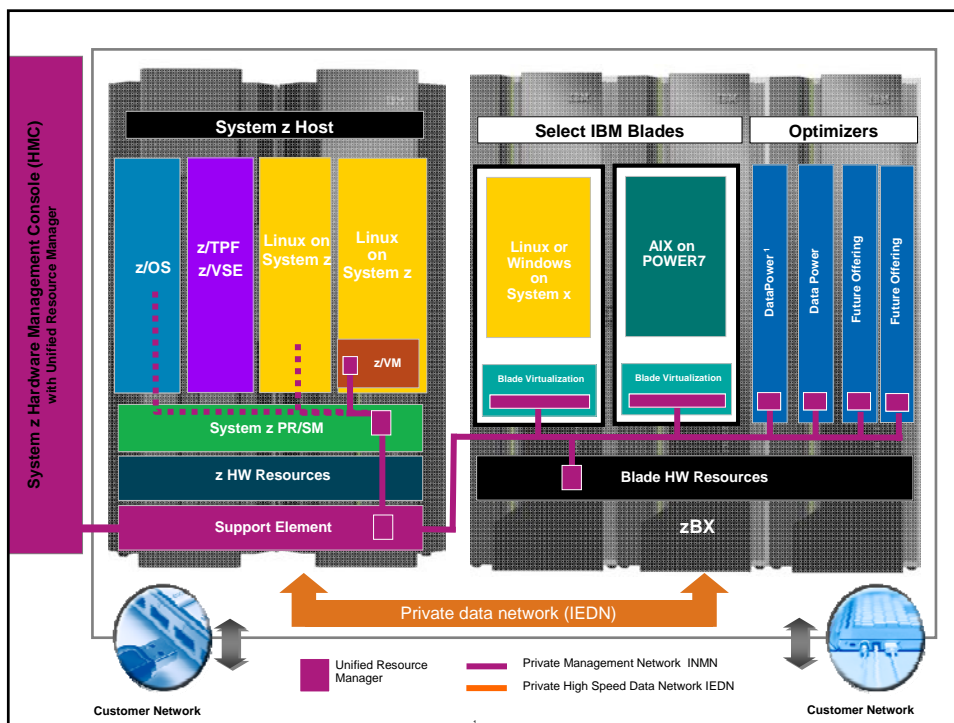SmarterComputing

IBM

Application Performance Management and Capacity Planning for IBM zEnterprise Hybrid Workloads

Glenn Anderson, IBM Lab Services and Training

Summer SHARE 2013
Session 14038

Agenda

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

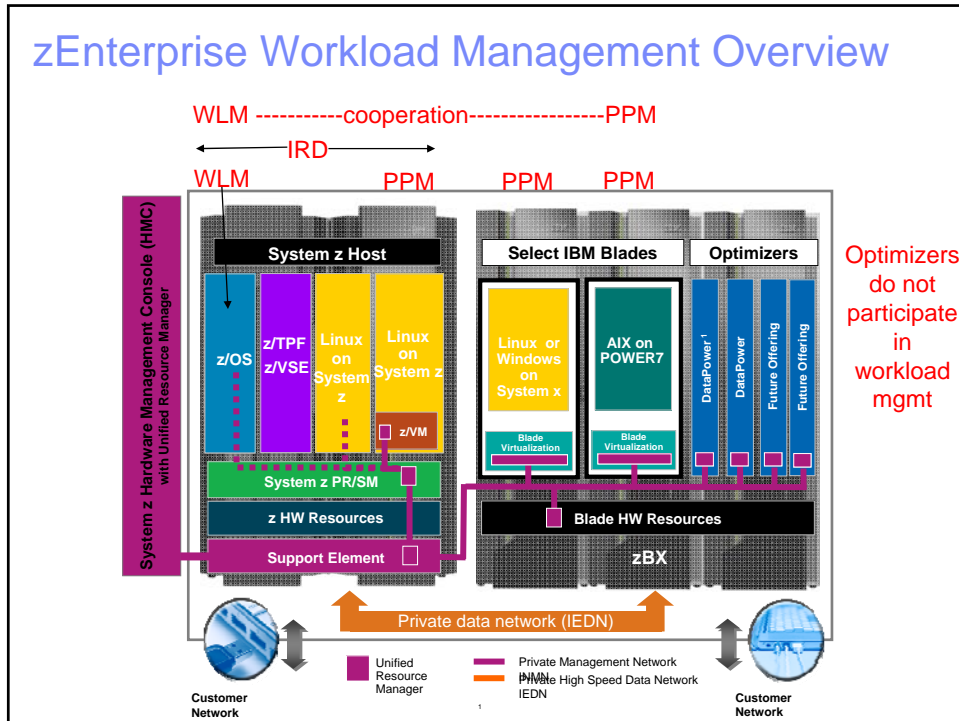- zEnterprise Resource Monitoring
  - z/OS Virtual Servers
    - RMF
  - z/VM Guests and zBX Blade Virtual Servers
    - PPM
    - RMF XP

