zEnterprise Platform Performance Manager: Overview and Deepdive

Hiren Shah (hiren@us.ibm.com)
IBM
August 10th, 2011
9709

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

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 therefrom.
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.
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
- Guest Platform Management Provider
- WLM and PPM relationship

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
    - System X86 Virtual images
    - IBM Smart Analytics Optimizer for DB2
    - IBM DataPower Appliance
  - A zEnterprise Node can be a member of at most one Ensemble
- zEnterprise Unified Resource Manager (zManager)
  - 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

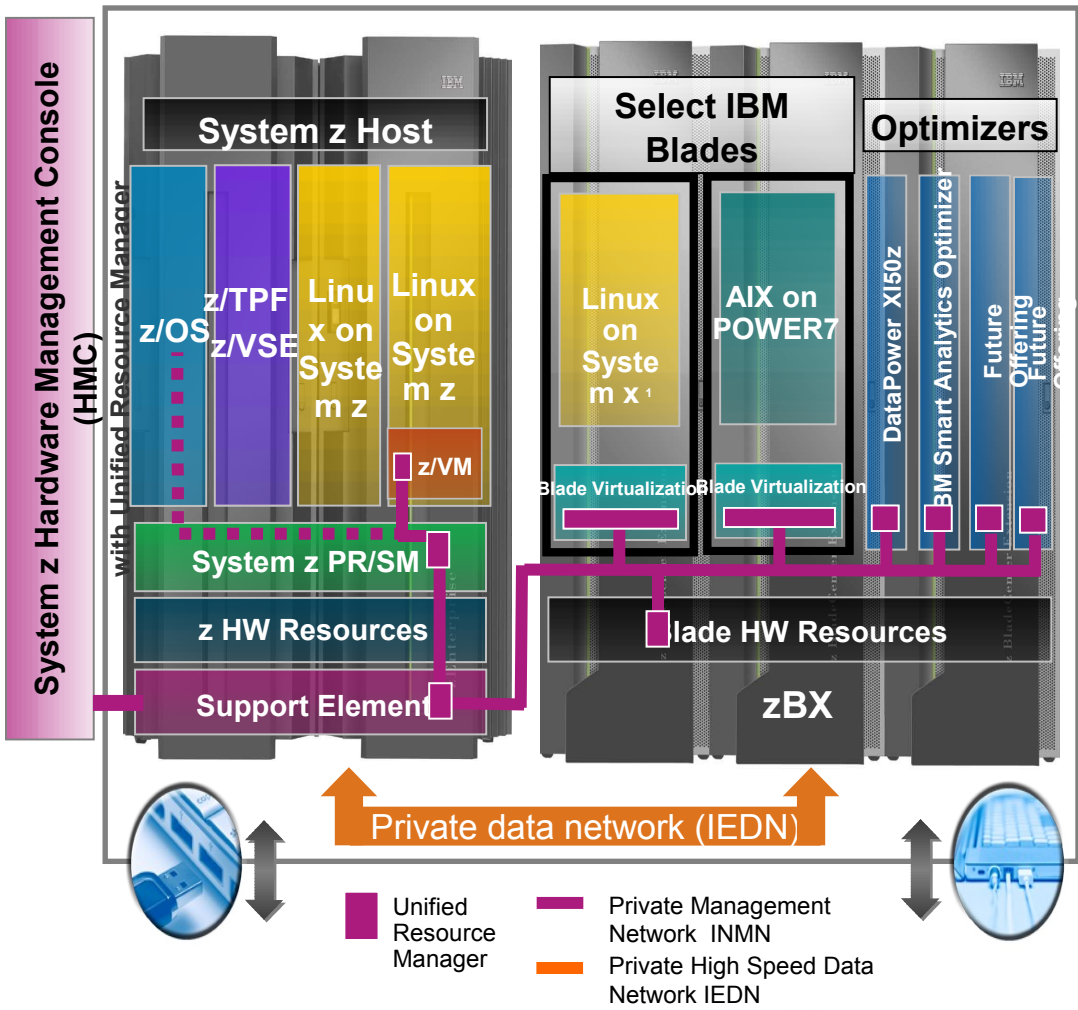

zEnterprise Platform Performance Manager

- 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
  - 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

Platform Performance Manager Structure

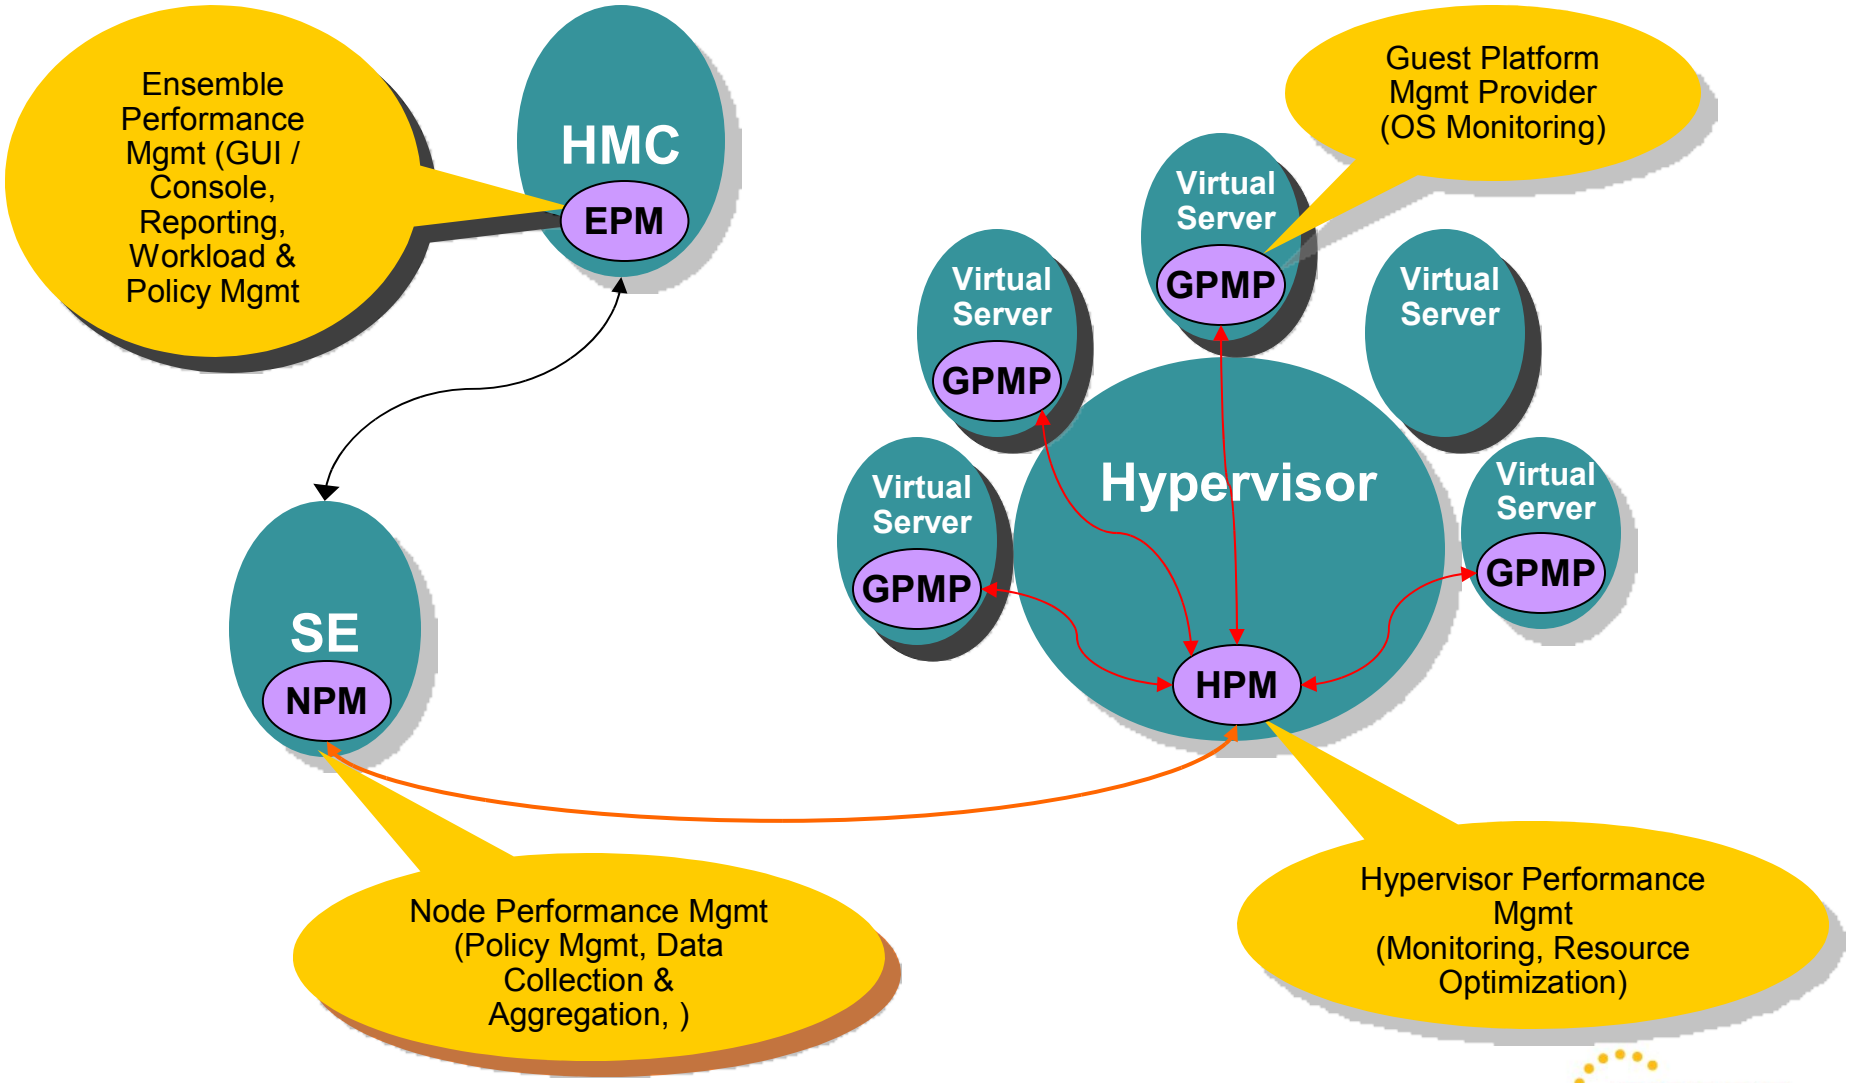

Platform Workload

- **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

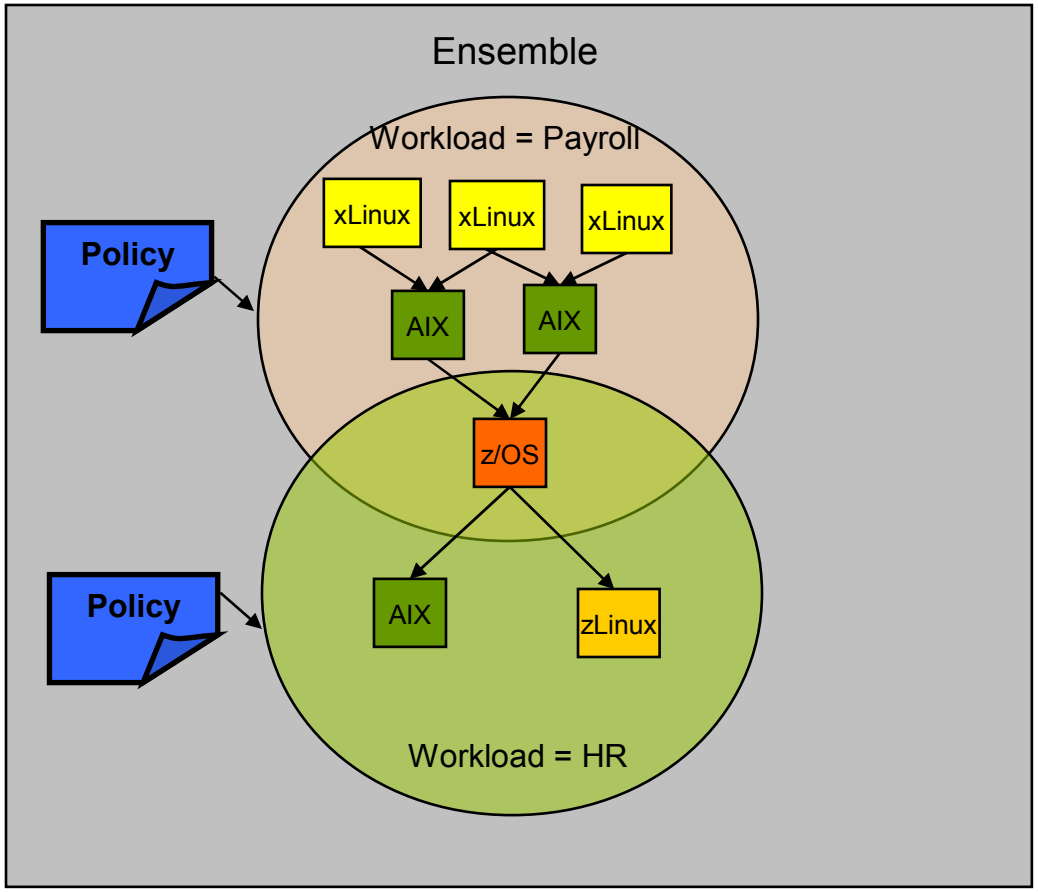

Workload Performance Policy

- 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…

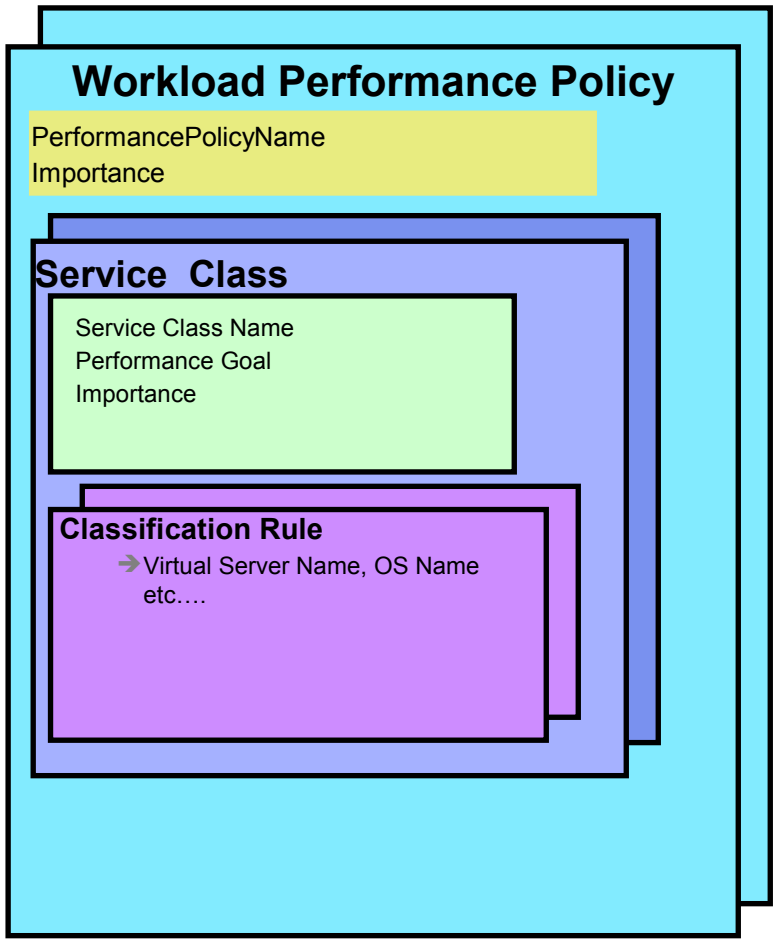

- Policy structure:
  - Policy contains a set of service classes
  - Classification rules map each virtual server within the workload to a service class
  - A service class assigns a performance goal and importance
- HMC as console for policy creation and editing
  - Wizard for policy creation
  - Repository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

