

SHARE, August 2011, Orlando



zEnterprise eXposed! Part 1: The Intersection of WLM, RMF, and zManager Performance Management

Glenn Anderson, IBM Technical Training

Session 10002



© 2011 IBM Corporation




Agenda

- **zEnterprise Workload Management**
 - z196 / z114
 - z/OS Virtual Servers - WLM
 - IRD Clusters
 - 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


2

© 2011 IBM Corporation



zEnterprise Workload Management

zEnterprise Workload Management Overview




z/OS Workload Manager (WLM)

- A contract between the installation and the z/OS operating system
- Installation
 - Classifies work running on z/OS in distinct Service Classes
 - Defines goals that express the expectation of how work should perform
- WLM
 - Uses goal definitions to manage work across all systems of a sysplex through distribution of resources



z/OS WLM Workload Mgmt in an LPAR



© Copyright IBM Corporation 2011

Intelligent Resource Director (IRD)

What is IRD?

- Set of functions that distribute CPC resources based on business importance
- Problem areas being addressed:
 - Workloads may change over the course of a day, week, month, year...
 - Business priorities may change over course of a day, week, month, year...
 - E.g., online vs. batch, production vs. test, workload fluctuations, or periodic work
 - Single static configuration may be sub-optimal to handle different workload mixes
 - Distribute resources based on workload and service level agreements (WLM goals)
 - Reliability problems, e.g. caused by single points of failure

Consists of

- **LPAR CPU Management**
 - LPAR Weight Management
 - LPAR Vary CPU Management
- **Dynamic Channel Path Management (DCM)**
- **Channel Subsystem Priority Queuing (CSSPQ)**

Shared general purpose processors


Channel bandwidth

Channel subsystem

© Copyright IBM Corporation 2011

LPAR Clusters

- Scope of IRD management is the **LPAR Cluster (LPC)**
- LPC = "Set of LPARs on same CPC, which are part of same Sysplex"
 - For CPU management of zLinux systems, specify the sysplex name as CP management cluster name on **Customize Image Profiles** panel
 - For DCM, IOCLUSTER keyword in the IODF must match Sysplex name.
- LPARs with dedicated CPs can still join a cluster
 - But will not be enabled for WLM LPAR Weight and Vary CPU Management.
- Multi-image/Sysplex LPCs require a CF structure (except for CSSPQ)



© Copyright IBM Corporation 2011



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

10

© 2010 IBM Corporation

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 (Performance Management) architecture. On the left, a vertical purple bar represents the HMC (Hardware Management Console). To its right, a yellow box represents the Support Element (SE), which contains a 'Z CPU, Memory and IO' block. Above the SE is a 'PR/SM' block. The main part of the diagram shows a stack of operating systems and hypervisors: z/OS, z/OS, z/OS, Linux, Linux, and z/VM Mgmt Guest. Below these are Power7 and System x environments. Power7 includes AIX, AIX, VIOS, and Power VM. System x includes Linux, Linux, and KVM. An 'Optimizer' block is shown next to the System x environment. At the bottom, an 'AMM' (Advanced Memory Manager) block is shown, and a 'z Blade Extension' is indicated at the very bottom.

© 2010 IBM Corporation

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.

This diagram is identical to the one on slide 11, showing the PPM architecture with HMC, SE, and various operating systems/hypervisors.

© 2010 IBM Corporation



zManager Workload Mgmt Functions

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


13

© 2010 IBM Corporation



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



14

© 2011 IBM Corporation

IBM

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


15
© 2011 IBM Corporation

IBM


Elements of a Service Class

- **Performance Goal**
 - 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 hipervisors should be classified into a Service Class for resource monitoring purposes