zEnterprise Workload
Management

zEnterprise Workload Management Overview

Platform Performance Manager — IBM
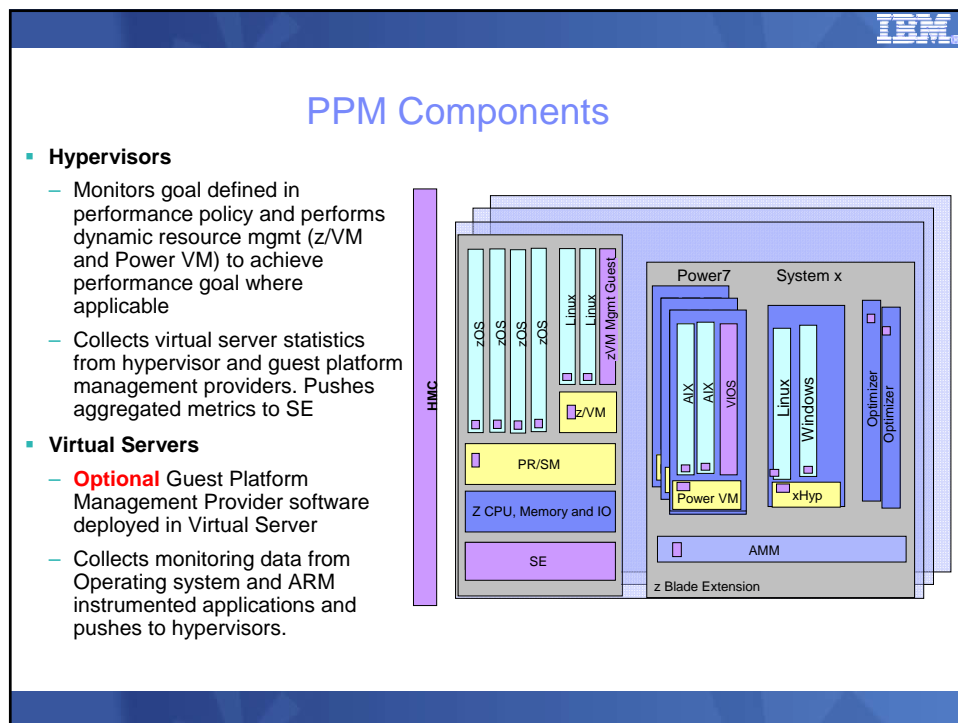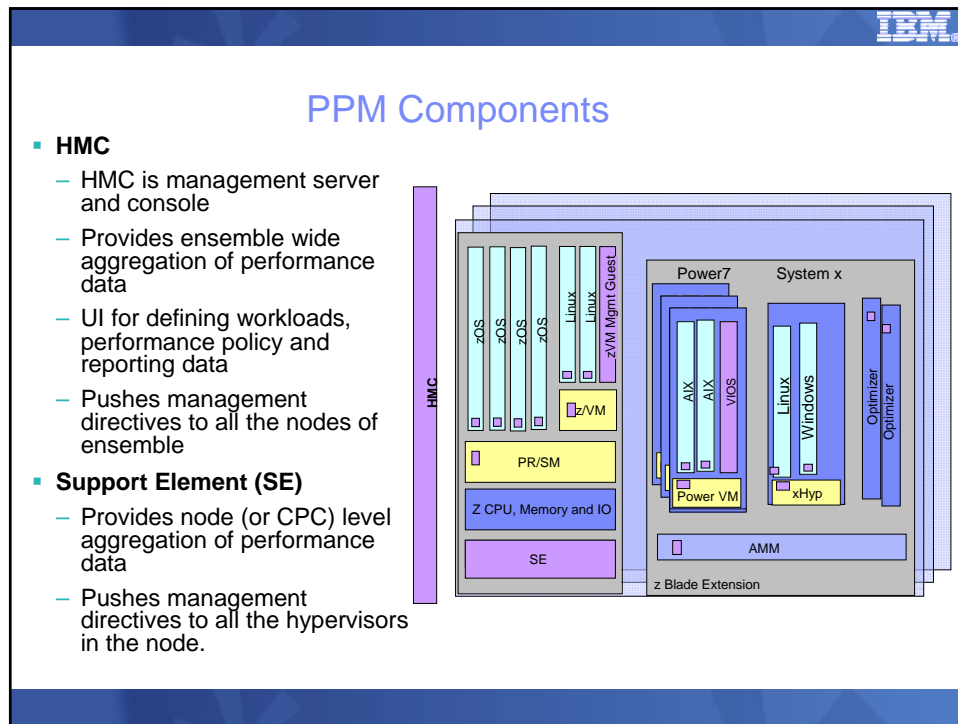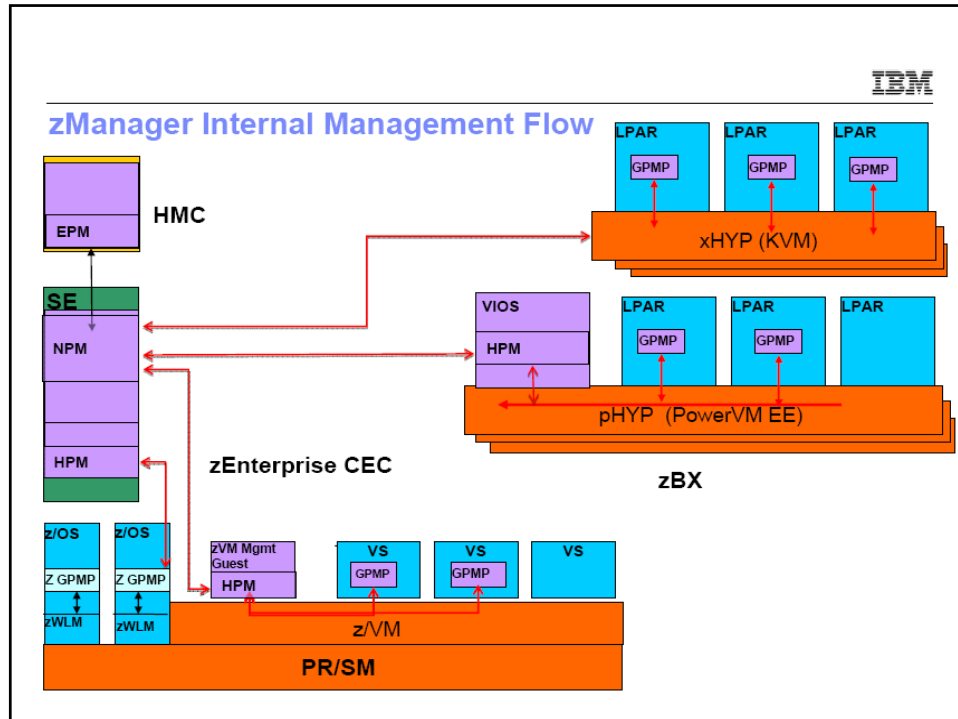
zEnterprise Platform Performance Manager

- **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

PPM Components

IBM

- **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.

PPM Components

IBM

- **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.

zManager Internal Management Flow

zManager CPU Resource Mgmt Function

- z/VM, PowerVM and System x (KVM based) Hypervisors
  - Virtual Server CPU Management provides the ability to manage CPU resources across virtual servers based on a goal-oriented performance policy.