Define a new Workload

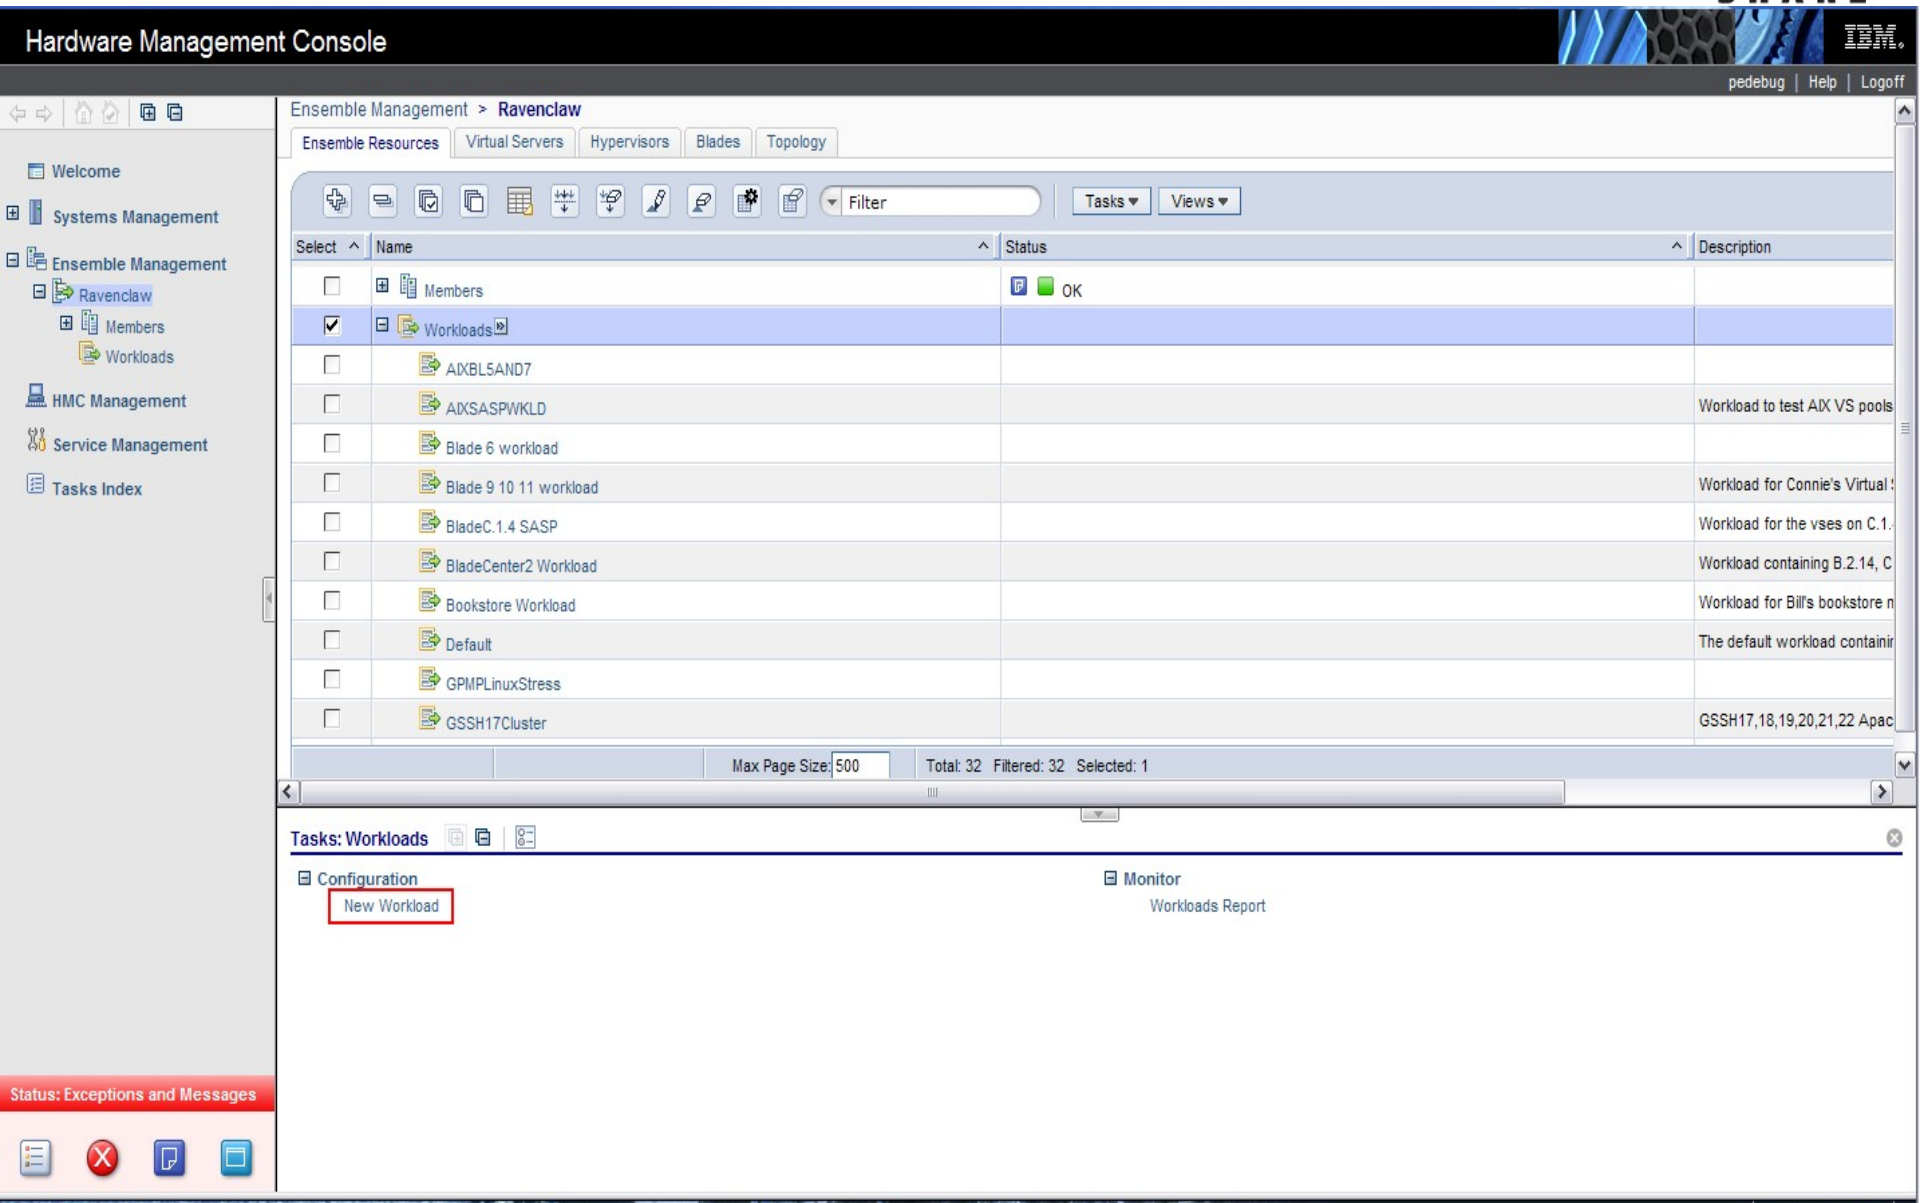

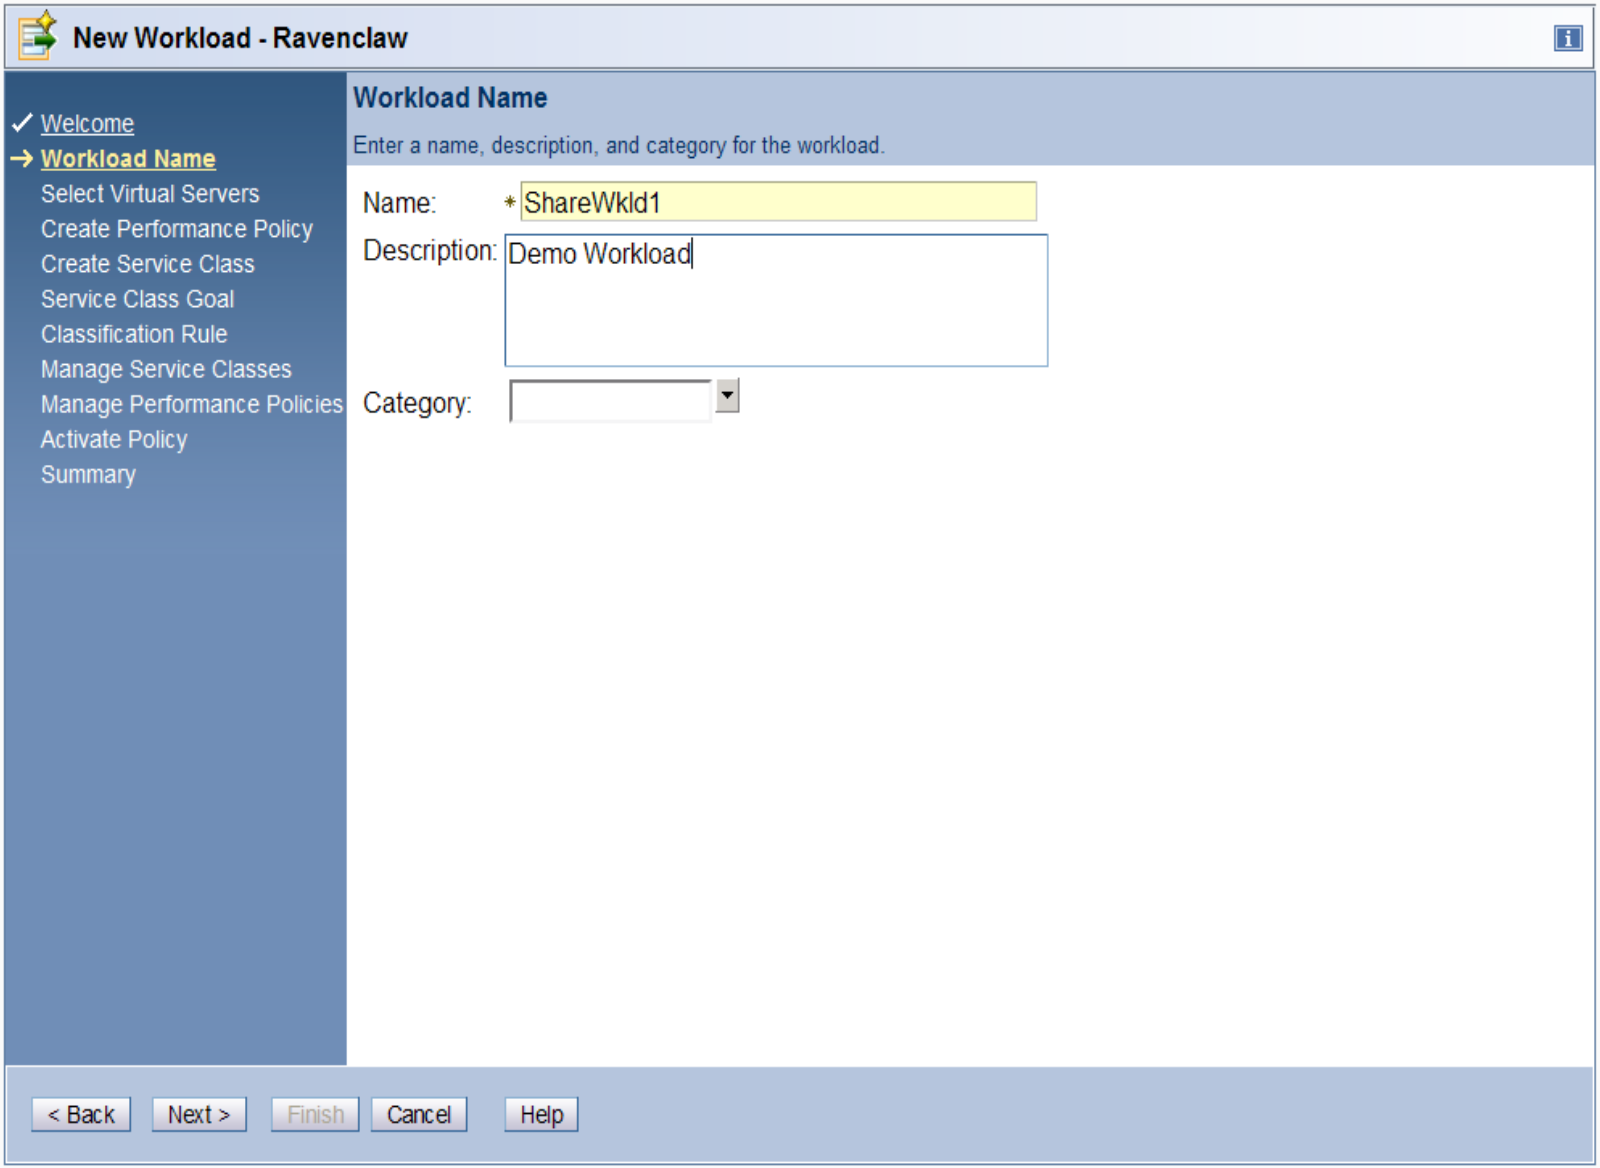

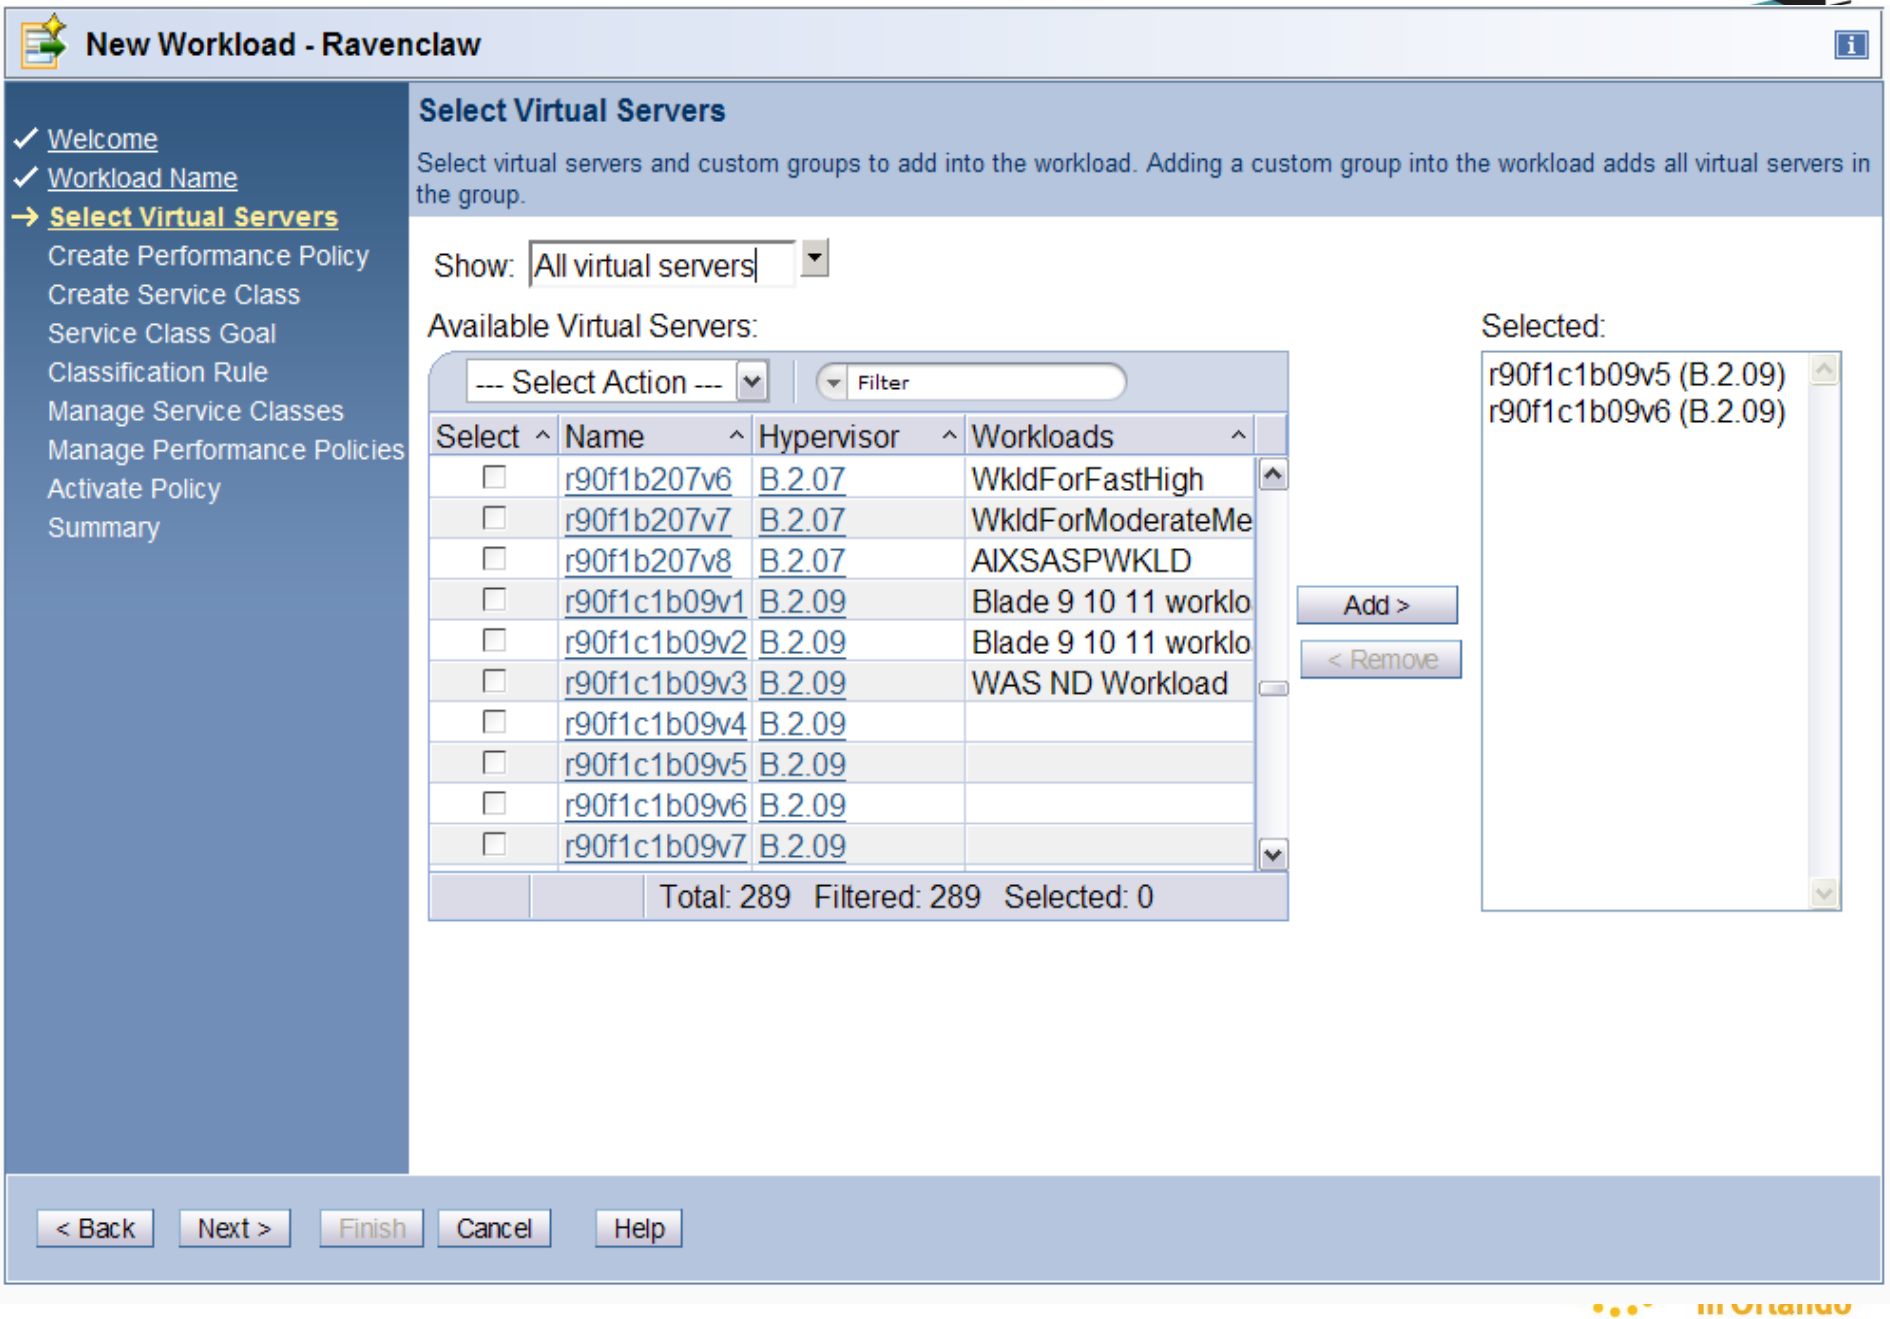
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Select Virtual Servers
Select virtual servers and custom groups to add into the workload. Adding a custom group into the workload adds all virtual servers in the group.
Show: All virtual servers
Available Virtual Servers:
--- Select Action ---
Filter
Select
Name
Hypervisor
Workloads
r90f1b207v6 B.2.07 WkldForFastHigh
r90f1b207v7 B.2.07 WkldForModerateMe
r90f1b207v8 B.2.07 AIXSASPWKLD
r90f1c1b09v1 B.2.09 Blade 9 10 11 worklo
r90f1c1b09v2 B.2.09 Blade 9 10 11 worklo
r90f1c1b09v3 B.2.09 WAS ND Workload
r90f1c1b09v4 B.2.09
r90f1c1b09v5 B.2.09
r90f1c1b09v6 B.2.09
r90f1c1b09v7 B.2.09
Total: 289 Filtered: 289 Selected: 0
Add >
< Remove
Selected:
r90f1c1b09v5 (B.2.09)
r90f1c1b09v6 (B.2.09)
< Back
Next >
Finish
Cancel
Help

