AM
r90f1b207v7	WkldForModerateMedium	SrvClsForModerateMedium	0.45 (0.25)	r90f1b207v6	Default	1.36 (1.40)	Sep 29, 2010 11:33:19 AM
Total: 4 Filtered: 4							

Failed Adjustments

Done

© Copyright IBM Corporation 2012

Benefits of GPMP

- Guest Platform Management Provider (GPMP) is a lightweight component of PPM that provides additional monitoring data
- Allows cooperative management with z/OS WLM
- Allows virtual server to be classified using additional attributes such as HostName, SystemName, OS Level etc.
- GPMP provides data for recommendation for load balancing function
- With instrumented middleware support, GPMP provides metrics that allows detailed transaction topology as transaction hops through heterogeneous platforms in zEnterprise

© Copyright IBM Corporation 2012

Benefits of Middleware instrumentation


- Transaction response time reporting
- Multi-tiered work request flow across environments
- Relationship to server resources being consumed
- Same reasoning lead to instrumentation of z/OS subsystems (CICS, IMS, DB2, etc) for z/OS WLM
- OpenGroup Application Response Measurement (ARM) standards based instrumentation.

© Copyright IBM Corporation 2012

Topology report with GPMP active

Name	Hop Number	Group Name	Successful Transactions	Failed Transactions	Skipped Transactions	Inflight Transactions	Queue Time (s)	Execution Time (s)	Successful Average Response Time (s)
WebSphere APPLICATION_SERVER	0	g1sr011	0	0	0	0	0.000	0.000	0.000
IBM DB2 Universal Database	0	db0smr1	0	0	0	0	0.000	0.000	0.000
WebSphere APPLICATION_SERVER	0	server1	31,195	0	0	11	0.010	0.000	0.157
IBM WebSphere Plugin	0	IBM_HTTP_Server	31,195	0	0	11	0.010	0.000	0.157
WebSphere APPLICATION_SERVER	1	g1sr012	33,307	0	0	7	0.000	0.019	0.053
WebSphere APPLICATION_SERVER	1	server1	33,307	0	0	7	0.000	0.019	0.053
WebSphere APPLICATION_SERVER	2	DB02LOG1	2,454,512	0	0	2	0.000	0.000	0.000
DB2	2	DB02LOG1	2,454,512	0	0	2	0.000	0.000	0.000
Total			13	Filtered	13				


© Copyright IBM Corporation 2012



Cross Platform Performance Monitoring with RMF XP

- The Common Information Model (aka CIM) instrumentation is available for almost all operating systems on this planet
- RMF has the infrastructure already in place to
 - combine performance data from multiple systems to a Sysplex wide view
 - display performance data by means of state-of-the-art graphical frontends
- RMF XP brings these two well-proven things together
- RMF XP supports the following operating systems:
 - AIX on System p
 - Linux on System x
 - Linux on System z