- PR/SM Hypervisor
  - Does not make resource management adjustments based on PPM Policy. Only IRD dynamically influences the PR/SM hypervisor

PPM Dynamic Processor Management

Resource Adjustment Based on Goals

Platform Workload

- A Platform Workload is a grouping mechanism and "management view" of virtual servers supporting a business function

- Provides the context within which associated platform resources are presented, monitored, reported, and managed

- Management policies are associated to Platform Workload
  - Performance Policy

- Workload can be defined by an administrator
  - Requires 'Automate' level enabled
  - Requires appropriate role

Ensemble

Workload = Payroll

xLinux Windows zLinux

Policy

AIX AIX

z/OS

Policy

AIX zLinux

Workload = HR

Workload Performance Policy…

Workload Performance Policy

PerformancePolicyName
Importance

Service Class

Service Class Name
Performance Goal
Importance

Classification Rule
→ 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

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

IBM

Performance Policies
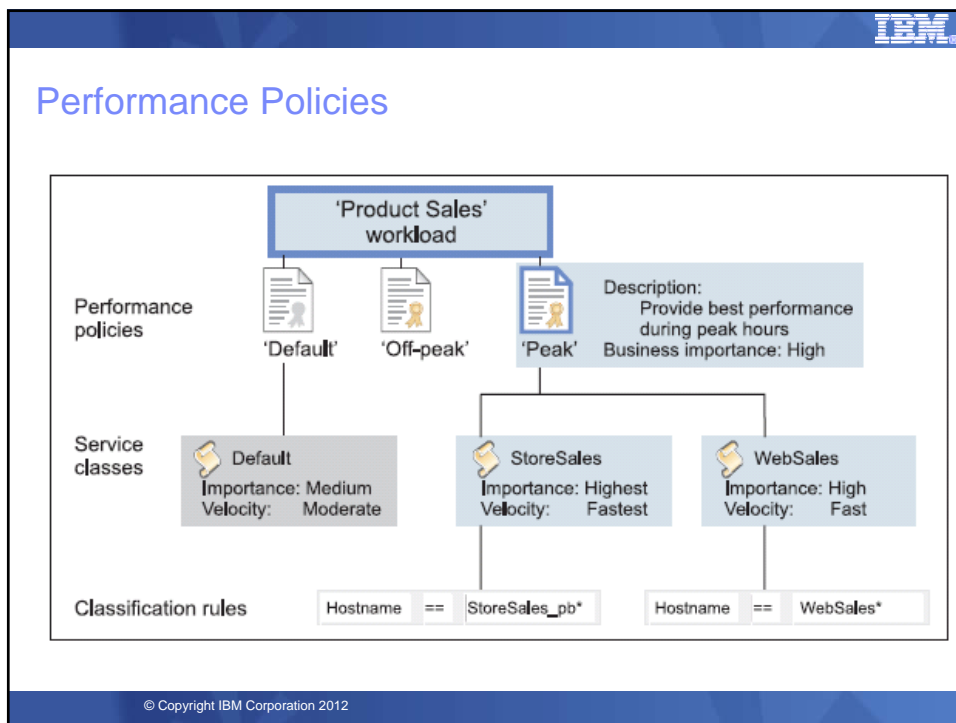
Managing Resources across Virtual Servers on P7 blade

- 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
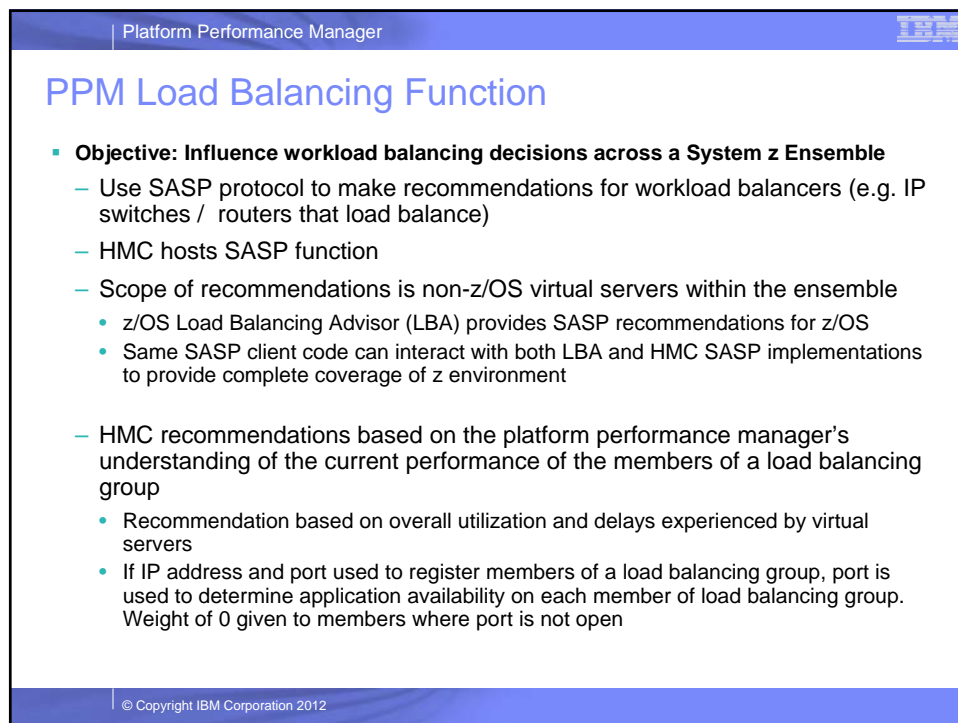 - If good trade-off based on policy, redistribute resources
  - Initially support CPU management

CPU Management causes the hypervisor to move processing units (pu) from the donor to the receiver virtual server