16
© 2011 IBM Corporation



Performance Policies




IBM

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

19 © 2011 IBM Corporation



Managing Resources across z/VM Virtual Machines

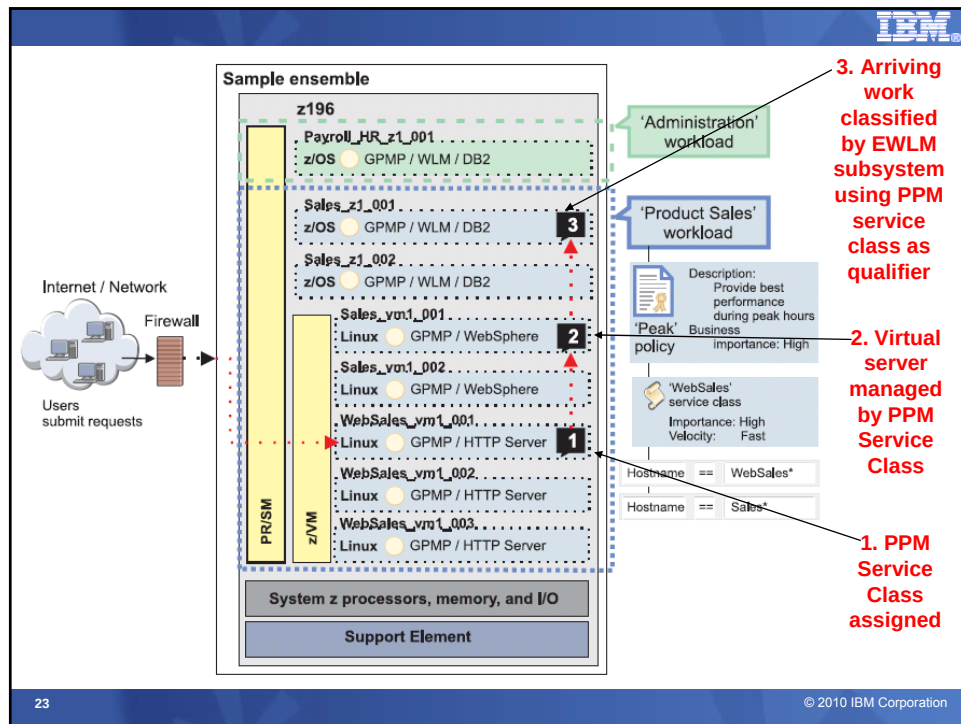
- Manage resources across z/VM virtual machines
 - Detect that a virtual machine that is part of Workload is not achieving goals
 - Determine that the virtual machine performance can be improved with additional resources
 - Project impact on all effected Workloads of moving resources to virtual machine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

21 © 2011 IBM Corporation

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.

22 © 2010 IBM Corporation



Setup for co-operative mgmt with z/OS WLM

```

Session B - [24 x 80]
File Edit View Communication Actions Window Help
Subsystem-Type Xref Notes Options Help
-----
Command ==> Modify Rules for the Subsystem Type Row 3 to 10 of 16
Scroll ==> CSR
Subsystem Type . : EWLM Fold qualifier names? Y (Y or N)
Description . . : Rules for testing PPM/GPMP RJD

Action codes: A=After C=Copy M=Move I=Insert rule
               B=Before D=Delete row R=Repeat IS=Insert Sub-rule
               More ==>

Action  Type  Qualifier  Name  Start  Service  Report
-----
1  ESC  SrvClsFo  1  1  EWLMEFA  1
2  ESC  rFastest  9  2  GPFSHTST  1
3  ESC  Highest  17  3  GPFSHTST  1
1  ESC  SrvClsFo  1  1  GPFSHTST  1
2  ESC  rFastHig  9  2  GPFSHTST  1
3  ESC  h  17  3  GPFSHTST  1
1  ESC  SrvClsFo  1  1  GPFSHTST  1
2  ESC  rModerat  9  2  GPFSHTST  1
  
```

21/040

Connected to remote server/host 9.12.41.91 using lu/pool TCP00114 and port 23

24

© 2010 IBM Corporation




zEnterprise eXposed Part 3: zManager and z/OS Workload Manager (Brad Snyder)



Tuesday, 3:00pm

25

© 2010 IBM Corporation

zEnterprise Workload Management







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

29

© 2011 IBM Corporation



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

30

© 2011 IBM Corporation



Application Programming Interfaces (APIs) to Unified Resource Manager (July 12, 2011 announcements)

IBM intends to offer Application Program Interfaces (APIs) to Unified Resource Manager. These APIs provide access to the same underlying functions that support the Unified Resource Manager user interface and can be exploited to enable Discovery, Monitoring, and Provisioning use cases.