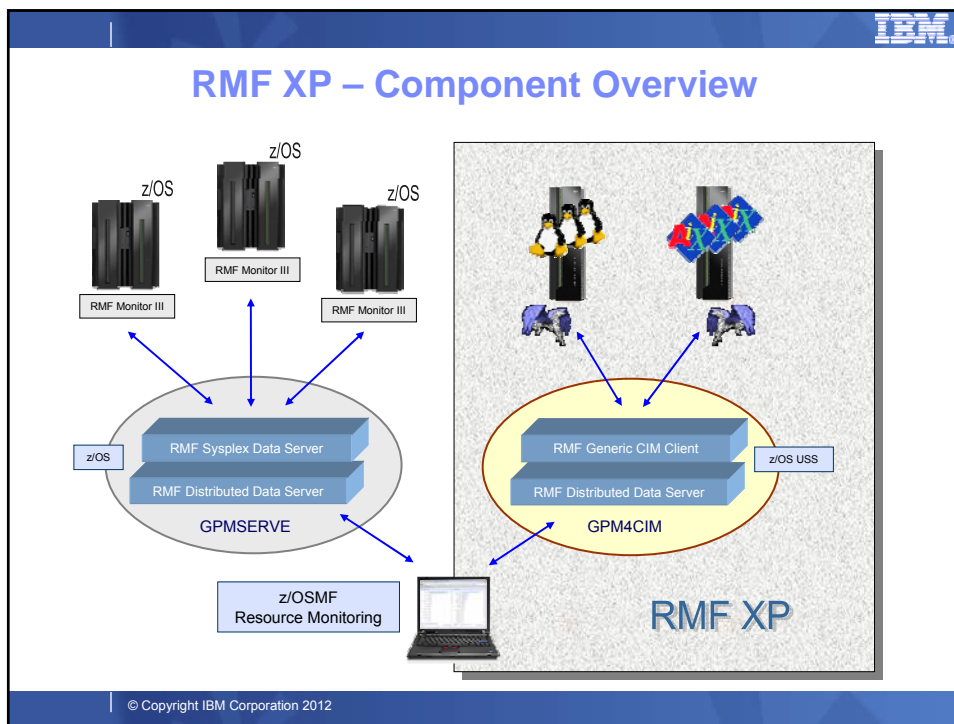
© Copyright IBM Corporation 2012



RMF XP

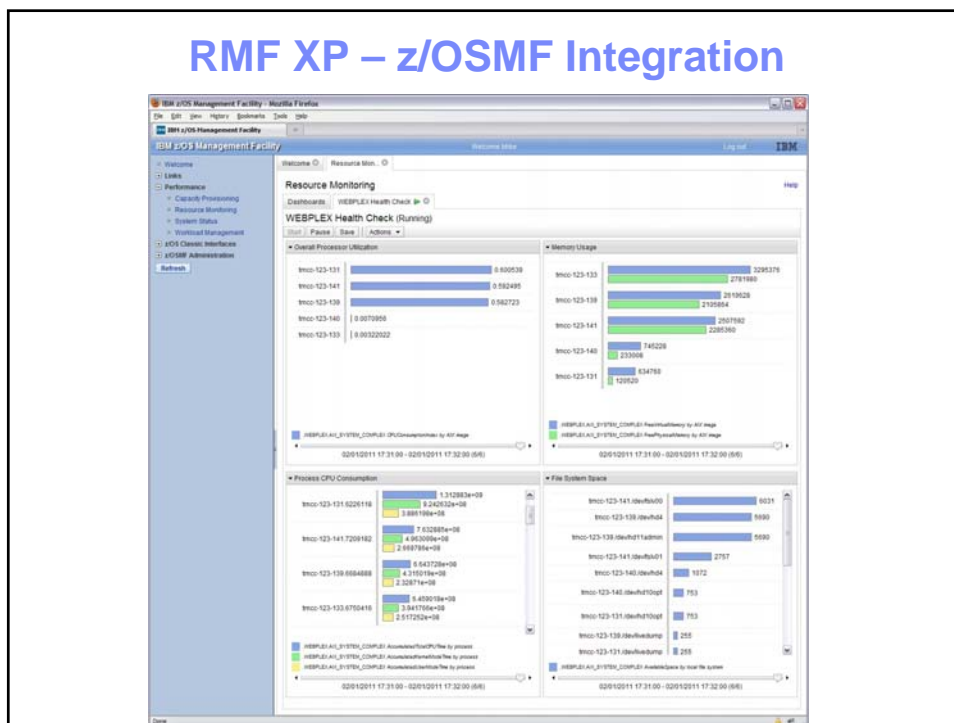
- Seamless performance monitoring solution for z/OS and distributed platforms
- z/OS as management platform for distributed environments
- Easy to setup, almost no customization needed
- Two graphical frontends
 - ▶ Instant access via web browser
 - ▶ z/OSMF with advanced capabilities
- zIIP exploitation helps to reduce costs
- Available with z/OS V1R13 RMF and z/OS V1R12 RMF (APAR OA36030)