Platform Performance Manager

Managing Resources across Virtual Servers on System x Blade

- **Manage processor 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
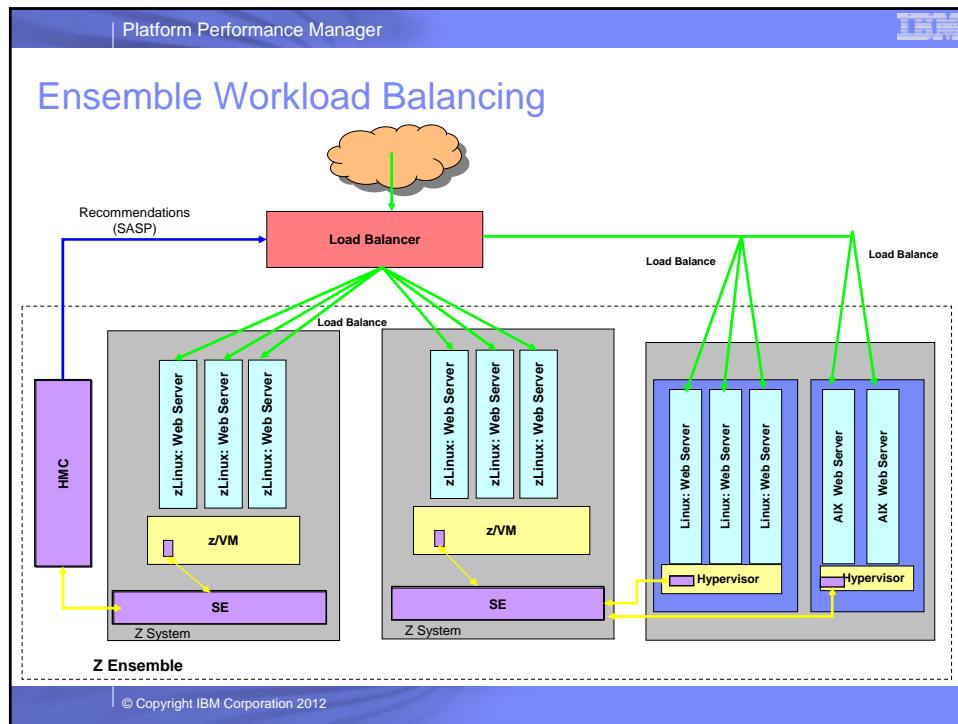 - If good trade-off based on policy, redistribute processor resources by altering **CPU shares**

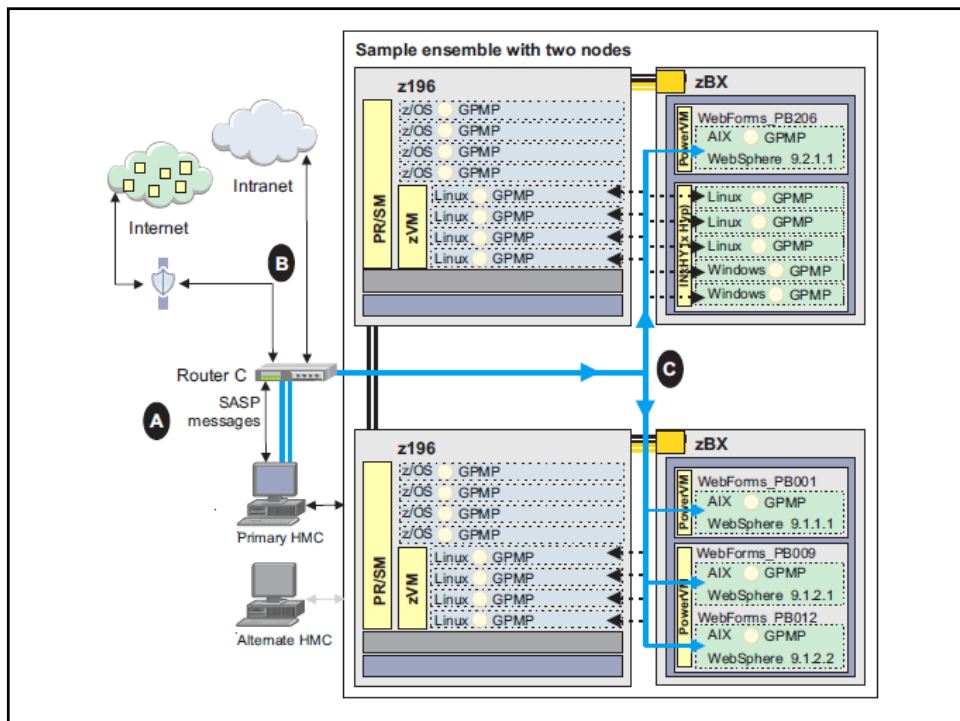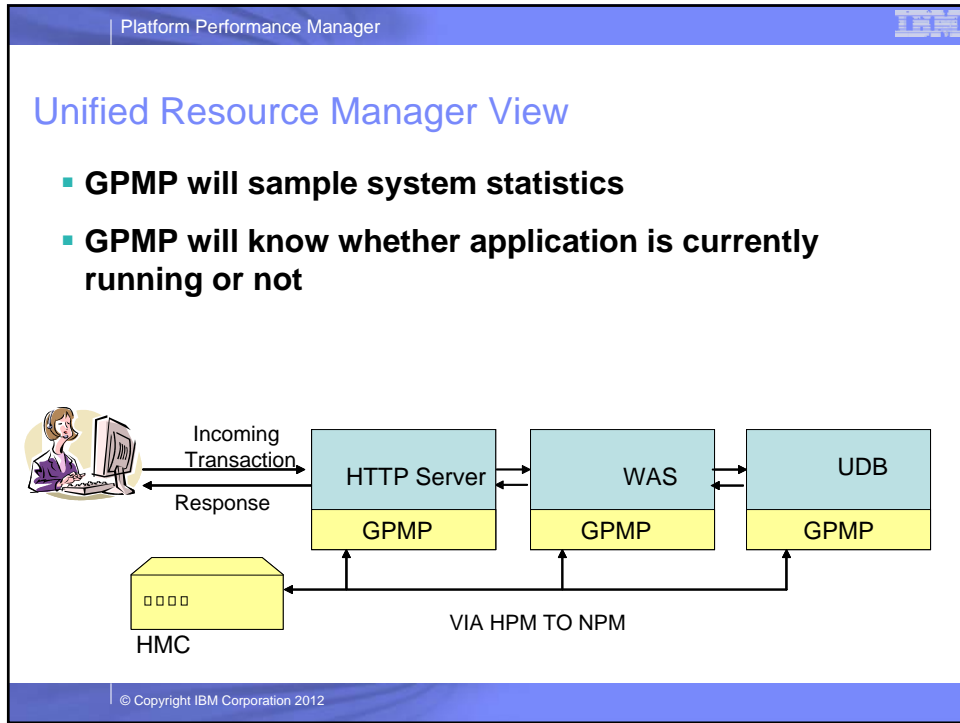© Copyright IBM Corporation 2012   © 2011 IBM Corporation