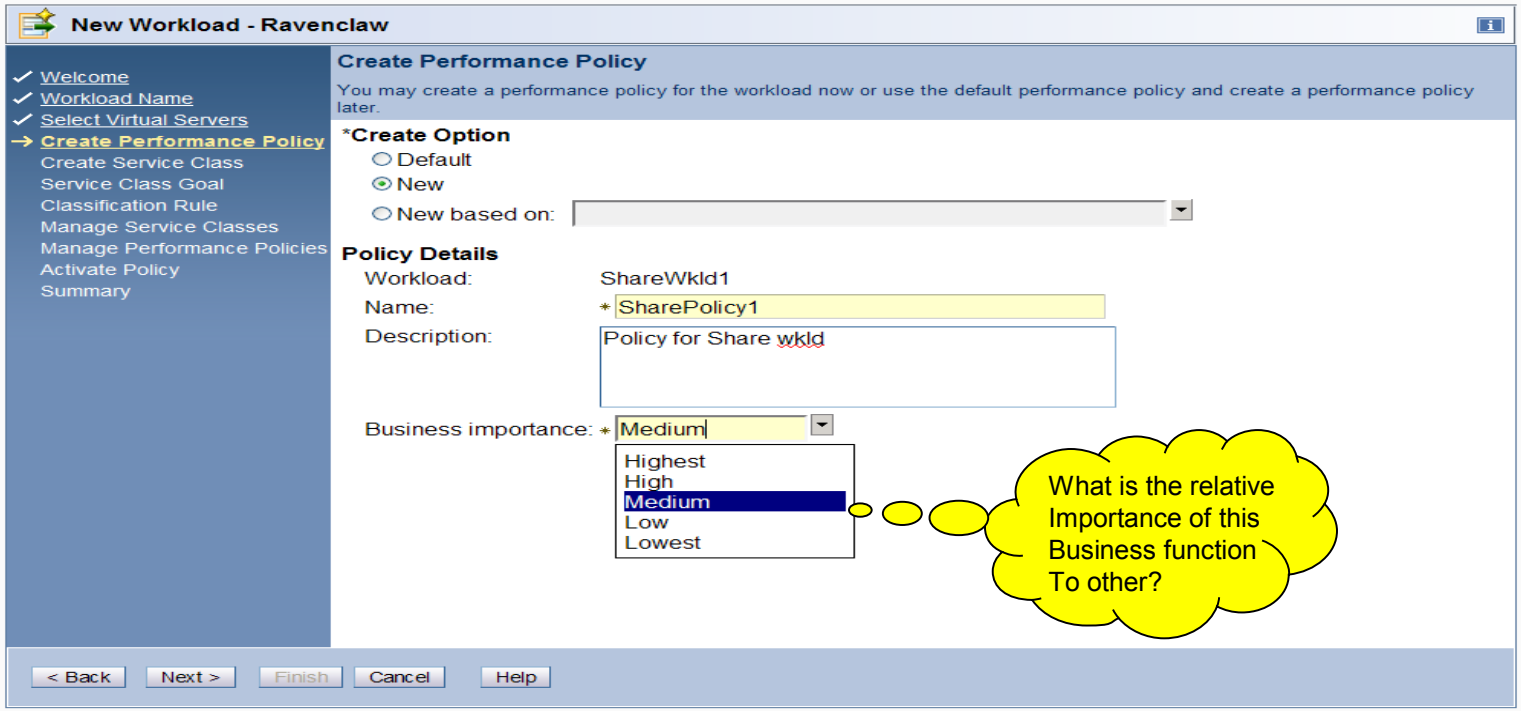
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Create Performance Policy
You may create a performance policy for the workload now or use the default performance policy and create a performance policy later.
*Create Option
Default
New
New based on:
Policy Details
Workload: ShareWkld1
Name: * SharePolicy1
Description: Policy for Share wkld
Business importance: * Medium
Highest
High
Medium
Low
Lowest
What is the relative Importance of this Business function To other?
< Back
Next >
Finish
Cancel
Help

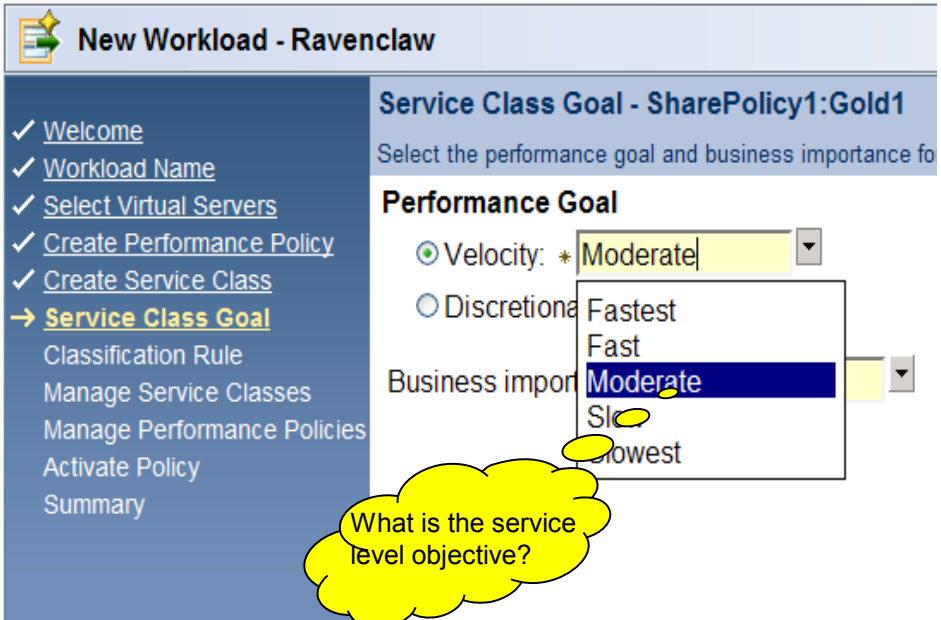
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Service Class Goal - SharePolicy1:Gold1
Select the performance goal and business importance fo
Performance Goal
Velocity: * Moderate
Discretiona
Business impor
Fastest
Fast
Moderate
Slo
owest
What is the service level objective?

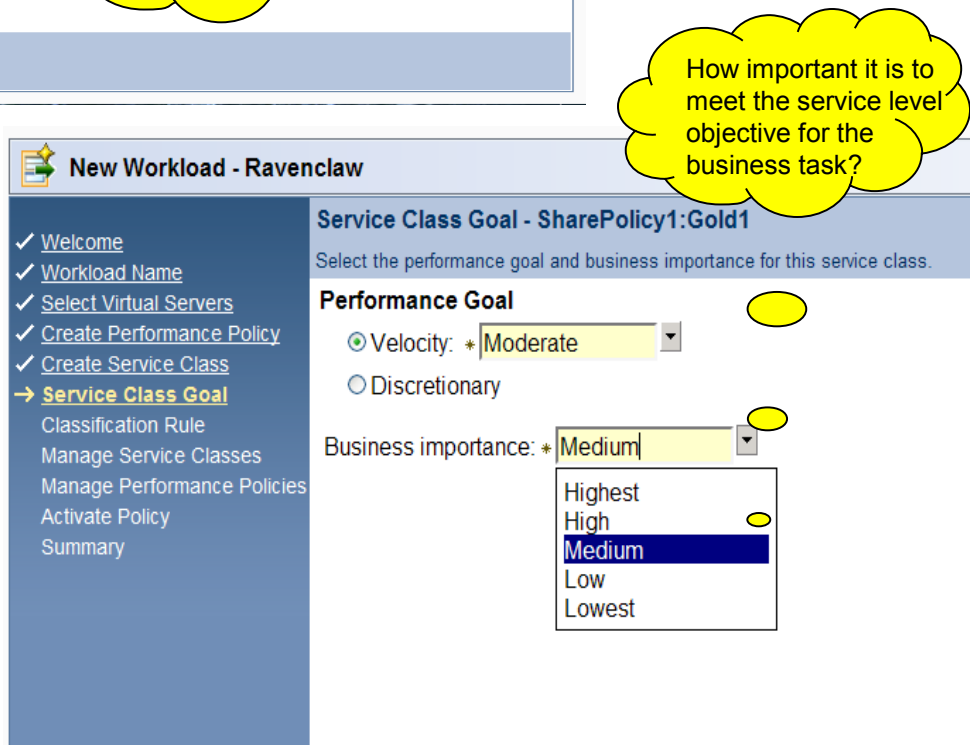
How important it is to meet the service level objective for the business task?
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Service Class Goal - SharePolicy1:Gold1
Select the performance goal and business importance for this service class.
Performance Goal
Velocity: * Moderate
Discretionary
Business importance: * Medium
Highest
High
Medium
Low
Lowest

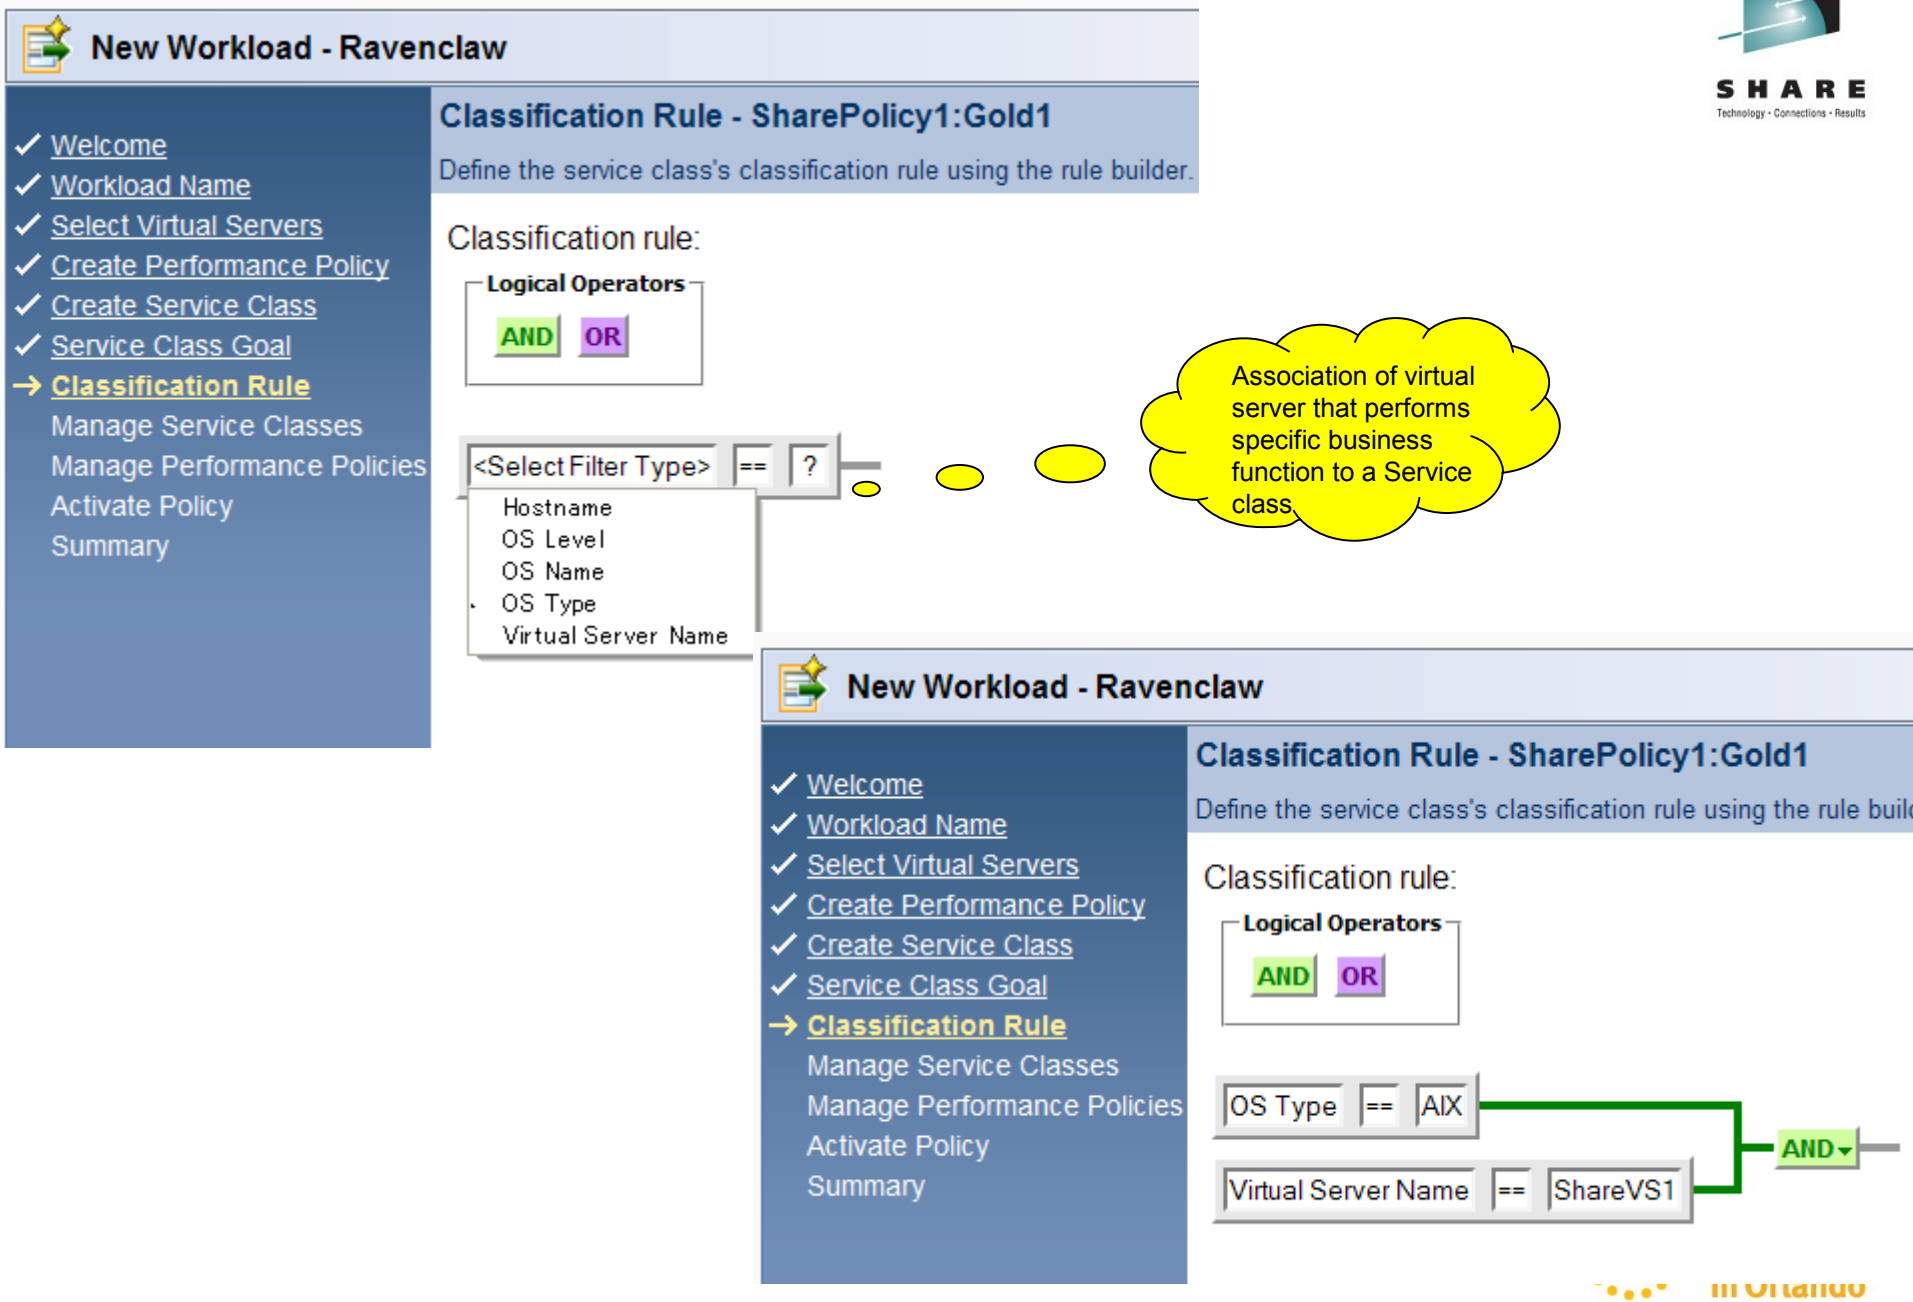
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Classification Rule - SharePolicy1:Gold1
Define the service class's classification rule using the rule builder.
Classification rule:
Logical Operators
AND
OR
<Select Filter Type>
==
?
Hostname
OS Level
OS Name
OS Type
Virtual Server Name
Association of virtual server that performs specific business function to a Service class
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Classification Rule - SharePolicy1:Gold1
Define the service class's classification rule using the rule build
Classification rule:
Logical Operators
AND
OR
OS Type
==
AIX
Virtual Server Name
==
ShareVS1
AND