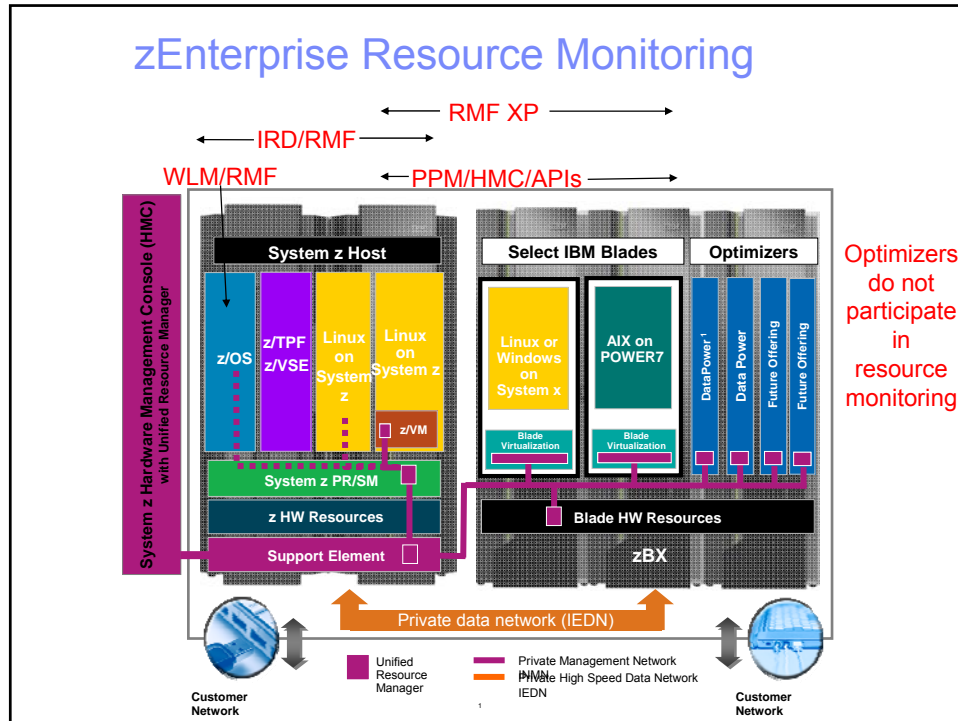
© Copyright IBM Corporation 2012



© Copyright IBM Corporation 2012

RMF XP – z/OSMF Integration

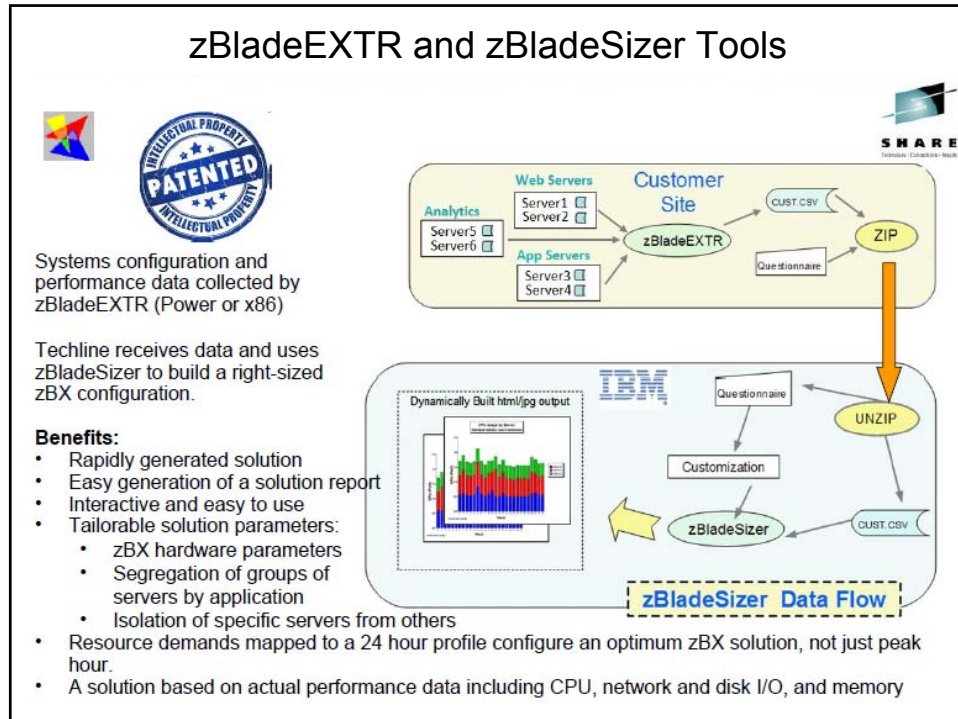




zEnterprise Ensemble Workload Selection Process

- Topology of the workload**
 - Does the workload have System z (z/OS or zLinux) components? The ensemble is designed for efficient interaction between IBM blades, optimizers, and zEnterprise.
 - What are the connection protocols?
 - The opportunity for added value in running the solution in an ensemble, such as the centralized control through the Unified Resource Manager?
- Workload environments supported?**
 - Do the workload components run in a zBX supported environment?
 - Do supporting workloads (e.g. management tools) also run in the supported environments?
 - Can the workloads run in the zBX supported virtualization environments?
 - PowerVM
 - KVM
- Workload sizing**
 - What are the resources that the solution consumes, such as:
 - Processing capacity and memory
 - Storage resources, I/O activity and bandwidth
 - Networking
 - Can the workload fit?
 - zBX Capacity Planning Tool (zBladeSizer) and Extractor (zBladeEXTR)**
 - Performance data sources: NMON data for System p and SAR for System x server running Linux

© 2013 IBM Corporation



Other Share Sessions of Interest



- **zBX Capacity Sizing Using IBM zBladeSizer and IBM zBladeEXTR**
 - Speaker: Chuck Hackett, IBM
 - Thursday, 4:30pm
- **z/OS Tuning Basics: Exploring the World of zEnterprise Hybrid, Blades and the zManager**
 - Speaker: Glenn Anderson, IBM
 - Thursday, 3pm
- **z/OS Hybrid Batch Processing on the zEnterprise**
 - Speakers: Steve Goetze and Kirk Wolf, Dovetailed Technologies
 - Tuesday, 3pm
- **Unified Resource Manager Ensemble Labs**
 - Speaker: Hiren Shah, IBM
 - Wednesday, 3pm

© 2013 IBM Corporation



Please complete the evaluation for Session # 12946



© 2013 IBM Corporation