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
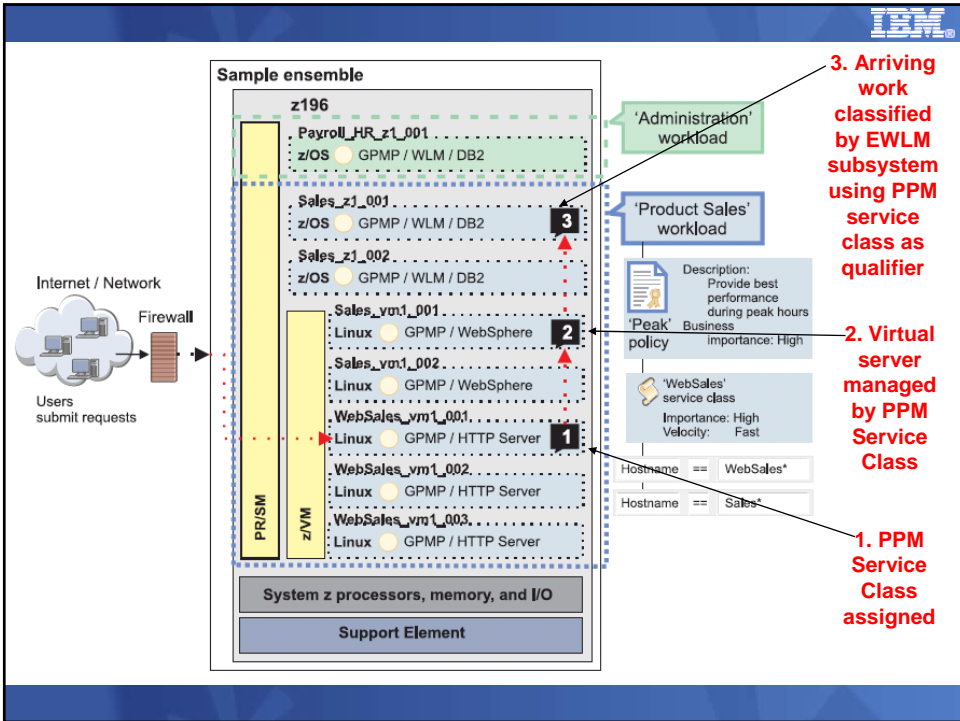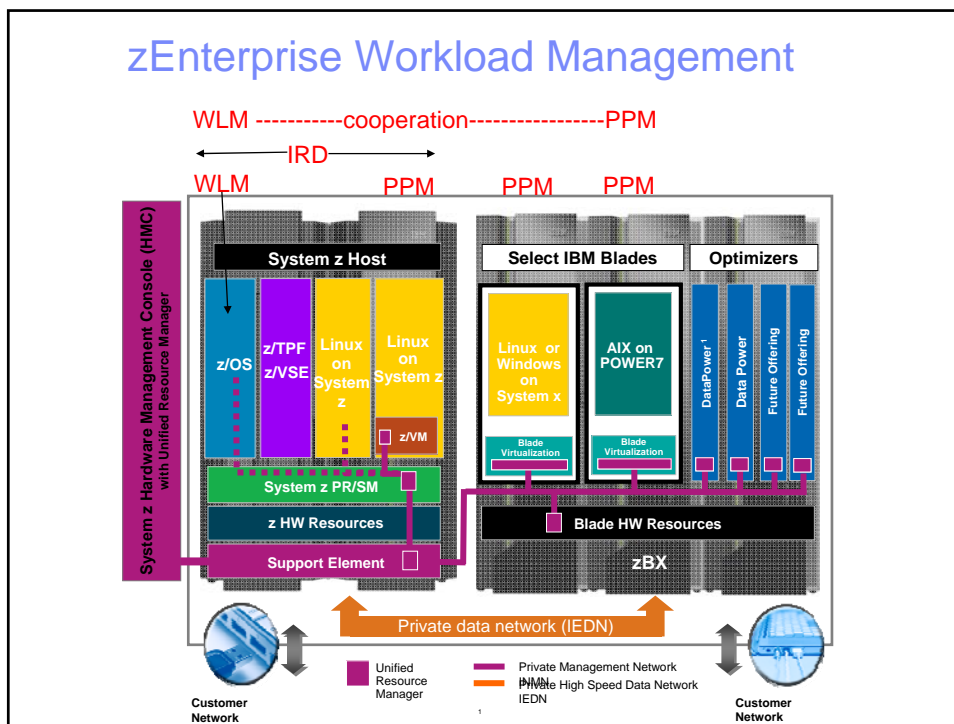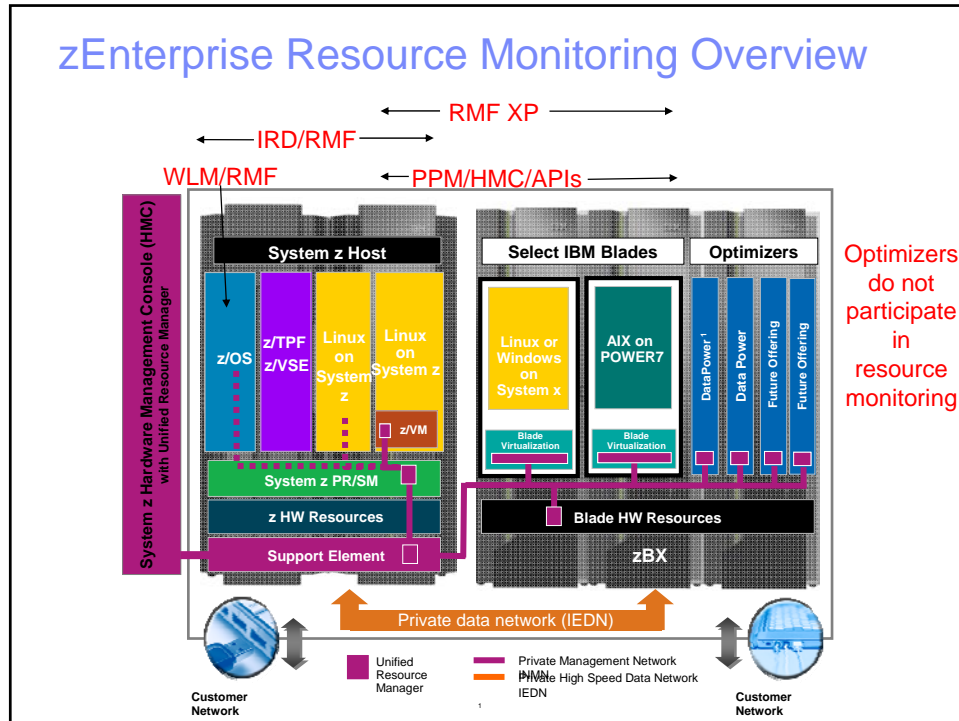
© Copyright IBM Corporation 2012