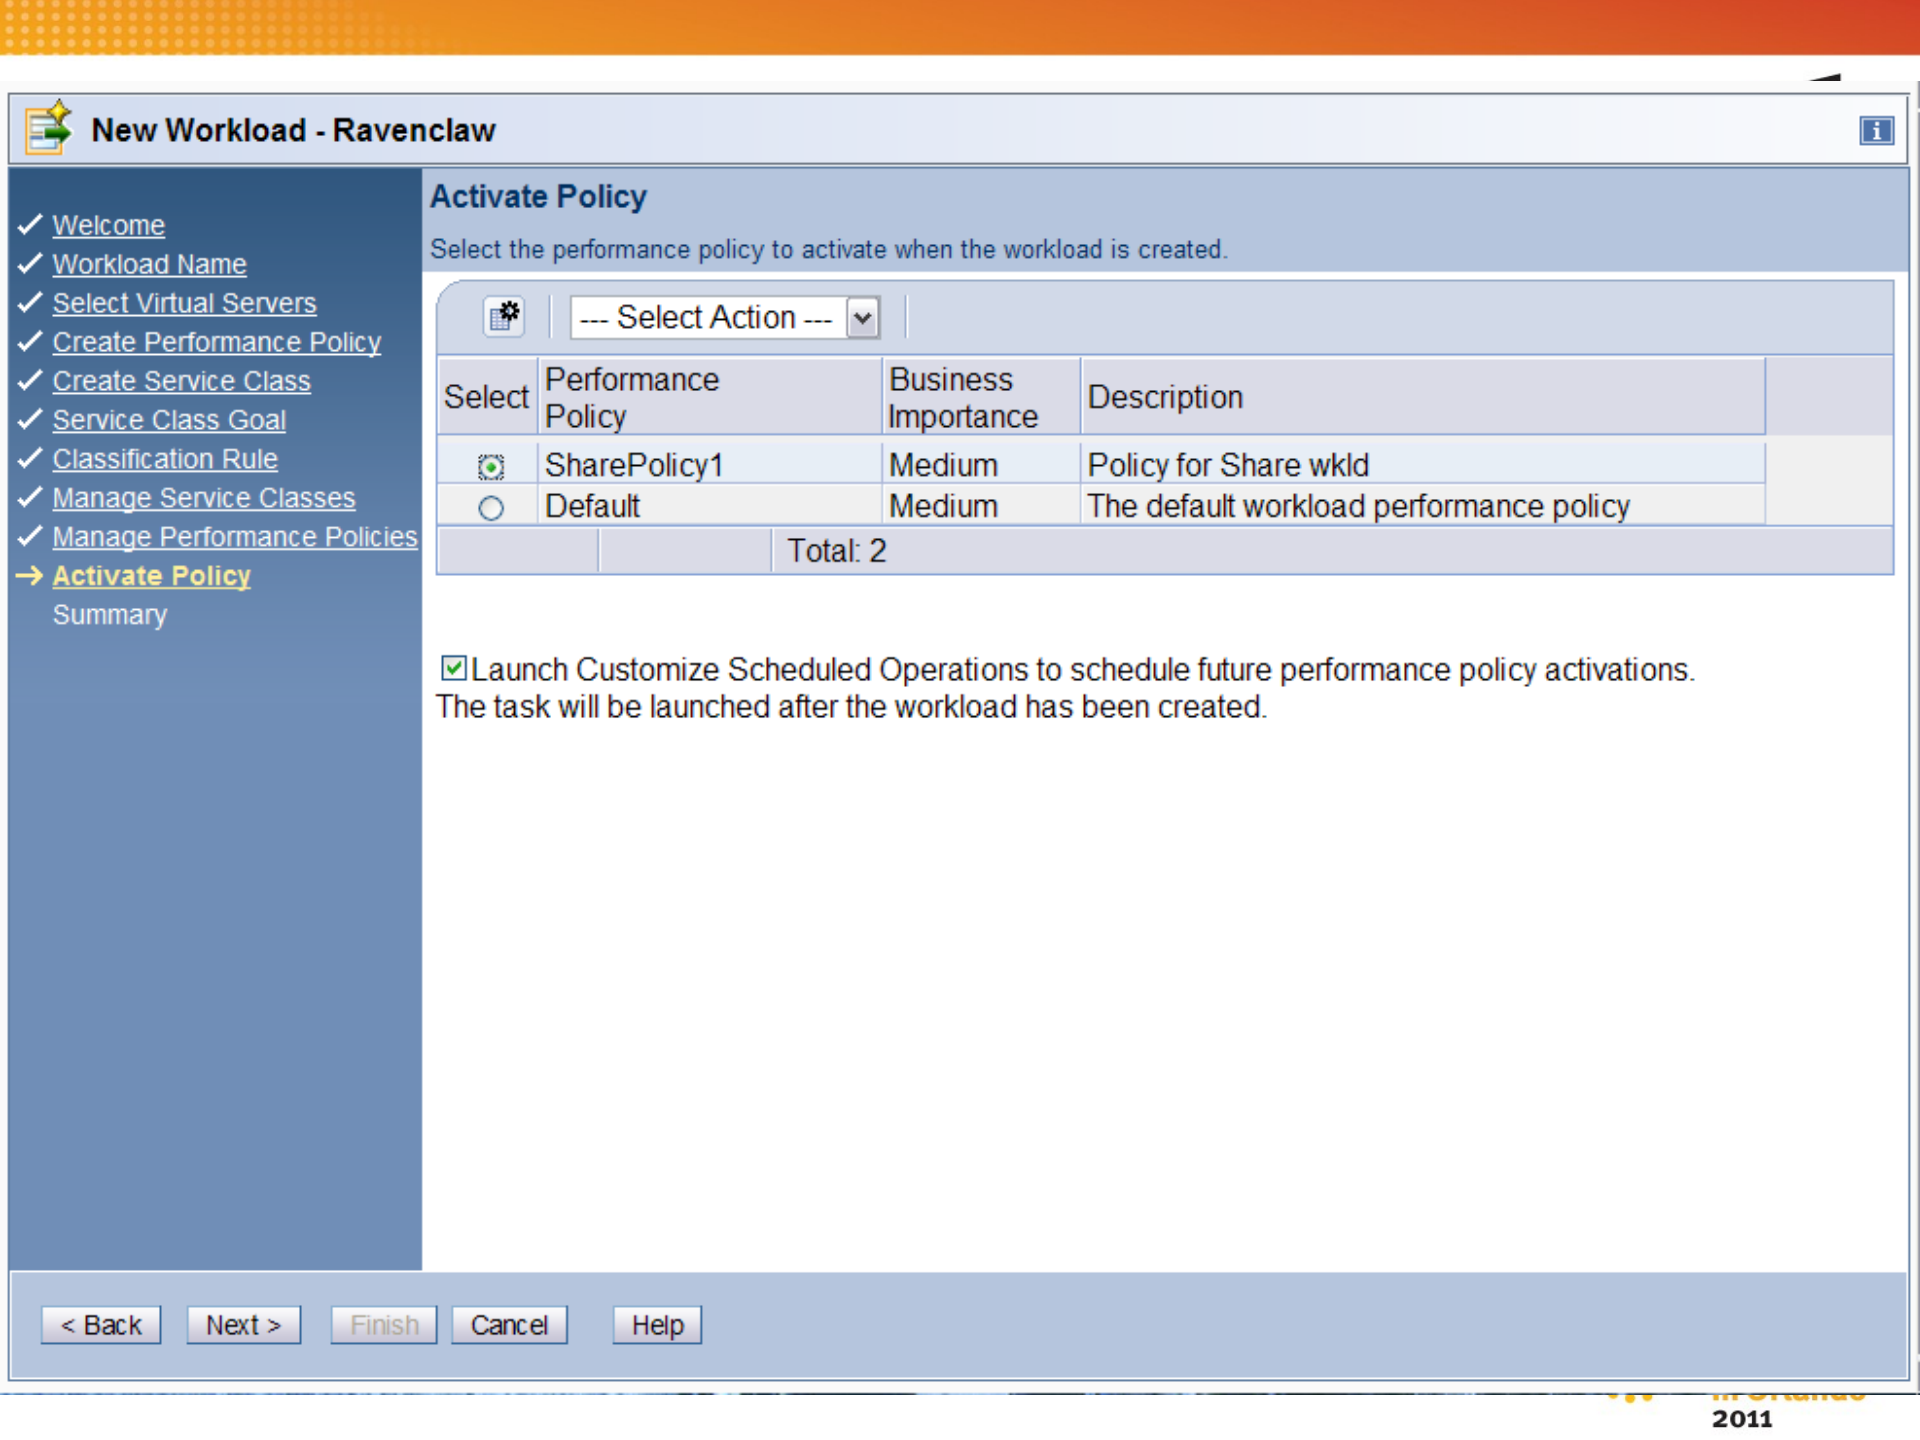
New Workload - Ravenclaw
Welcome
Workload Name
Select Virtual Servers
Create Performance Policy
Create Service Class
Service Class Goal
Classification Rule
Manage Service Classes
Manage Performance Policies
Activate Policy
Summary
Activate Policy
Select the performance policy to activate when the workload is created.
--- Select Action ---
Select
Performance Policy
Business Importance
Description
SharePolicy1
Medium
Policy for Share wkld
Default
Medium
The default workload performance policy
Total: 2
Launch Customize Scheduled Operations to schedule future performance policy activations.
The task will be launched after the workload has been created.
< Back
Next >
Finish
Cancel
Help

New Workload - Ravenclaw

- ✓ Welcome
- ✓ Workload Name
- ✓ Select Virtual Servers
- ✓ Create Performance Policy
- ✓ Create Service Class
- ✓ Service Class Goal
- ✓ Classification Rule
- ✓ Manage Service Classes
- ✓ Manage Performance Policies
- ✓ Activate Policy
- → **Summary**

Summary

Click Finish to create the workload, its performance policies and their service classes and activate the selected policy.

Workload

Name: ShareWkld1
Active performance policy: SharePolicy1
Description: Demo Workload
Category:

Virtual servers:
Custom groups:

Review the policy
Before activation

Performance Policies

SharePolicy1

Description: Policy for Share wkld
Business importance: Medium

Service Classes

Gold1

Description: Gold Service class for ShareWkld1
Performance goal: Velocity - Moderate
Business importance: Medium
Classification rule: (OS Type == "AIX"
AND Virtual Server Name == "ShareVS1")

Default

Description: The default workload performance policy service class.
Performance goal: Velocity - Moderate
Business importance: Medium

< Back | Next > | Finish | Cancel | Help

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 Based Monitoring and Reporting

- 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 Monitoring Overview

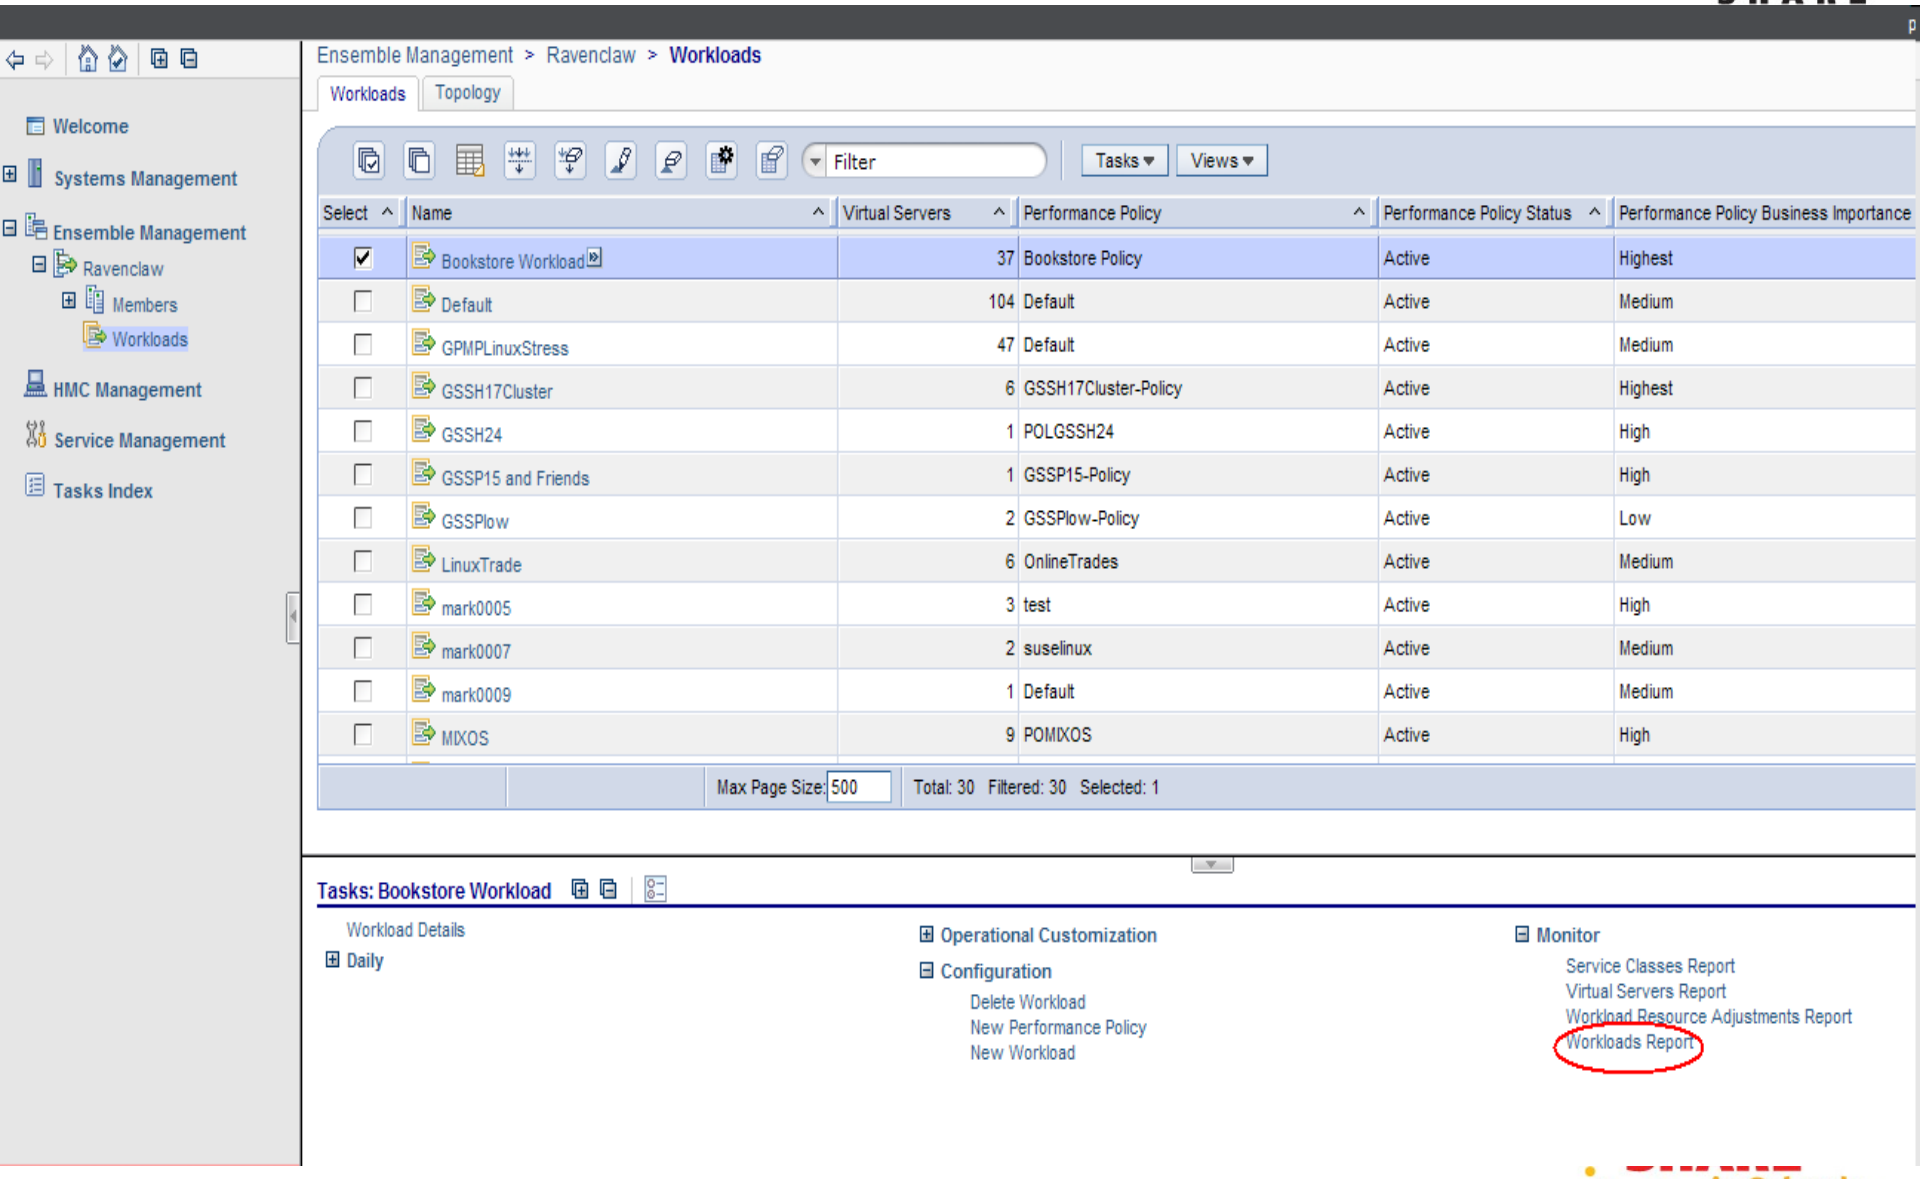

Workload Report

- 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

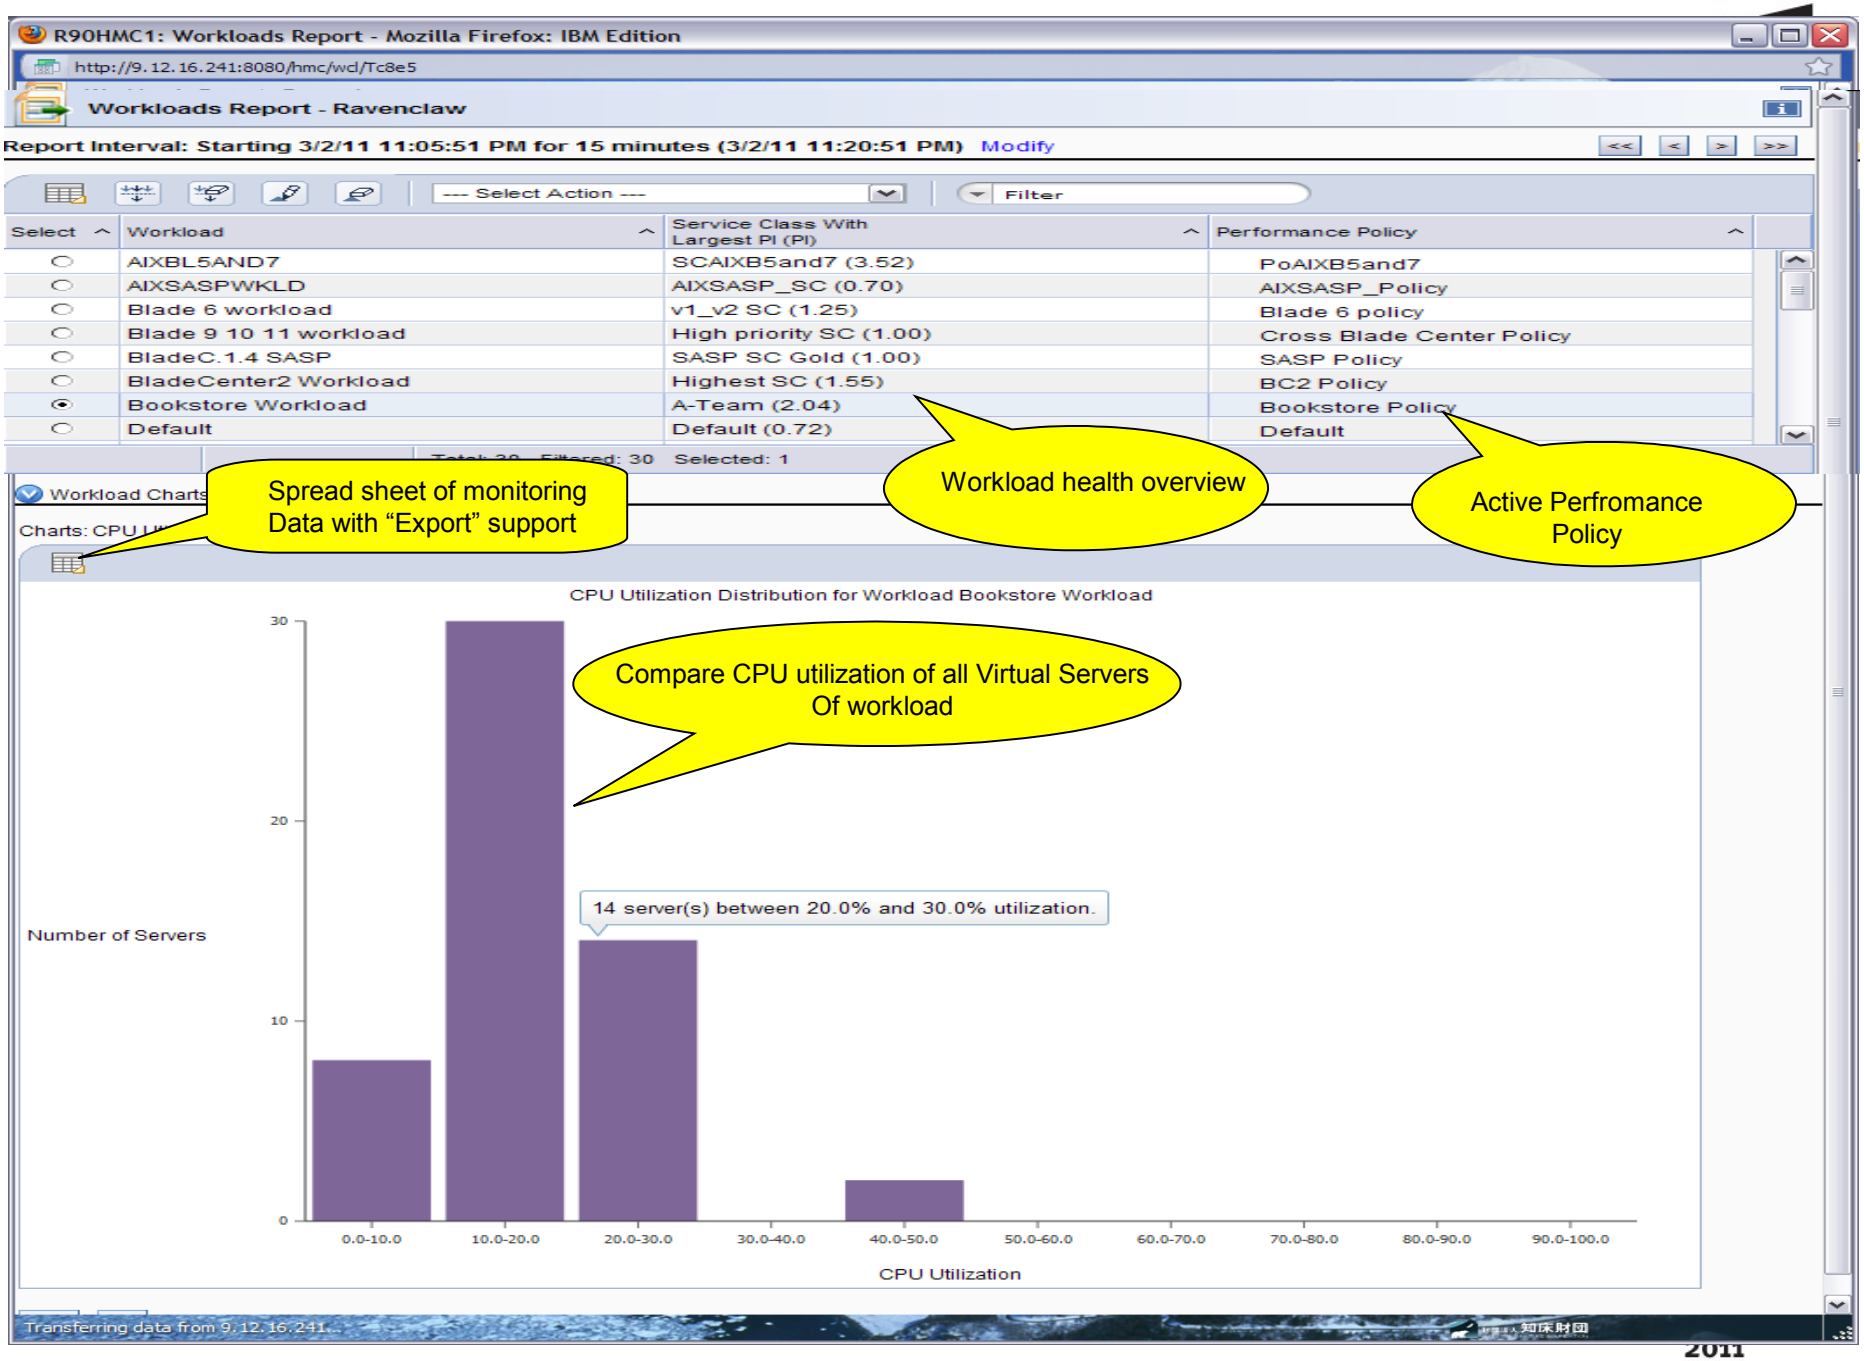
R90HMC1: Workloads Report - Mozilla Firefox: IBM Edition
http://9.12.16.241:8080/hmc/wcl/Tc8e5
Workloads Report - Ravenclaw
Report Interval: Starting 3/2/11 11:05:51 PM for 15 minutes (3/2/11 11:20:51 PM) Modify
--- Select Action ---
Filter
Select
Workload
Service Class With Largest PI (PI)
Performance Policy
AIXBL5AND7
SCAIXB5and7 (3.52)
PoAIXB5and7
AIXSASPWKLD
AIXSASP_SC (0.70)
AIXSASP_Policy
Blade 6 workload
v1_v2 SC (1.25)
Blade 6 policy
Blade 9 10 11 workload
High priority SC (1.00)
Cross Blade Center Policy
BladeC.1.4 SASP
SASP SC Gold (1.00)
SASP Policy
BladeCenter2 Workload
Highest SC (1.55)
BC2 Policy
Bookstore Workload
A-Team (2.04)
Bookstore Policy
Default
Default (0.72)
Default
Selected: 1
Workload Charts
Charts: CPU Ut
Spread sheet of monitoring Data with “Export” support
Workload health overview
Active Perfromance Policy
CPU Utilization Distribution for Workload Bookstore Workload
Compare CPU utilization of all Virtual Servers Of workload
14 server(s) between 20.0% and 30.0% utilization.
Number of Servers
30
20
10
0
0.0-10.0
10.0-20.0
20.0-30.0
30.0-40.0
40.0-50.0
50.0-60.0
60.0-70.0
70.0-80.0
80.0-90.0
90.0-100.0
CPU Utilization
Transferring data from 9.12.16.241...