IBM intends to extend the Tivoli® Integrated Service Management for System portfolio to take advantage of the zEnterprise ensemble monitoring and management capabilities provided by the Unified Resource Manager APIs.

31

© 2011 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


32

© 2011 IBM Corporation



33

© 2011 IBM Corporation



34


© 2011 IBM Corporation

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

35

© 2011 IBM Corporation



36


© 2011 IBM Corporation

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

37

© 2011 IBM Corporation



38

© 2011 IBM Corporation

IBM

HMC1: Workloads Report - Mozilla Firefox

http://9.60.31.154:8080/hmc/vcd/T1400

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

39 © 2011 IBM Corporation

PPM Processor Mgmt results (SC Report)

IBM

R90HMC1: Service Classes Report - Mozilla Firefox

https://9.12.16.241/hmc/vcd/T5737stableTop_e30e3c

Service Class Resource Adjustments Report - SrvClsForFastHigh

Report Interval: Starting 9/29/10 12:33:02 PM for 15 minutes (9/29/10 12:48:02 PM) [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
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.75 (0.71)	r90f1b207v5	Default	1.39 (1.40)	Sep 29, 2010 12:43:36 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.75 (0.71)	r90f1b207v8	Default	1.39 (1.40)	Sep 29, 2010 12:43:36 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.78 (0.72)	r90f1b207v2	Default	1.78 (1.80)	Sep 29, 2010 12:45:06 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.78 (0.72)	r90f1b207v5	Default	1.38 (1.40)	Sep 29, 2010 12:45:06 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.78 (0.72)	r90f1b207v8	Default	1.38 (1.40)	Sep 29, 2010 12:45:06 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.78 (0.72)	r90f1b207v2	Default	1.78 (1.80)	Sep 29, 2010 12:47:06 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.78 (0.72)	r90f1b207v5	Default	1.38 (1.40)	Sep 29, 2010 12:47:06 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	0.78 (0.72)	r90f1b207v8	Default	1.38 (1.40)	Sep 29, 2010 12:47:06 PM
			Total 31	Filtered: 31			

Failed Adjustments

Receiver Virtual Servers	Receiver Workload	Receiver Service Class	Failure Reason	Time
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	Unknown	Sep 29, 2010 12:38:51 PM
r90f1b207v6	WkldForFastHigh	SrvClsForFastHigh	Unknown	Sep 29, 2010 12:47:06 PM
			Total 2	Filtered: 2

Close Help

Done

40 © 2010 IBM Corporation

PPM Hypervisor Report

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

Failed Adjustments

Done

41 © 2010 IBM Corporation


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.
- With instrumented middleware support, GPMP provides metrics that allows detailed transaction topology as transaction hops through heterogeneous platforms in zEnterprise

42 © 2010 IBM Corporation

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.



43

© 2010 IBM Corporation

Topology report with GPMP active



44

© 2010 IBM Corporation



zEnterprise eXposed Part 2: Experiences with the zManager Guest Platform Management Providers (Mary Astley)

Tuesday, 1:30pm

45

© 2010 IBM Corporation



Cross Platform Performance Monitoring with RMF XP

- The Common Information Model (aka CIM) instrumentation is available for almost all operating systems of 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

46


© 2010 IBM Corporation

IBM

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)

47
07/26/11
© 2011 IBM Corporation





Invocation

- Started Task: SYS1.PROCLIB(GPM4CIM)
- Runs in USS Environment via BPXBATCH
- Multiple instances can run in parallel: one STC per platform
 - ▶ S GPM4CIM.GPM4A,OS=A
 - ▶ S GPM4CIM.GPM4X,OS=X
 - ▶ S GPM4CIM.GPM4Z,OS=Z

```
//GPM4CIM PROC OS=X
//STEP1 EXEC PGM=BPXBATCH,TIME=NOLIMIT,REGION=0M,
//      PARM='PGM /usr/lpp/gpm/bin/gpm4cim cfg=/etc/gpm/gpm4OS..cfg'
//STDENV DD PATH='/etc/gpm/gpm4cim.env'
//STDOUT DD PATH='/var/gpm/logs/gpm4cim&OS..out',
//      PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
//      PATHMODE=(SIRUSR,SIWUSR,SIRGRP),
//STDERR DD PATH='/var/gpm/logs/gpm4cim&OS..trc',
//      PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
//      PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//      PEND
```