Ensemble Workload Balancing

PPM Load Balancing Function

- **Objective: Influence workload balancing decisions across a System z Ensemble**
  - Use SASP protocol to make recommendations for workload balancers (e.g. IP switches / routers that load balance)
  - HMC hosts SASP function
  - Scope of recommendations is non-z/OS virtual servers within the ensemble
    - z/OS Load Balancing Advisor (LBA) provides SASP recommendations for z/OS
    - Same SASP client code can interact with both LBA and HMC SASP implementations to provide complete coverage of z environment
  - HMC recommendations based on the platform performance manager's understanding of the current performance of the members of a load balancing group
    - Recommendation based on overall utilization and delays experienced by virtual servers
    - If IP address and port used to register members of a load balancing group, port is used to determine application availability on each member of load balancing group. Weight of 0 given to members where port is not open

Platform Performance Manager

Unified Resource Manager View

- **GPMP will sample system statistics**
- **GPMP will know whether application is currently running or not**

© Copyright IBM Corporation 2012

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.

zEnterprise Workload Management

WLM -----------cooperation-----------------PPM

IRD

WLM   PPM   PPM   PPM

System z Hardware Management Console (HMC) with Unified Resource Manager

System z Host

Select IBM Blades

Optimizers

z/OS

z/TPF z/VSE

Linux on System z

Linux on System z

z/VM

Linux or Windows on System x

AIX on POWER7

DataPower [1]

Data Power

Future Offering

Future Offering

Blade Virtualization

Blade Virtualization

System z PR/SM

z HW Resources

Blade HW Resources

Support Element

zBX

Private data network (IEDN)

Unified Resource Manager

Private Management Network INMN

Private High Speed Data Network IEDN

Customer Network

Customer Network

1

---

zEnterprise Resource Monitoring

zEnterprise Resource Monitoring Overview

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

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

---

zManager APIs for Integrated Service Management

IBM Tivoli Service Management

Integrates zEnterprise monitoring into the IBM Tivoli Monitoring Infrastructure to provides real time and historical reporting data

Provides zEnterprise resource automation policies and commands in the same standard fashion as existing resources

Discovery of zEnterprise resources for use in mapping application resources and dependencies

Monitoring
ITM 6.2.3
zEnterprise Monitoring Agent and
Enterprise Common Collector

Availability
SA z/OS V3.4
SA Application Manager V3.2.2

Discovery
TADDM (in Beta*)
and
Enterprise Common Collector

APIs

zManager zHMC

* IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

© Copyright IBM Corporation 2012 IBM Confidential   © 2013 IBM Corporation

---

Component Architecture

Enterprise Common Collector (ECC)
- Captures data from Unified Resource Manager via the zManager APIs
- Provides a ReST API to other Tivoli products for access its collected data
- Can connect to multiple zHMCs to consolidate monitoring

IBM Tivoli zEnterprise Monitoring Agent
- Displays captured data in the Tivoli Enterprise Portal
- Records data into the Tivoli Data Warehouse for reporting/trending usage
- Situations to TEP and/or EIF events to Netcool/OMNIbus/TBSM, NetView on z/OS, etc.

ECC and zEnterprise agent can be located on the same or separate servers

© Copyright IBM Corporation 2012IBM 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
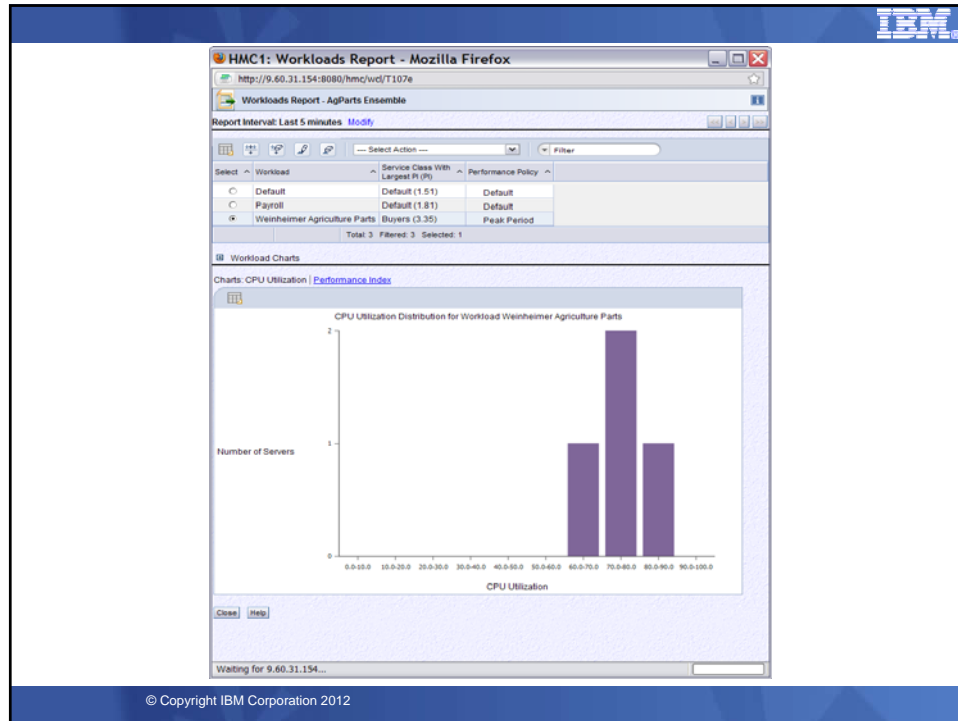
Workload Monitoring Overview…