Service Class Report

- 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

R90HMC1: Workloads Report - Mozilla Firefox: IBM Edition

http://9.12.16.241:8080/hmc/wcl/Te81f

Service Classes Report - Bookstore Workload

Interval can be adjusted

Report Interval: Starting 2/26/11 8:02:01 PM for 15 minutes (2/26/11 8:17:01 PM) Modify

--- Select Action --- Filter

Select	Service Class	Performance Policy	Current Performance	PI	Goal	Business importance
◉	A-Team	Bookstore Policy	Moderate	1.90	Velocity - Fast	High
○	B-Team	Bookstore Policy	Slow	1.46	Velocity - Moderate	Medium
○	BuyNode	Bookstore Policy			Velocity - Fastest	Highest
○	C-Team	Bookstore Policy	Slow	1.00	Velocity - Slow	Low
○	Dead Meat	Bookstore Policy	Slow	0.37	Velocity - Slowest	Lowest
○	Default	Bookstore Policy	Fastest	0.40	Velocity - Moderate	Medium

Total: 6 Filtered: 6 Selected: 1

Service Class Charts

Spread sheet of monitoring Data with “Export” support

Service Class Performance Index (PI) for Workload Bookstore Workload

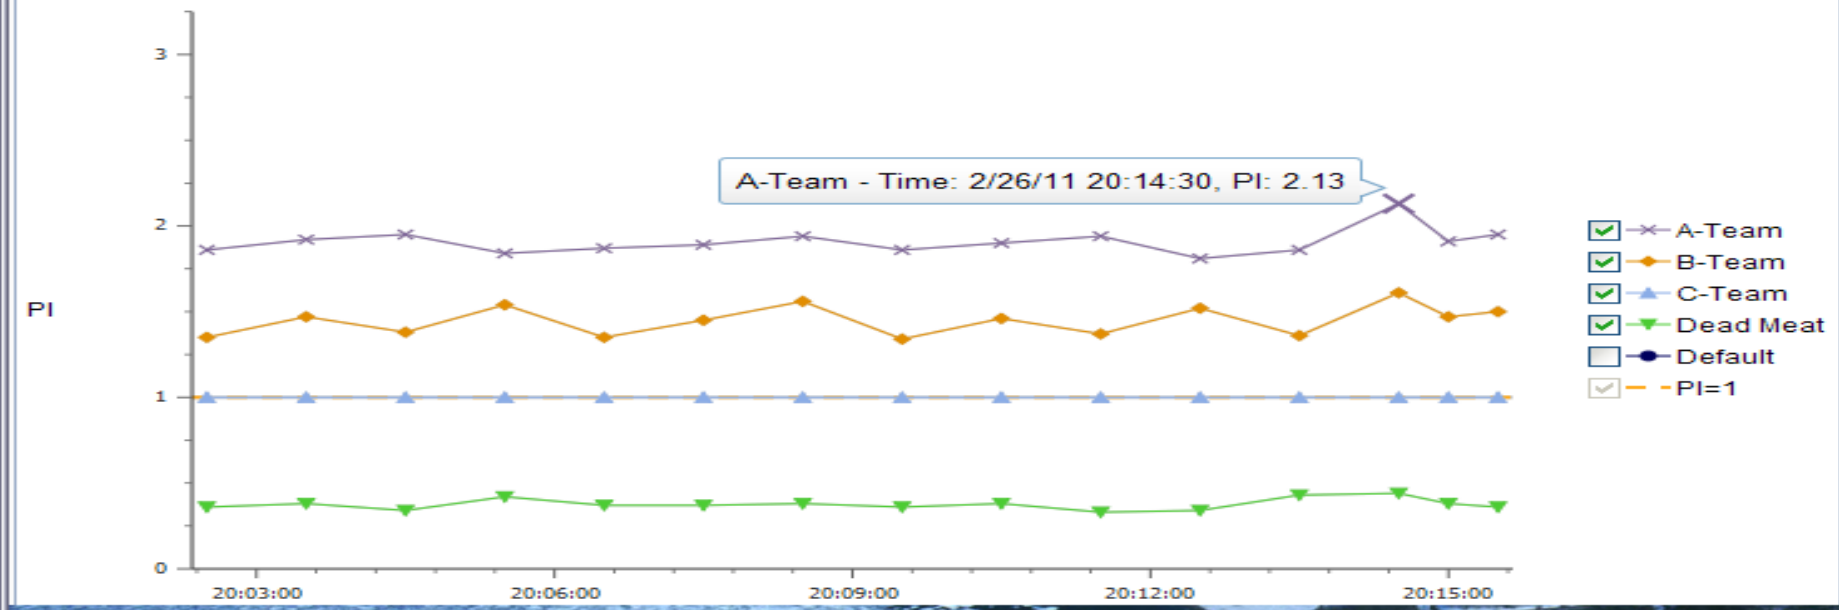

Read 9.12.16.241

Monitoring Events

- Monitoring Events (Alerts)
  - Leverage HMC event monitoring
  - Send e-mail when selected metrics reach threshold
    - Service Class PI threshold
    - Virtual Server CPU Utilization threshold

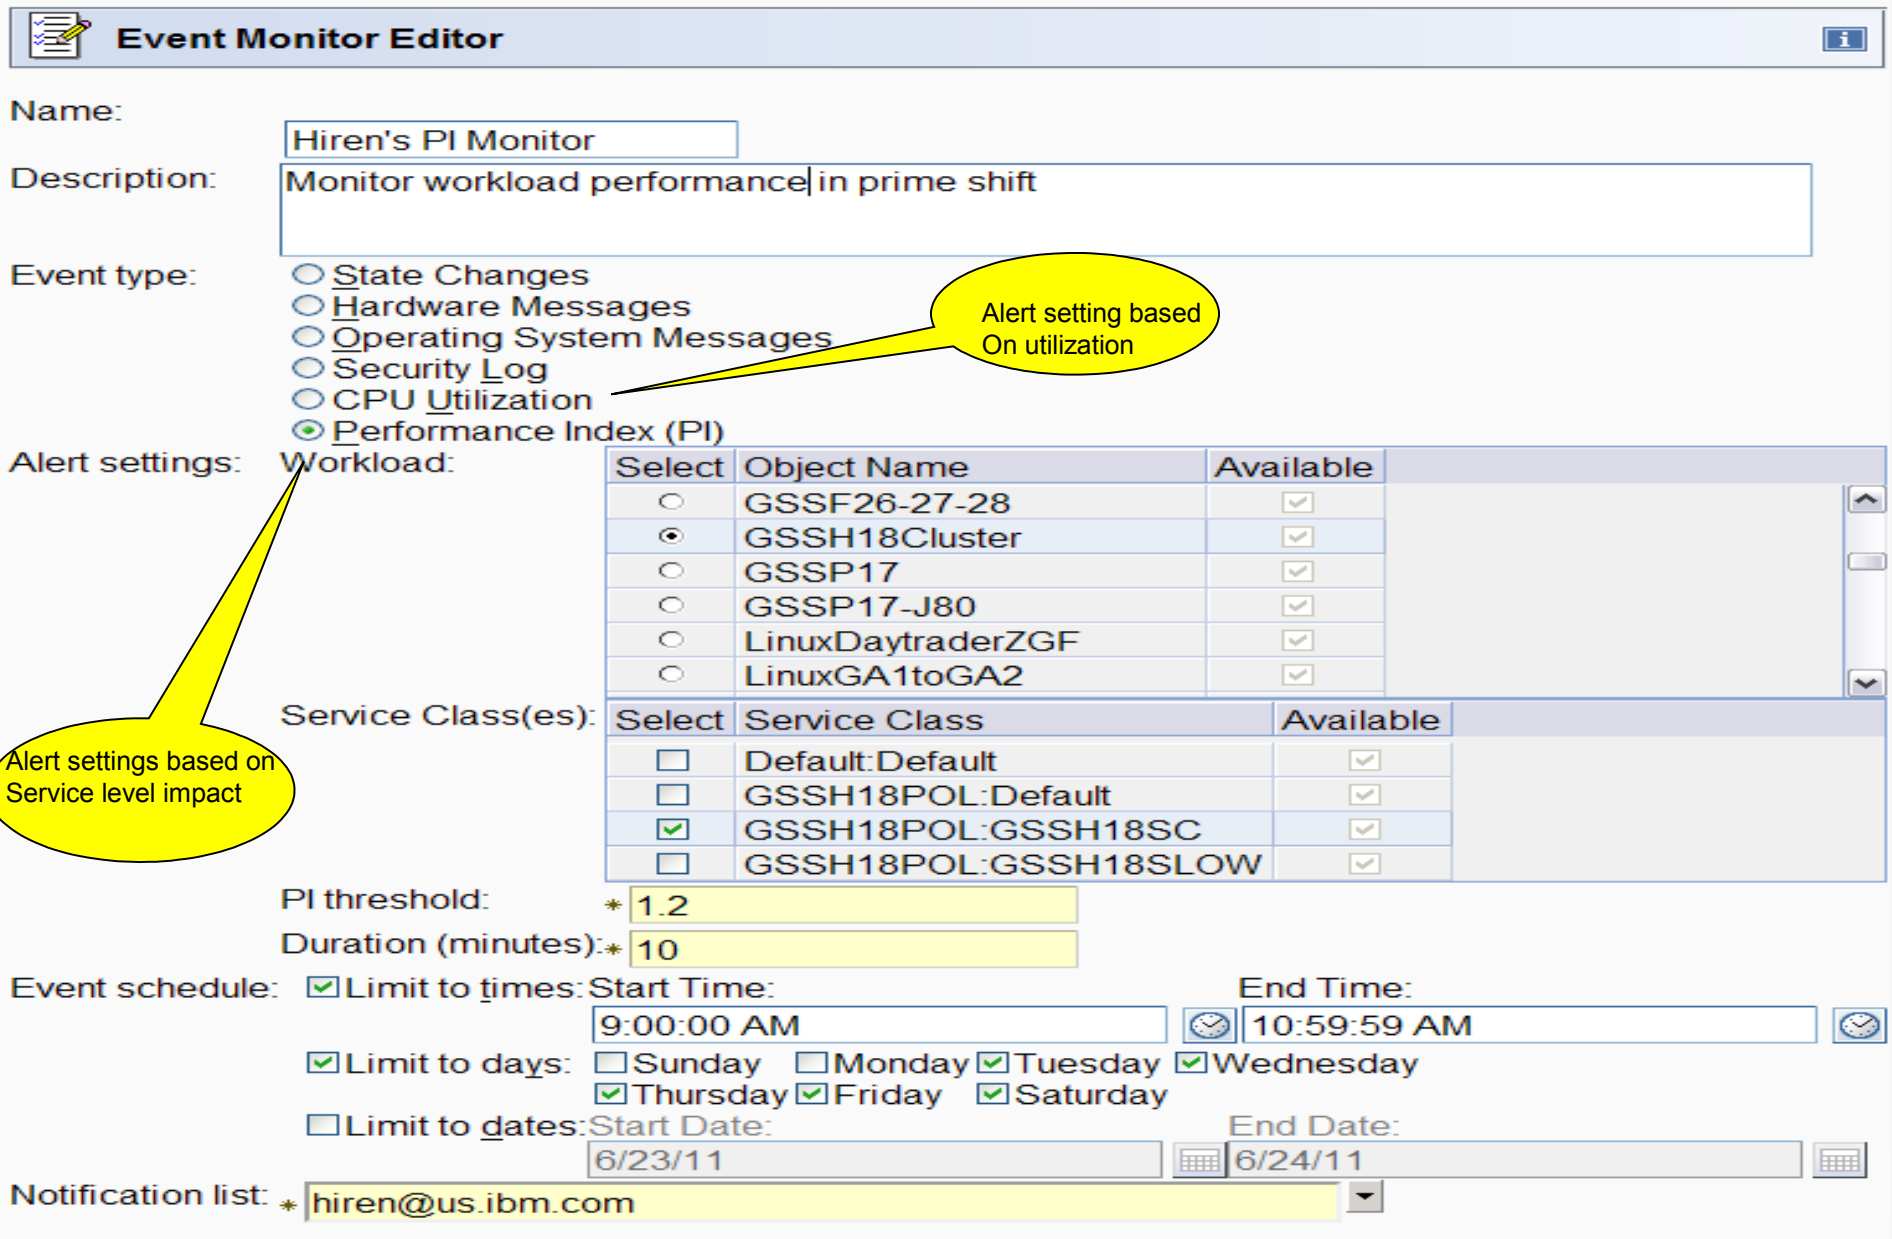
Event Monitor Editor
Name:
Hiren's PI Monitor
Description:
Monitor workload performance in prime shift
Event type:
State Changes
Hardware Messages
Operating System Messages
Security Log
CPU Utilization
Performance Index (PI)
Alert setting based On utilization
Alert settings:
Workload:
Select
Object Name
Available
GSSF26-27-28
GSSH18Cluster
GSSP17
GSSP17-J80
LinuxDaytraderZGF
LinuxGA1toGA2
Service Class(es):
Alert settings based on Service level impact
Select
Service Class
Available
Default:Default
GSSH18POL:Default
GSSH18POL:GSSH18SC
GSSH18POL:GSSH18SLOW
PI threshold:
1.2
Duration (minutes):
10
Event schedule:
Limit to times:
Start Time:
9:00:00 AM
End Time:
10:59:59 AM
Limit to days:
Sunday
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday
Limit to dates:
Start Date:
6/23/11
End Date:
6/24/11
Notification list:
hiren@us.ibm.com