49

07/26/11

© 2011 IBM Corporation



Resource Models



50

07/26/11

© 2011 IBM Corporation

RMF XP – Resource Tree

The screenshots show the RMF Performance Data Portal interface with the following data:

RMF Performance Data Portal - Overview
Welcome, you are connected to: .WEBPLEX.AIX_SYST

Children of: .WEBPLEX.AIX_SYSTEM_COMPLEX

Icon	Resource	Metrics	Attributes	Res-Type
	.WEBPLEX.AIX_SYSTEM_COMPLEX	Metrics	N/A	AIX_RANGE
	RMF DDS-Genr GPACM...	Metrics	N/A	AIX_RANGE

Children of: tmcc-123-141.*.ALL_NETWORK_PORTS

Icon	Resource	Metrics	Attributes	Res-Type
	tmcc-123-141.*.ALL_NETWORK_PORTS	Metrics	N/A	ALL_NETWORK_PORTS
	tmcc-123-141.*.ALL_NETWORK_PORTS	Metrics	N/A	ALL_NETWORK_PORTS
	tmcc-123-141.*.ALL_NETWORK_PORTS	Metrics	N/A	ALL_NETWORK_PORTS

Children of: .tmcc-123-141.AIX_IMAGE

Icon	Resource	Metrics	Attributes	Res-Type
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ACTIVE_MEMORY_EXPANSION
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ACTIVE_MEMORY_SHARING
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_DISKS
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_LOGICAL_PROCESSORS
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_NETWORK_PORTS
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_LOCAL_FILE_SYSTEMS
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_PROCESSES
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_SHARED_ETHERNET_ADAPTERS
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	ALL_VIRTUAL_TARGET_DEVICES
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	PARTITION
	.tmcc-123-141.AIX_IMAGE	Metrics	N/A	MEMORY

51

07/26/11

© 2011 IBM Corporation

RMF XP – z/OSMF Integration

The screenshot shows the IBM z/OSMF Management Facility interface with the following data:

Resource Monitoring
WEBPLEX Health Check (Running)

Overall Processor Utilization

Resource	Value
tmcc-123-131	0.000538
tmcc-123-141	0.002490
tmcc-123-138	0.002723
tmcc-123-140	0.007086
tmcc-123-132	0.0032222

Memory Usage

Resource	Value
tmcc-123-133	320376
tmcc-123-138	219180
tmcc-123-138	219180
tmcc-123-141	219180
tmcc-123-140	219180
tmcc-123-131	219180

Process CPU Consumption


Resource	Value
tmcc-123-131	1.312861e+08
tmcc-123-141	9.24252e+08
tmcc-123-138	1.00155e+08
tmcc-123-140	1.00155e+08
tmcc-123-132	1.00155e+08
tmcc-123-138	1.00155e+08
tmcc-123-140	1.00155e+08
tmcc-123-132	1.00155e+08

File System Space

Resource	Value
tmcc-123-141	9031
tmcc-123-138	9031
tmcc-123-141	9031
tmcc-123-140	9031
tmcc-123-131	9031
tmcc-123-138	9031
tmcc-123-140	9031
tmcc-123-131	9031

52

© 2011 IBM Corporation



z/OS Tuning Courses from IBM Training

- **Basic z/OS Tuning Using the Workload Manager (WLM)**
 - ES545
 - 4.5 Days, Hands-on Lab Exercises
- **Advanced z/OS Performance: WLM, Sysplex, Unix Services, Web**
 - ES851
 - 4.5 Days
- ibm.com/training