- 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
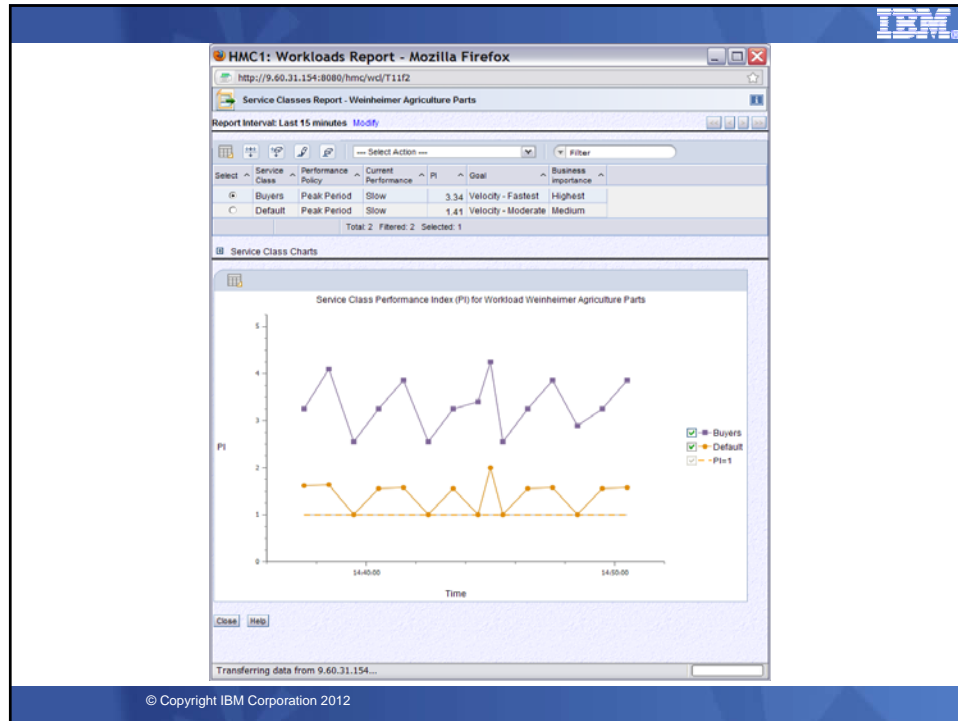
Workload Monitoring Overview…

- **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

© Copyright IBM Corporation 2012

© Copyright IBM Corporation 2012

PPM Hypervisor Report

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

8/6/2013

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.
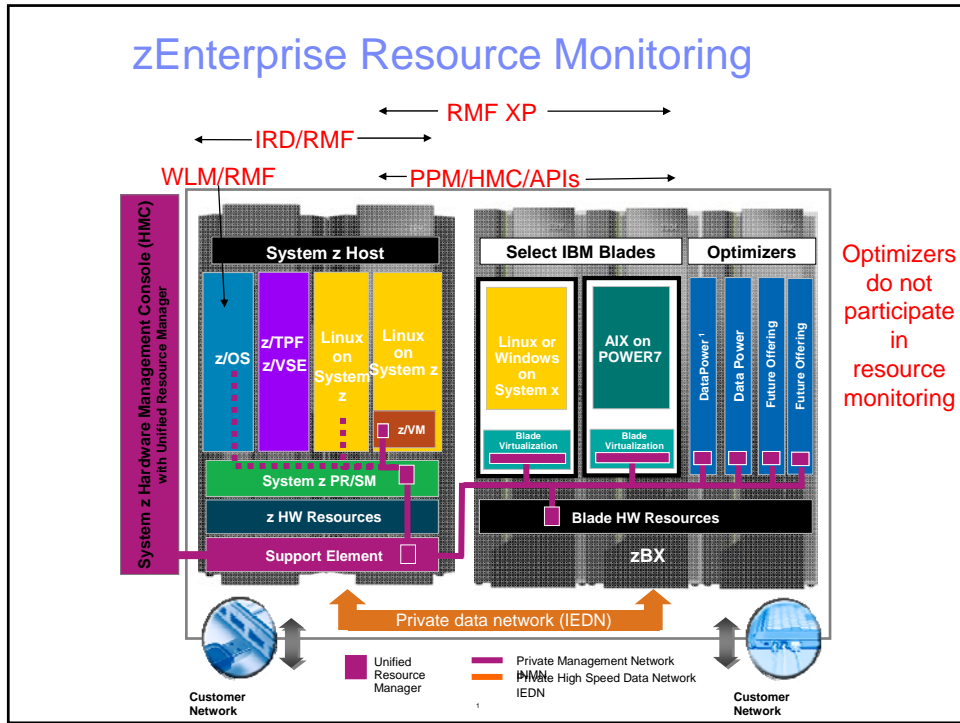
Topology report with GPMP active