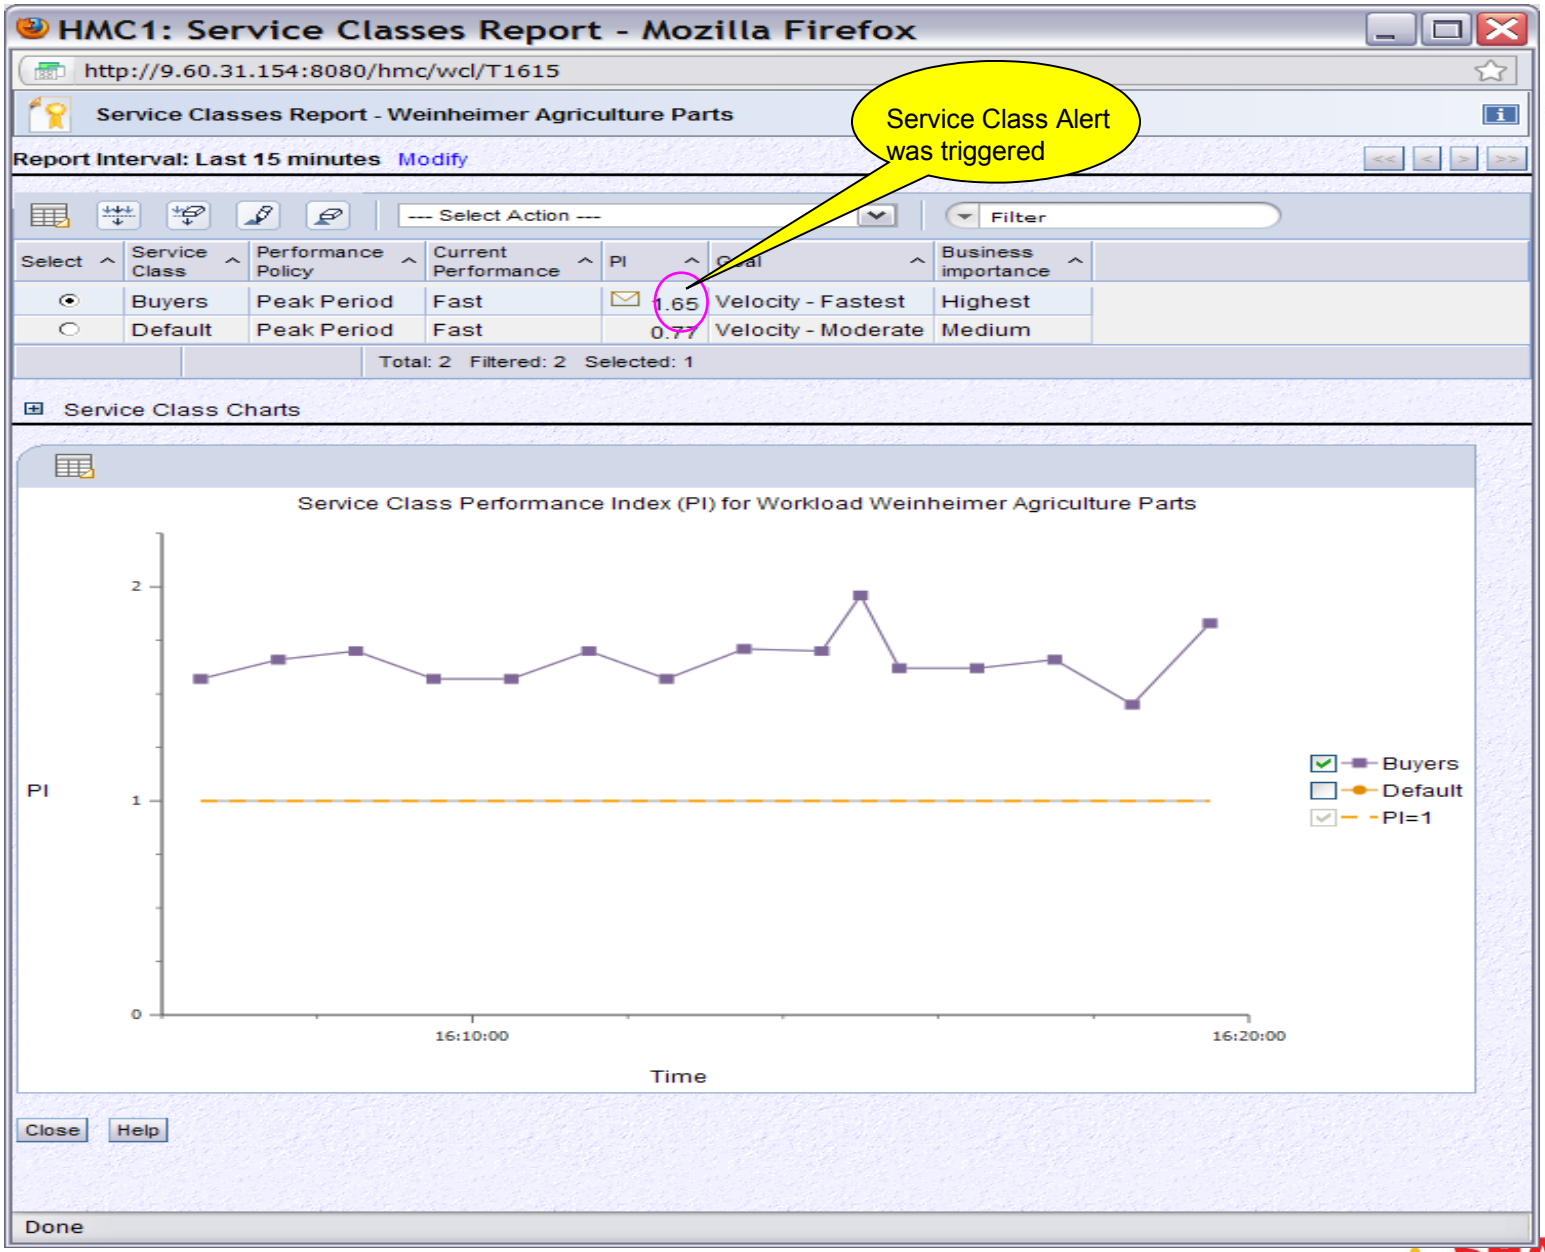

Virtual Server Report

- Virtual server report
  - List of virtual servers associated with the workload
    - Virtual server Service class PI
    - Resource usage
      - *Physical CPU utilization*
      - *OS view of CPU utilization*
      - *Physical memory used*
    - Hypervisor delay percentage
  - CPU Utilization trend for the selected interval
  - Launch Monitor dashboard
    - Provides hardware utilization data

Resource utilized by VS

Virtual Servers Report - Bookstore Workload

Report Interval: Starting 2/26/11 8:02:01 PM for 15 minutes (2/26/11 8:17:01 PM) Modify

--- Select Action --- Filter

Select	Virtual Server	OS Name	OS Type	OS Level	Virtual Processors	Allocated Memory (MB)	Physical CPU Utilization (%)	Hypervisor CPU Delay (%)	Other Time (%)	Service Class (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 (%)
◉	r90b208v1	r90f1c1b8v1	AIX	6.1.5.2	8	8,192	24.3	30.8		A-Team (1.90)	28.9	21.7	0.0	49.4
○	r90b208v2	r90f1c1b8v2	AIX	6.1.5.2	8	8,192	11.0	27.8		A-Team (1.90)	15.3	6.0	0.0	78.7
○	r90b208v3	r90f1c1b8v3	AIX	6.1.5.2	8	4,096	10.4	34.9		B-Team (1.46)	14.5	7.2	0.0	78.3
○	r90b208v4	r90f1c1b8v4	AIX	6.1.5.2	8	4,096	22.6	35.3		B-Team (1.46)	28.6	22.9	0.0	48.4
○	r90b208v5	r90f1c1b8v5	AIX	6.1.5.2	8	4,096	10.6	18.0		C-Team (1.00)	14.0	8.1	0.0	77.8
○	r90b208v6	r90f1c1b8v6	AIX	6.1.5.2	8	4,096	10.5	35.3		C-Team (1.00)	14.0	7.7	0.0	78.4
○	r90b208v7	r90f1c1b8v7	AIX	6.1.5.2	8	4,096	3.0	28.1		Dead Meat (0.37)	4.5	7.3	0.0	88.2
○	r90b208v8	r90f1c1b8v8	AIX	6.1.5.2	8	4,096	3.0	28.4		Dead Meat (0.37)	4.6		0.0	88.6

Page 1 of 1 Total: 37 Filtered: 37 Displayed: 37 Selected: 1

Monitoring data Provided by GPMP Running on VS

Virtual Server Charts

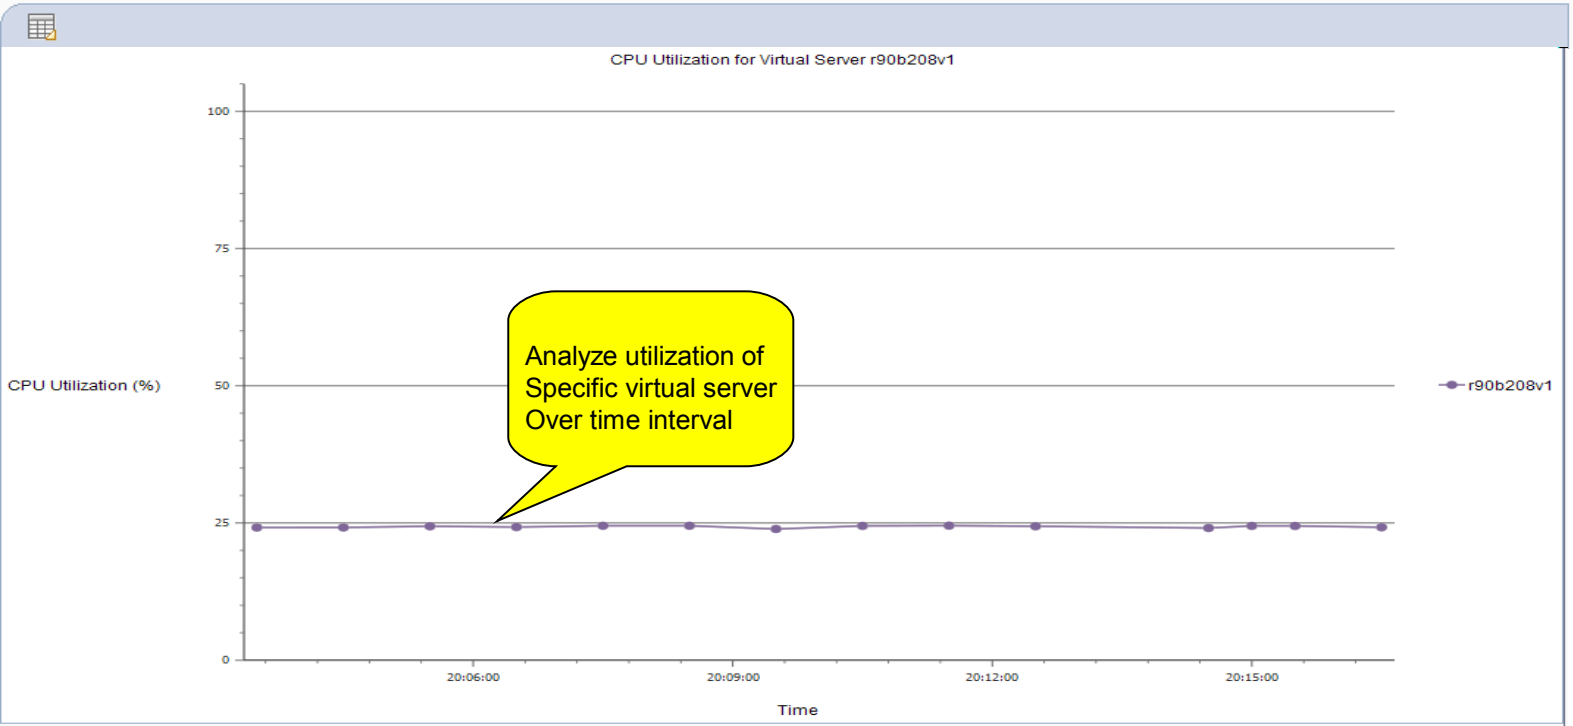

Analyze utilization of Specific virtual server Over time interval

SHARE in Orlando 2011

Hypervisor Report

- Hypervisor report
  - Hypervisor resource utilization
  - List of all virtual servers on hypervisor
    - Virtual server Resource allocations (e.g. Memory, CPU)
    - Virtual server current Resource usage
      - *Physical CPU utilization*
      - *Physical memory used*
    - Hypervisor delay percentage
  - Resource adjustment report
    - Resource adjustment actions taken over report interval

Hypervisor Report - r93f2c1b09v1

Report Interval: Starting 6/23/11 6:04:09 PM for 15 minutes (6/23/11 6:19:09 PM) Modify

Hypervisor Details:

Hypervisor:	C.1.09	Processor count:	8	Total memory allocated for virtual servers:	31,232 MB
Hypervisor type:	PowerVM	Total CPU consumption:	89.4%	Total memory:	32,768 MB
Total allocated processing units:	7.96				

Hypervisor Resource Allocation and utilization

Virtual Servers:

--- Select Action --- Filter

Virtual Server	Processor Management Status	Processor Management Reason	Virtual Processors	Min Virtual Processors	Max Virtual Processors	Consumed Processors	Hypervisor Processing Unit Delay (%)	Allocated Memory (MB)	Dedicated	Capped	Processing Units	Initial Processing Units	Min Processing Units	Max Processing Units	Min Memory (MB)	Max Memory (MB)
r93f2c1b09v1	Active	None	2	1	7	0.03	2.2	2,560	--	--	0.34	0.70	0.10	7.00	2,048	4,096
r93f2c1b09v10	Active	None	2	1	7	1.19	24.3	2,560	--	--	0.58	0.70	0.10	7.00	1,024	4,096
r93f2c1b09v2	Active	None	2	1	7	0.10	8.0	2,560	--	--	0.72	0.72	0.10	7.00	1,024	4,096
r93f2c1b09v3	Active	None	2	1	7	0.99	3.9	2,560	--	--	1.81	0.72	0.10	7.00	1,024	4,096
r93f2c1b09v4	Active	None	2	1	7	0.94	20.8	2,560	--	--	0.59	0.72	0.10	7.00	1,024	4,096
r93f2c1b09v5	Active	None	2	1	7	1.19	20.8	2,560	--	--	0.59	0.72	0.10	7.00	1,024	4,096
r93f2c1b09v6	Active	None	2	1	7	1.18	19.9	4,096	--	--	0.61	0.72	0.10	7.00	4,096	4,096
r93f2c1b09v7	Active	None	2	1	7	1.20	19.5	2,560	--	--	0.62	0.72	0.10	7.00	1,024	4,096
r93f2c1b09v8	Active	None	2	1	7	0.03	1.7	2,560	--	--	0.59	0.72	0.10	7.00	1,024	4,096
r93f2c1b09v9	Active	None	2	1	7	0.03	0.2	2,560	--	--	0.71	0.72	0.10	7.00	1,024	4,096

Page 1 of 1 Total: 10 Filtered: 10 Displayed: 10

Difference due to Dynamic CPU mgmt

Benefits of GPMP

- Guest Platform Management Provider (GPMP) is a lightweight component of PPM that provides additional monitoring data
- Allows virtual server to be classified with 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

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.

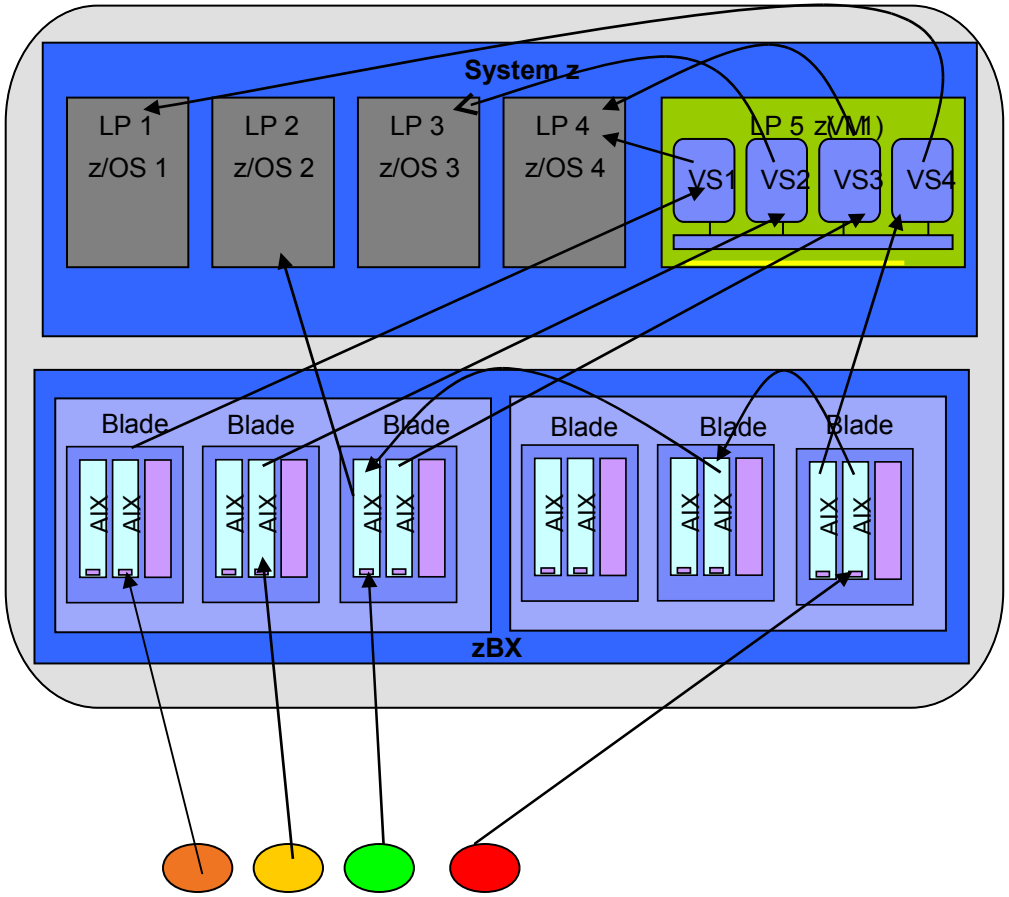

Basic ARM calls

- arm_register_application
- arm_register_transaction
- arm_start_application
- for (each transaction)
  - arm_start_transaction
  - arm_bind_thread
  - arm_blocked
    - Call downstream sub-transaction
  - arm_unblocked
  - arm_unbind_thread
  - arm_stop_transaction
- arm_stop_application
- arm_destroy_application

Application Response Measurement (ARM)

Standards Based Application Instrumentation

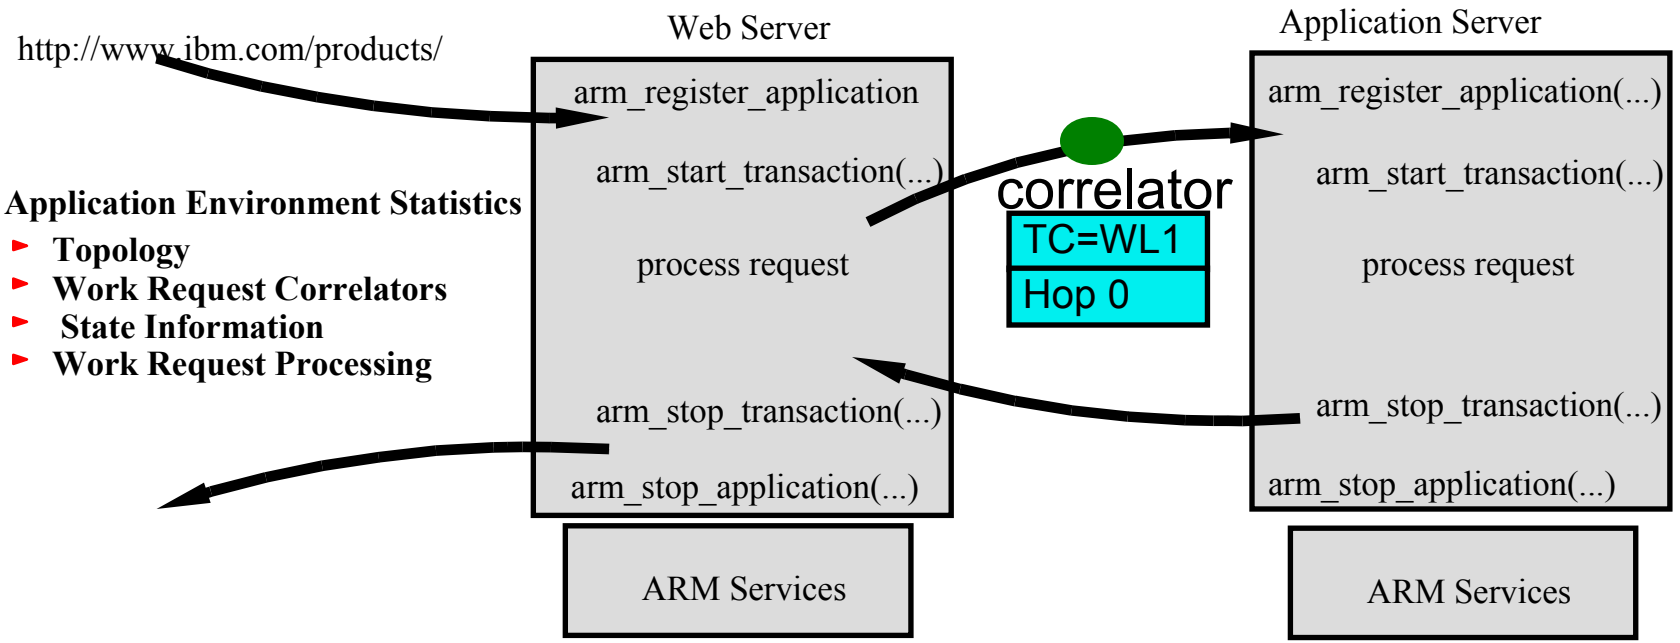

- **Application Environment Statistics**
  - **Topology**
  - **Work Request Correlators**
  - **State Information**
  - **Work Request Processing**

- **The Open Group ARM Standard V4.0**
  - Process registration, deregistration
  - Work request classification, start, and stop

ARM 4.0 Instrumented Middleware

- Web Server support:
  - WebSphere provided plugin
    - IHS/Apache
    - IIS
    - Domino
    - iPlanet
- WebSphere Application Server
  - WAS 6.0, WAS 7.0
- DB2 Universal Database

Enable ARM Services on Middleware Applications

- Ensure Guest Platform Management Provider status in HMC
- Ensure that user account under which the application will run is authorized to the ARM services
- Enable ARM services on the middleware application
- Verify that ARM is enabled
  - “lsarm –a” command

Workload Monitoring Overview…

Transaction Hops and topology report

- Different hops involved in processing of business transaction
  - Based on OpenGroup ARM Standard instrumentation
  - Middleware instrumentation to ARM APIs
  - Guest Platform Management Provider (GPMP) collects the transaction statistics
  - Provides detail view of resources consumed by instrumented applications

Hops and Topology report with GPMP active

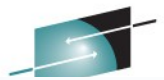

Hops Report - Blade46Medium in Workload WkldForBlade4and6

Report Interval: Starting 6/23/11 5:41:43 PM for 15 minutes (6/23/11 5:56:43 PM) Modify

Details for Blade46Medium:

Workload:	WkldForBlade4and6*	Performance goal:	Velocity - Moderate	PI:	0.75
Performance policy:	Blade46Po	Business importance:	Medium	Performance:	Fast

--- Select Action --- Filter

Name	Hop Number	Group Name	Successful Transactions	Failed Transactions	Stopped Transactions	Inflight Transactions	Queue Time (s)	Execution Time (s)	Successful Average Response Time (s)	Inflight Average Response Time (s)
Hop 0	0		57,744	0	0	230	0.000000	0.000260	3.742155	1.862512
IBM DB2 Universal Database	0	db2inst1	126	0	0	0	0.000000	0.000047	0.000047	0.000000
IBM Webserving Plugin	0	IBM_HTTP_Server	57,618	0	0	230	0.000000	0.000261	3.750339	1.862512
r93f2c1b06v2	0		57,618	0	0	230	0.000000	0.000261	3.750339	1.862512
Hop 1	1		4,175	0	0	45	0.000000	0.005133	0.710093	0.582881
WebSphere:APPLICATION_SERVER	1	server1	4,175	0	0	45	0.000000	0.005133	0.710093	0.582881
r93f2c1b06v1	1		4,175	0	0	45	0.000000	0.005133	0.710093	0.582881
Hop 2	2		63,408	0	0	2	0.000000	0.000579	0.000579	0.003900
IBM DB2 Universal Database	2	db2inst1	63,408	0	0	2	0.000000	0.000579	0.000579	0.003900
r93f2c1b06v1	2		63,408	0	0	2	0.000000	0.000579	0.000579	0.003900

Page 1 of 1 Total: 10 Filtered: 10 Displayed: 10

Detail hop report With transaction avg Response time

Do we have any failure?

How much time it took for transaction execution?

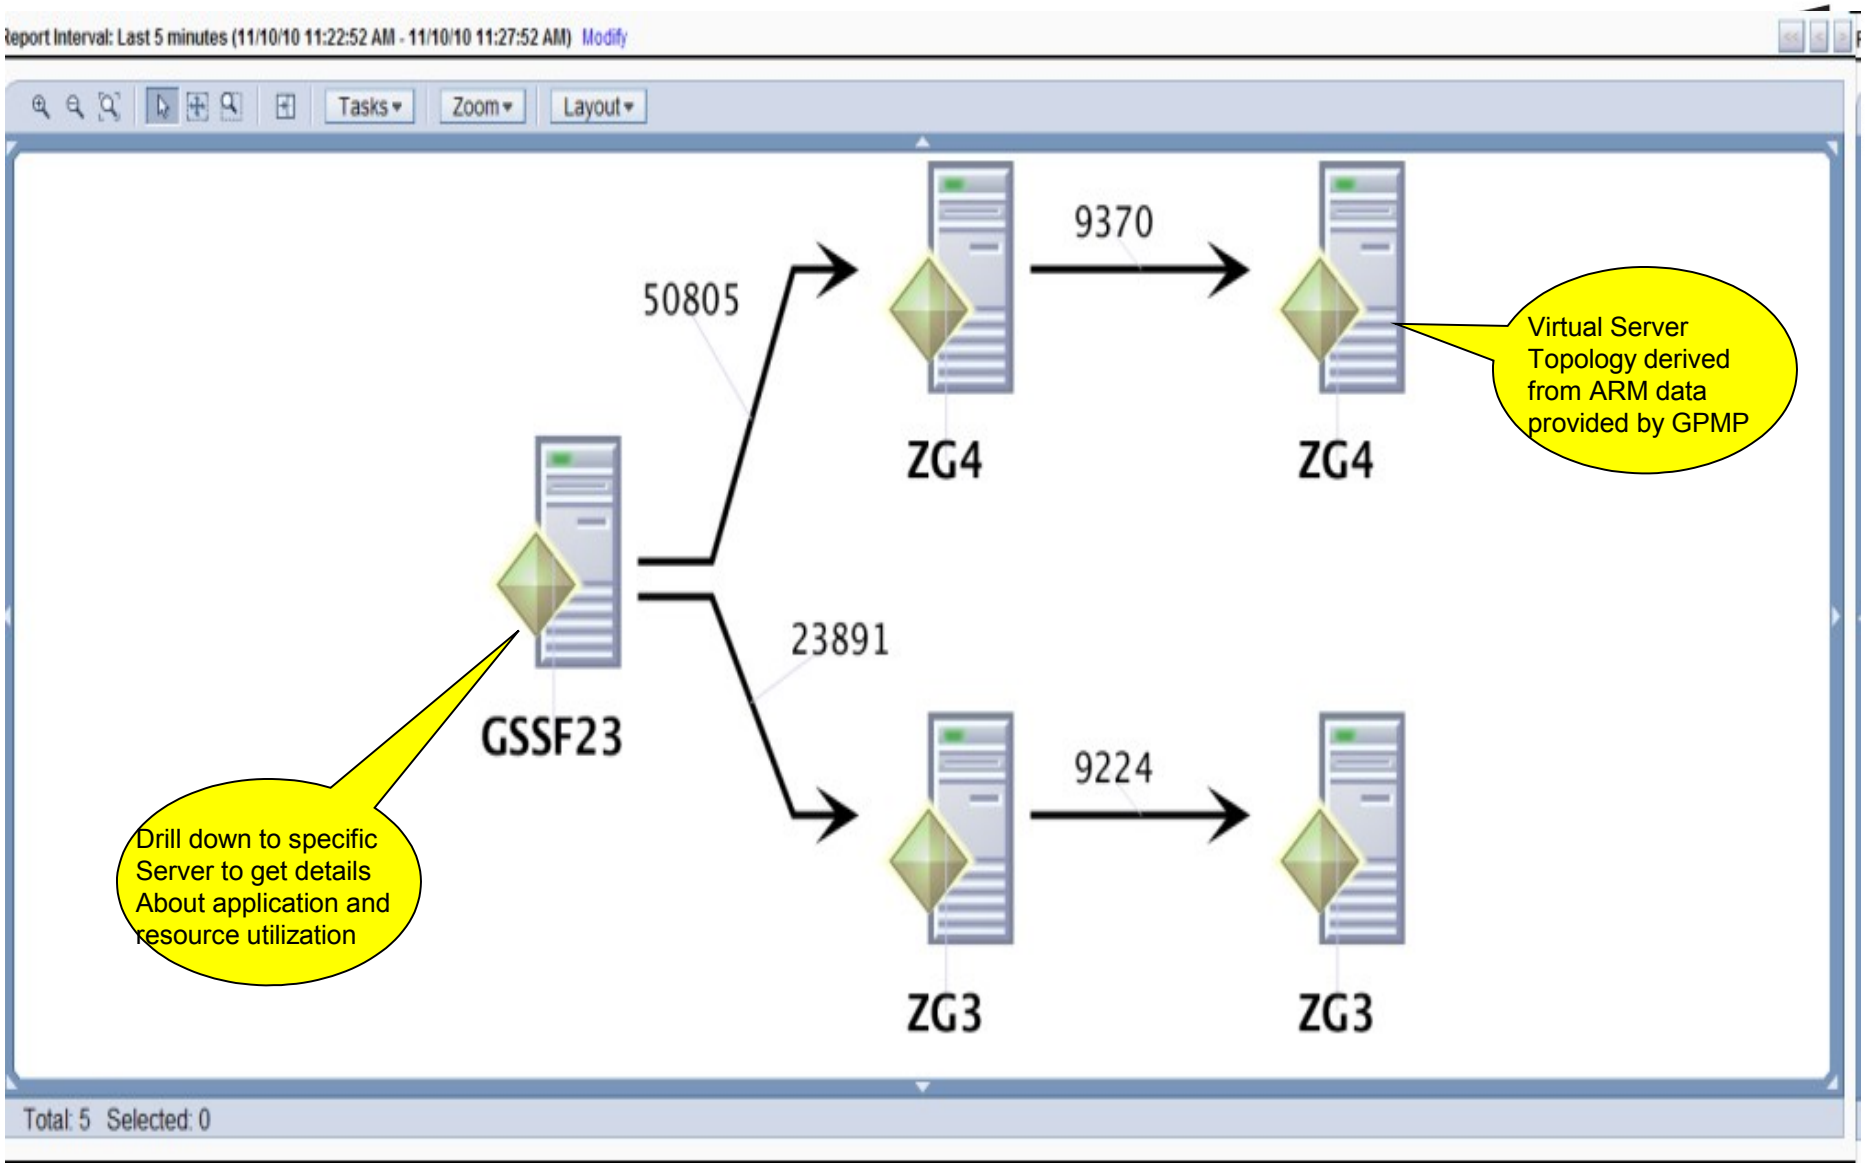
Report Interval: Last 5 minutes (11/10/10 11:22:52 AM - 11/10/10 11:27:52 AM) Modify
Tasks
Zoom
Layout
50805
9370
ZG4
ZG4
Virtual Server Topology derived from ARM data provided by GPMP
GSSF23
23891
9224
ZG3
ZG3
Drill down to specific Server to get details About application and resource utilization
Total: 5 Selected: 0

View Statistics - ZNTC25

Report Interval: Starting 6/22/11 3:50:14 PM for 15 minutes (6/22/11 4:05:14 PM)

Virtual Server Statistics:

Physical CPU utilization: 25.2% Hypervisor CPU delay: 27.3%
Idle time: 32.2% Other time: 0.0%

OS Processes Totals:

CPU using samples: 7.0% CPU delay samples: 92.9%
Page delay samples: 0.0% I/O delay samples: 0.0%

Application Environment Server Response Time Data:

--- Select Action --- Filter

Application Environment	Group Name	Successful Transactions	Failed Transactions	Stopped Transactions	Inflight Transactions	Queue Time (s)	Execution Time (s)	Successful Average Response Time (s)	Inflight Average Response Time (s)
IBM Webserving Plugin	Apache/2.2.3 (Linux/SUSE)	409,359	6,034	18,728	0	0.000000	0.000494	0.000494	0.000000

Total: 1 Filtered: 1

Application Environment Server Utilization:

--- Select Action --- Filter

Application Environment	Group Name	CPU Time (s)	CPU Using Samples (%)	CPU Delay Samples (%)	Page Delay Samples (%)	I/O Delay Samples (%)
IBM Webserving Plugin	Apache/2.2.3 (Linux/SUSE)	0.000000	0.0	0.0	0.0	0.0
IBM Webserving Plugin	Apache/2.2.3 (Linux/SUSE)	445.264282	7.9	91.0	0.0	1.2

Total: 2 Filtered: 2

CPU Time used by Apache

Management Functions

Managing Resources across Virtual Servers on Power Blade

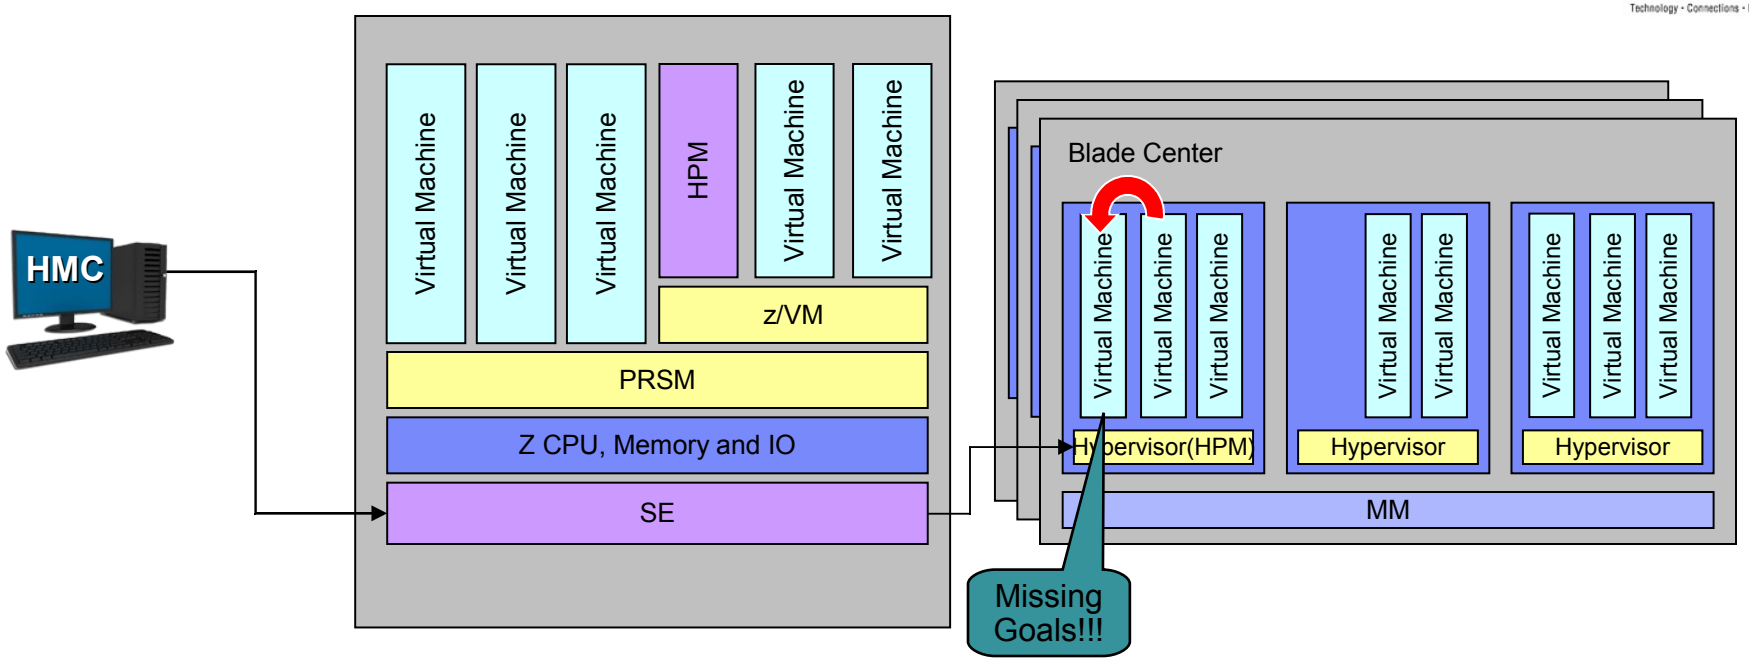

- Manage processor resources across virtual servers to achieve workload goals
  - Detect that a virtual server is part of Workload not achieving goals
  - Determine that the virtual server performance can be improved with additional resources
  - Project impact on all effected Workloads of moving resources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

Managing Resources across z/VM Virtual Machines

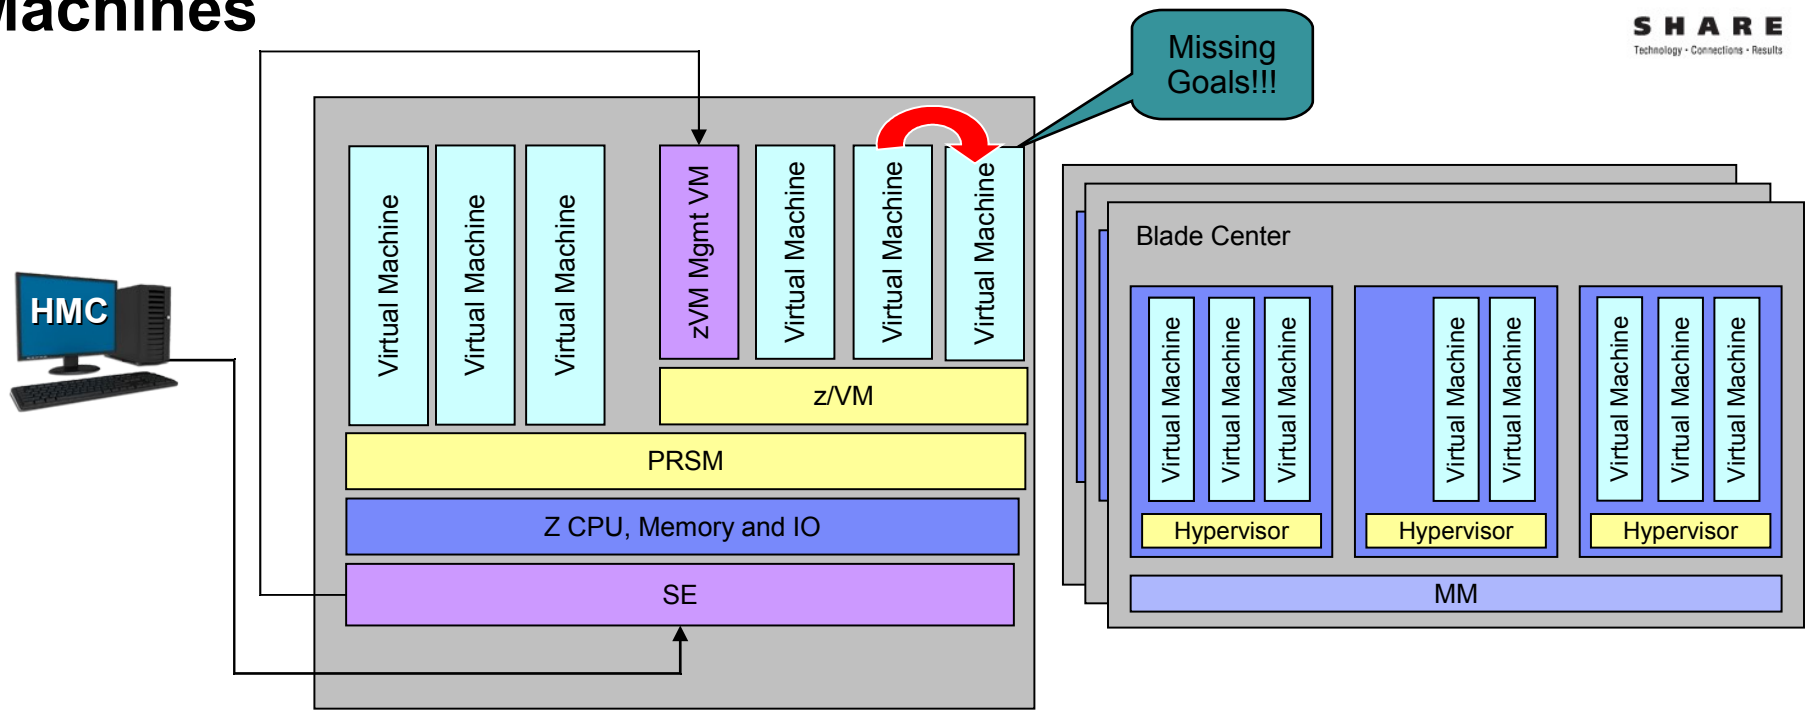

- Manage processor resources across z/VM virtual machines
  - Detect that a virtual machine that is part of Workload is not achieving goals
  - Determine that the virtual machine performance can be improved with additional resources
  - Project impact on all effected Workloads of moving resources to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