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
  - Windows on System x

RMF XP – Component Overview

RMF XP – z/OSMF Integration

zEnterprise Resource Monitoring

IBM
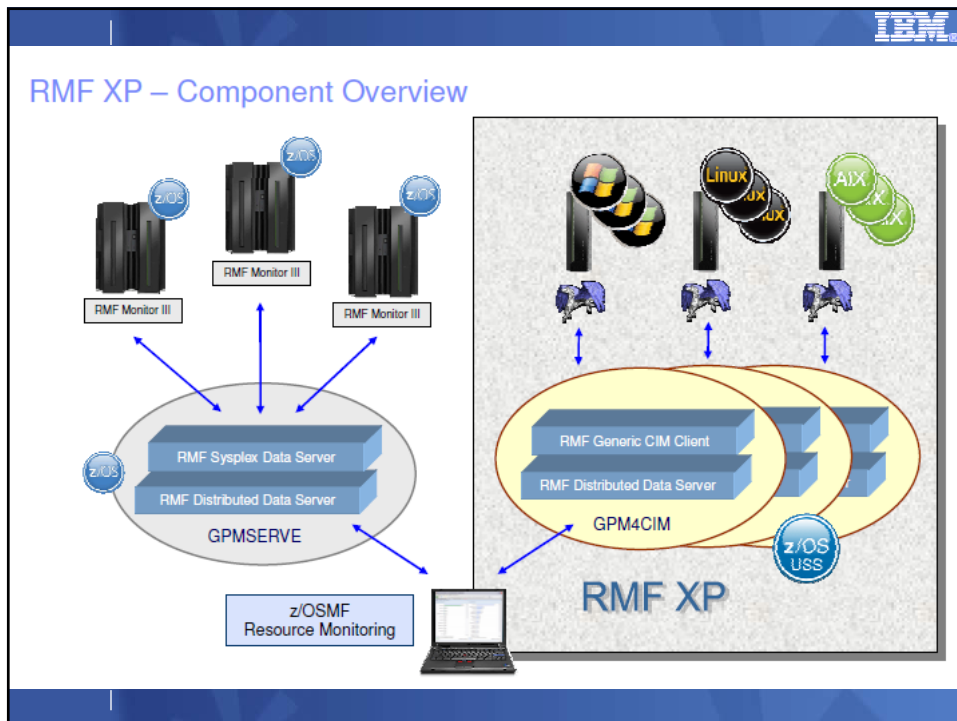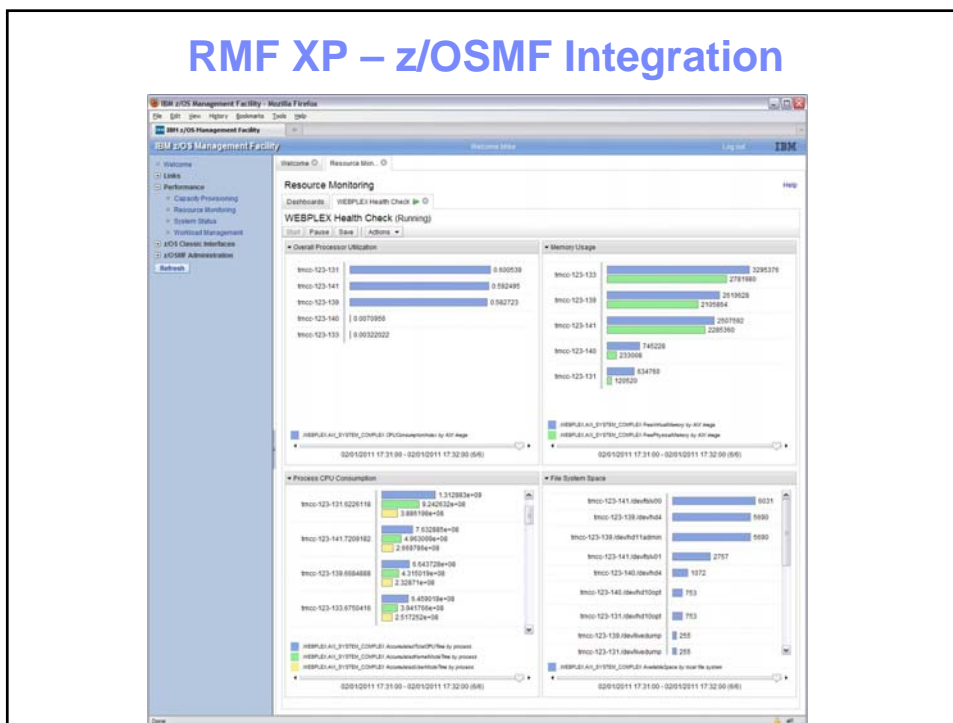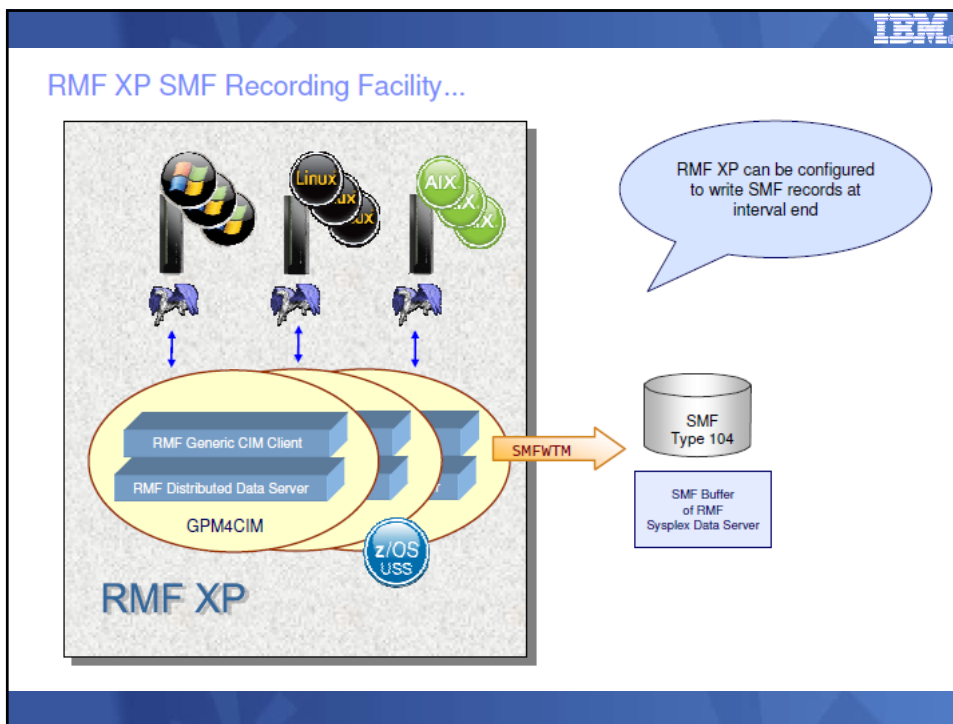
zEnterprise Ensemble Workload Selection Process

- **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)
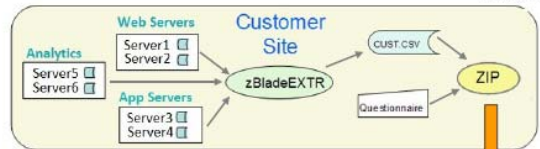   Performance data sources: NMON data for System p and SAR for System x server running Linux
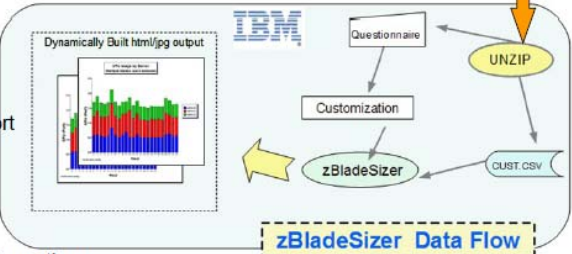
© 2013 IBM Corporation

zBladeEXTR and zBladeSizer Tools

Systems configuration and performance data collected by zBladeEXTR (Power or x86)

Techline receives data and uses zBladeSizer to build a right-sized zBX configuration.

Benefits:
- Rapidly generated solution
- Easy generation of a solution report
- Interactive and easy to use
- Tailorable solution parameters:
  - zBX hardware parameters
  - Segregation of groups of servers by application
  - Isolation of specific servers from others
- Resource demands mapped to a 24 hour profile configure an optimum zBX solution, not just peak hour.
- A solution based on actual performance data including CPU, network and disk I/O, and memory

---

SmarterComputing

IBM

Application Performance Management and Capacity Planning for IBM zEnterprise Hybrid Workloads

Glenn Anderson, IBM Lab Services and Training

Summer SHARE 2013
Session 14038

© 2013 IBM Corporation