Resource Adjustment Report

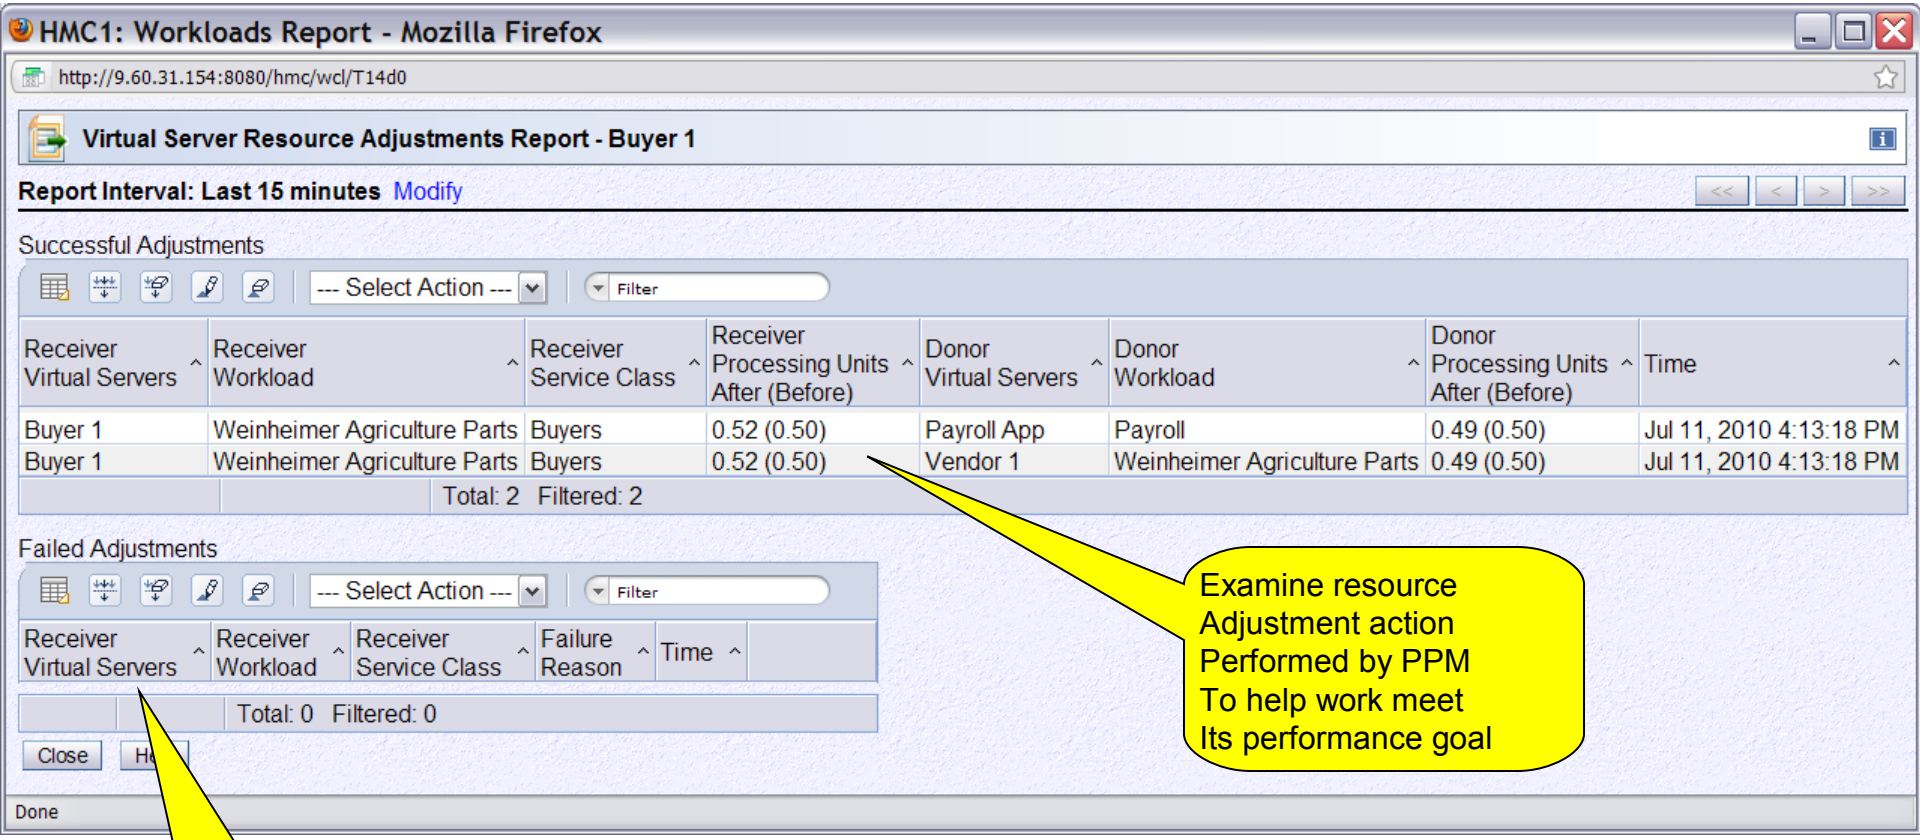

Examine resource Adjustment action Performed by PPM To help work meet Its performance goal

Explains why resource Adjusment action was Not perfromed

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
- WLM service definition needs to map PPM service classes to z/OS WLM service classes
- 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.

Setup for co-operative mgmt with z/OS WLM

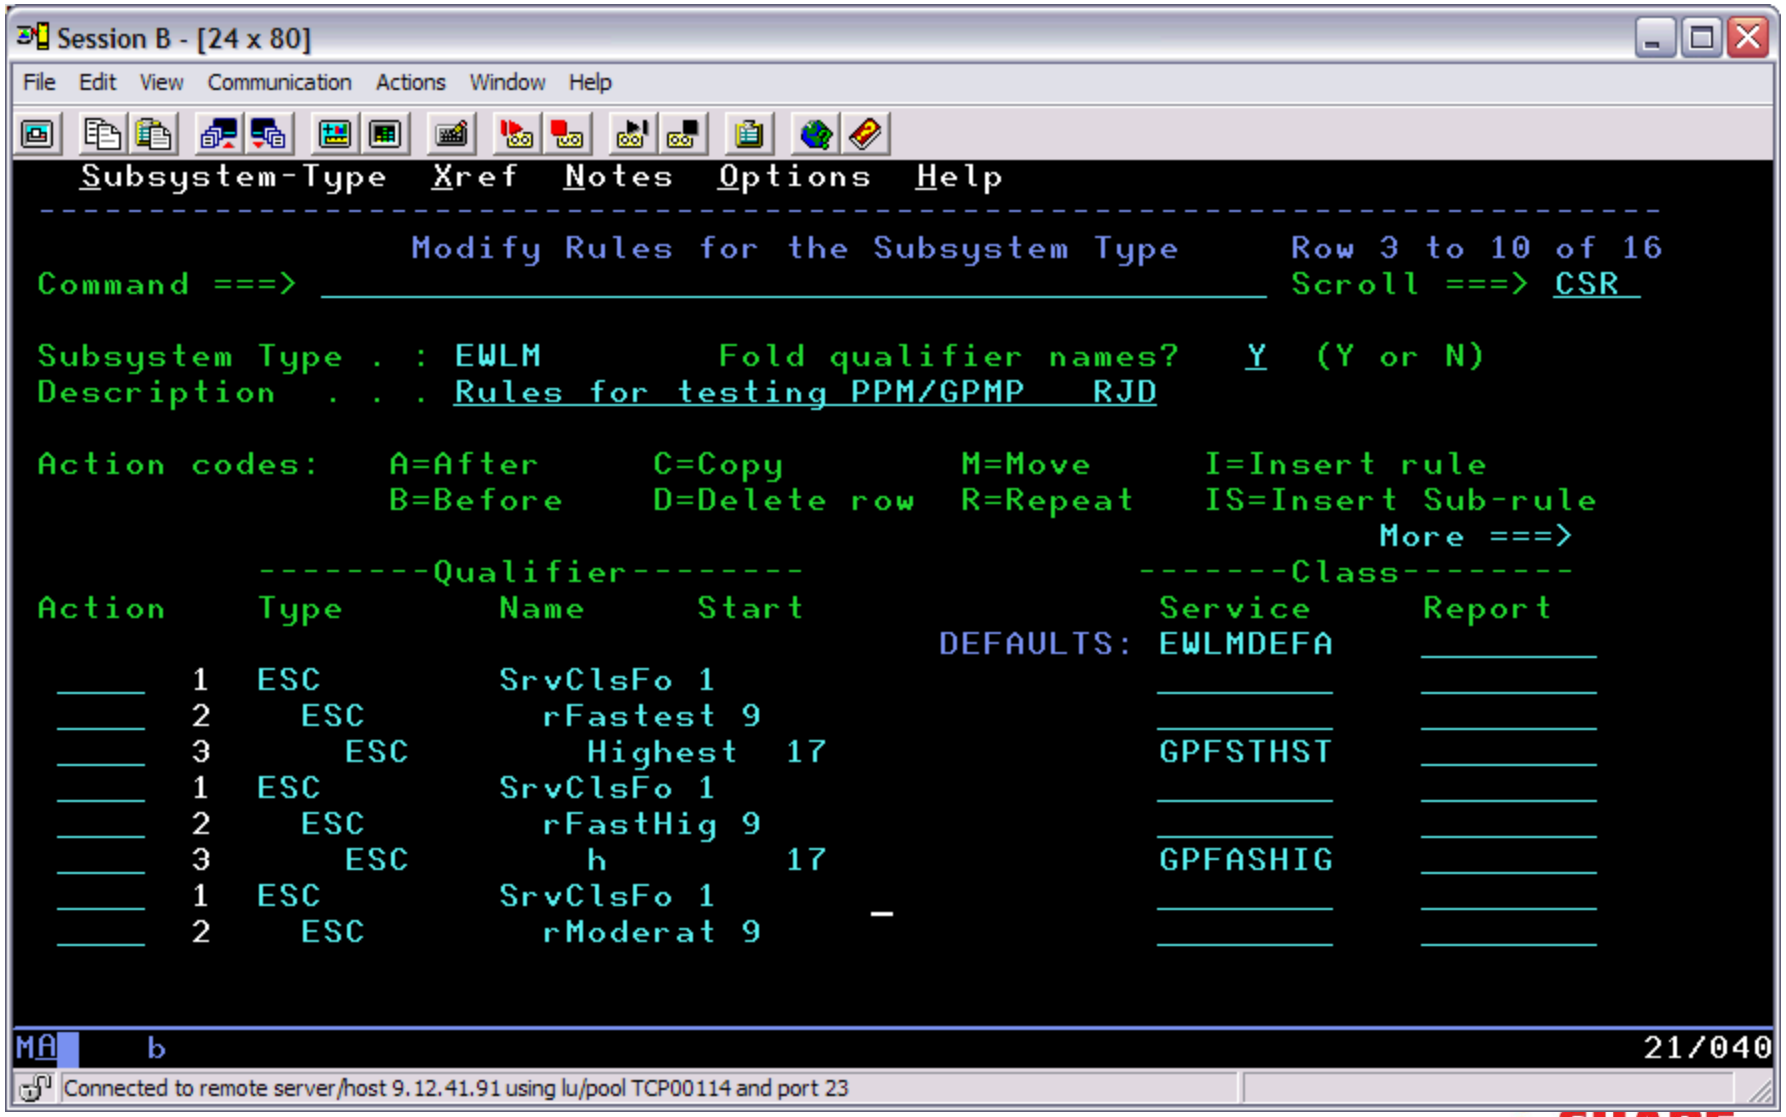

Monitoring with RMF

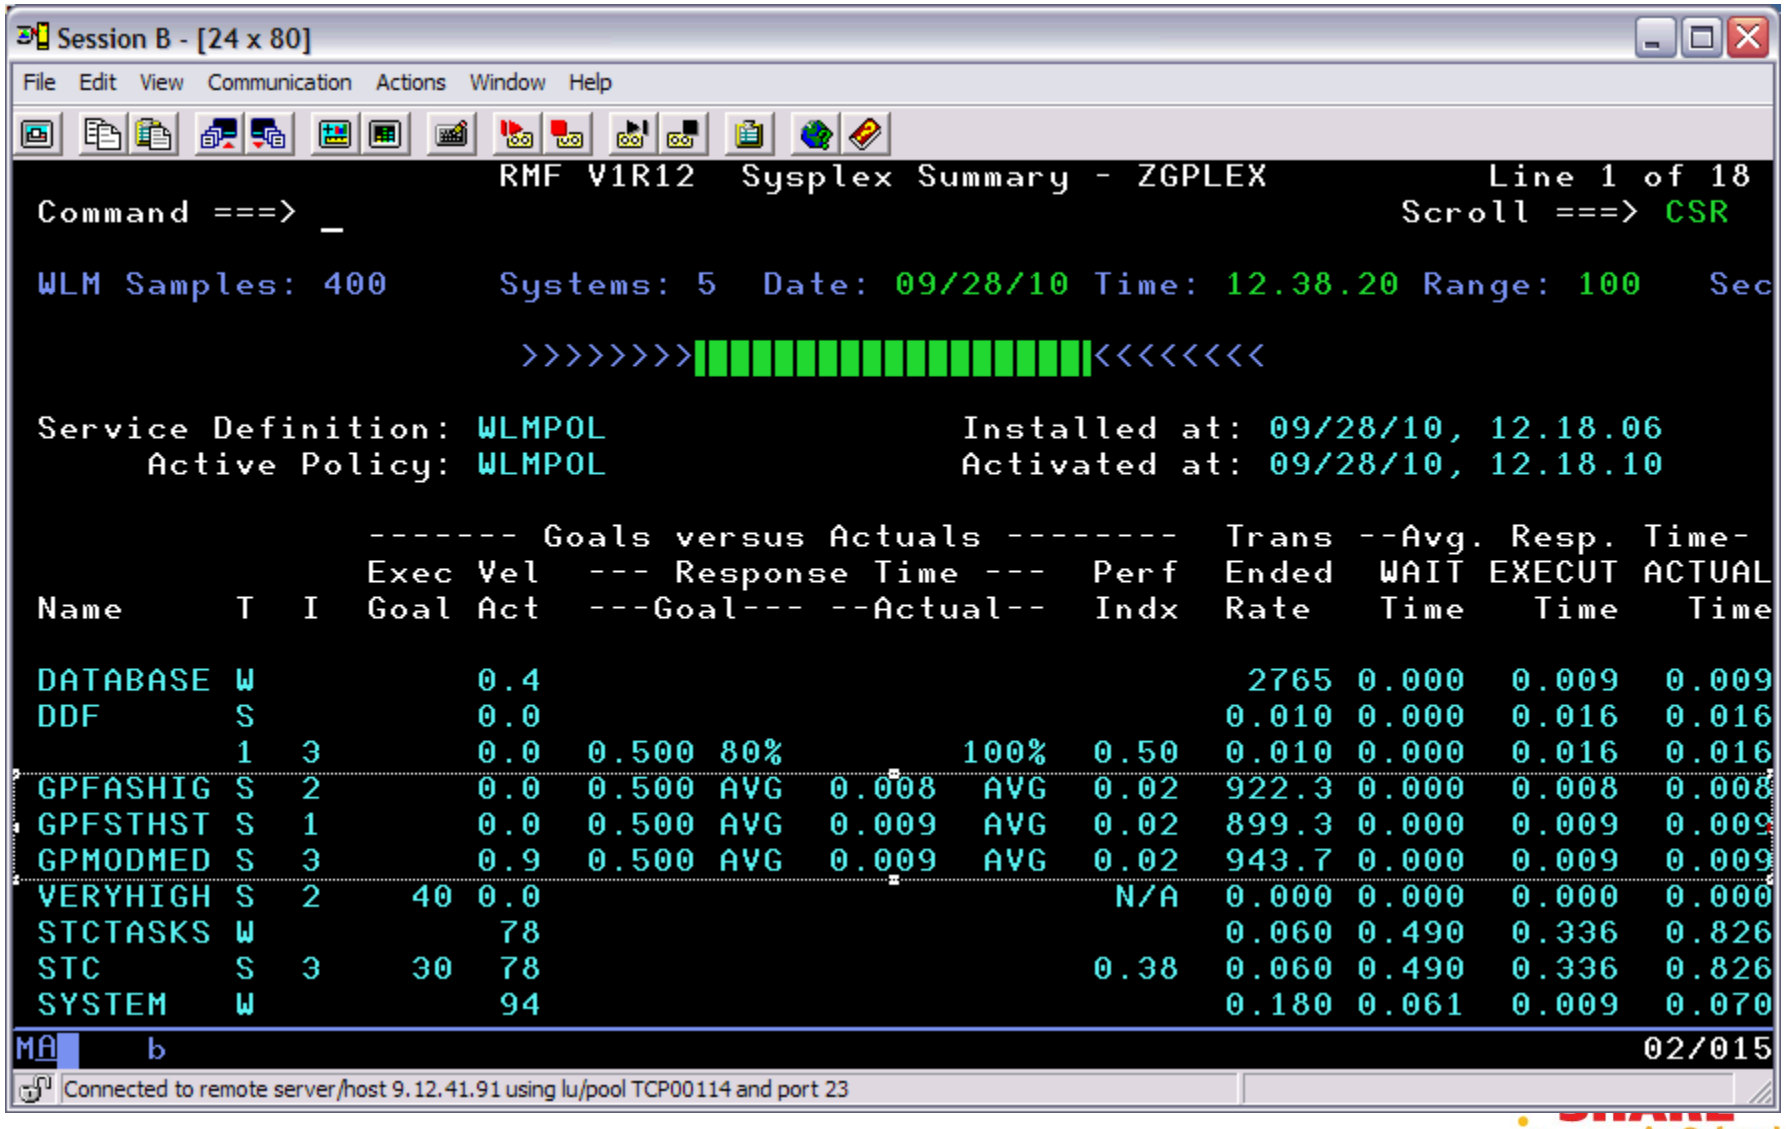

Platform Performance Manager Summary

- 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

Thank You