zEnterprise eXposed! Experiences with zEnterprise Unified Resource Manager

Session 11603

Brad Snyder

Mary Astley

Advanced Technical Skills
IBM Corporation

Trademarks

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

AIX*	POWER*	System z
BladeCenter*	POWER7*	System z10
CICS*	PowerVM	VTAM*
DB2*	RACF*	WebSphere*
Datapower*	Redbooks*	z/Architecture*
HiperSockets	RMF	z/OS*
IBM*	Rational	z/VM*
MVS	System p*	zEnterprise
Parallel Sysplex*	System x*	zSeries*
PR/SM		
Processor Resource/Systems Manager		

* Registered trademarks of IBM Corporation

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.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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.

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.

First Half Agenda

zEnterprise Ensemble Overview

Platform Performance Management

- Role
- Guest Platform Management Provider
- Application Response Measurement

z/OS and Unified Resource Manager Workload Policies

- z/OS Workload Manager Policy
- Unified Resource Manager Workload

Classification of zEnterprise Work in z/OS Workload Manager

- Service classes for zEnterprise work
- Example

IBM zEnterprise Ensemble

Ensemble

- Up to eight nodes
- Each z196 or z114 CPC is a node
- Node may optionally have attached zBX

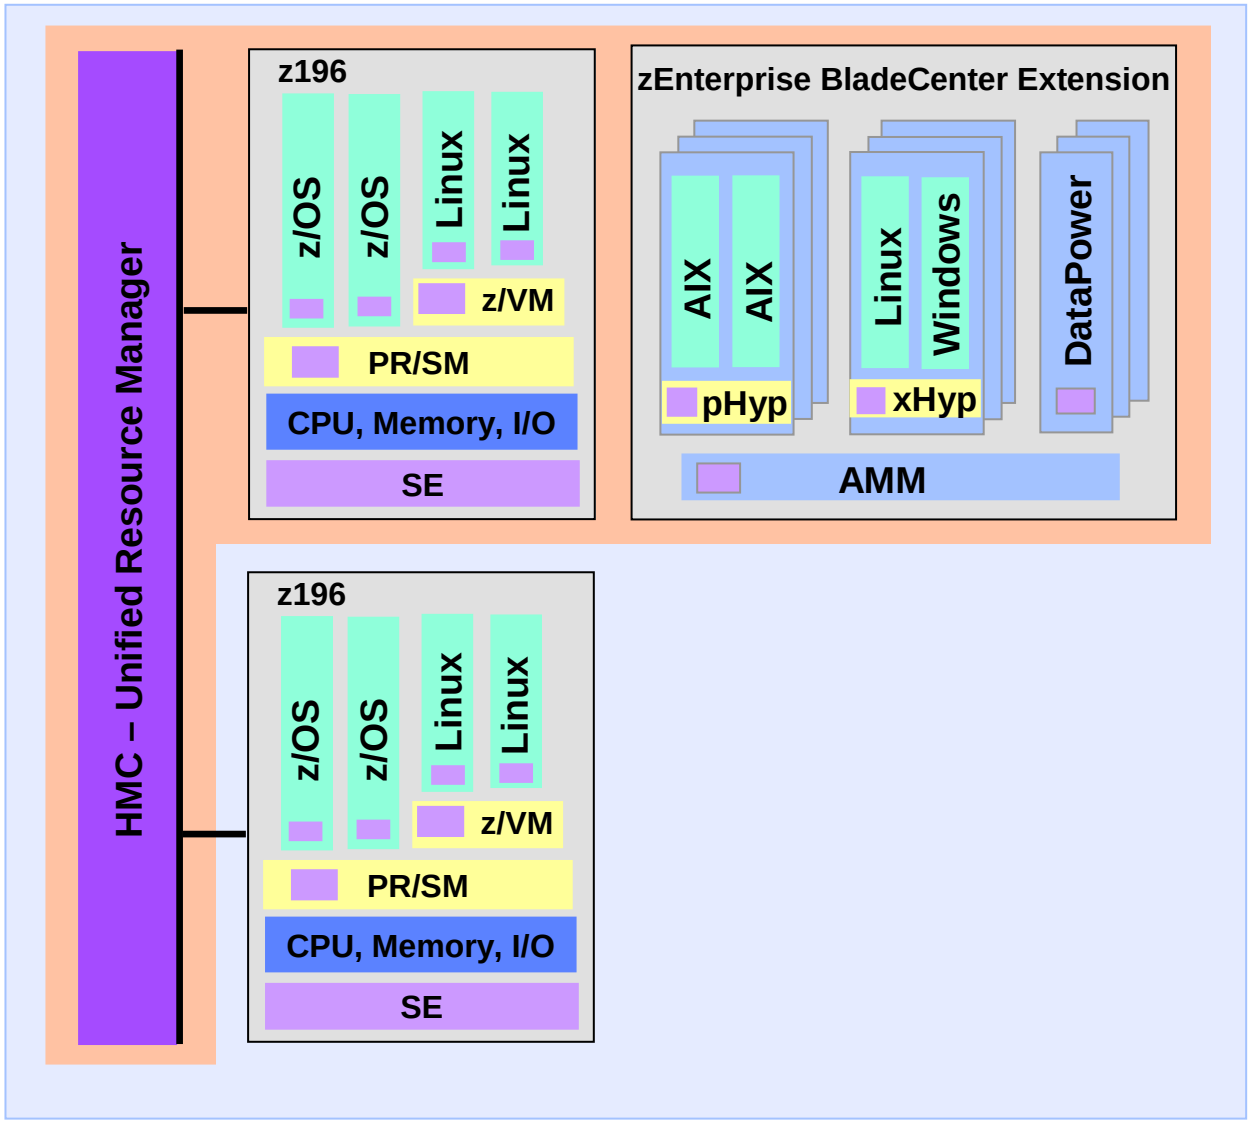

Unified Resource Manager (zManager) Platform Performance Management

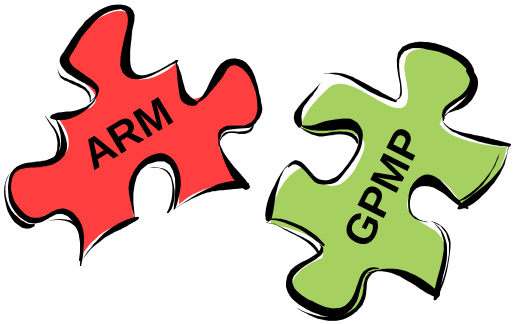

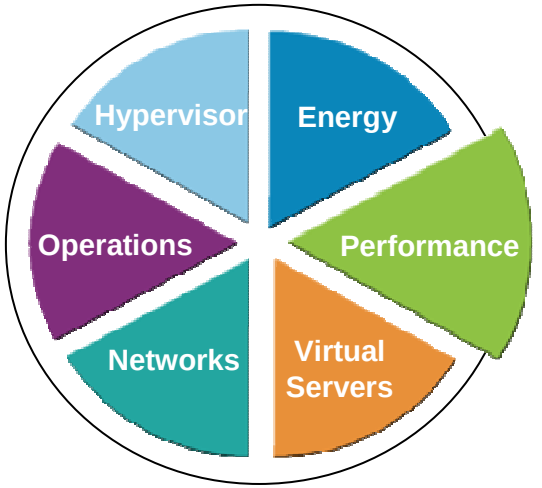

Platform Performance Management

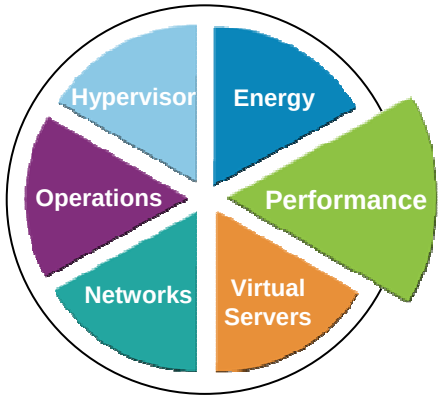

- zManager component responsible for monitoring, reporting, and management of resources used by virtual servers
- Scope is the ensemble
- User interface is the ensemble HMC
- Workload goals specified in workload performance policy
- Based on goals in workload policy, PPM may adjust processor resources across virtual servers running under the same hypervisor - AIX and z/VM virtual servers
- Reports available showing virtual server resource usage
- Optional Guest Platform Management Provider

Guest Platform Management Provider

GPMP – Link between operating system and zManager

- Lightweight component of PPM
- Collects performance data for work running on a virtual server and passes it to zManager
- User installed on operating systems
- GPMP provides additional monitoring data
- Option of mapping zManager workload service classes to z/OS WLM service classes
- With Application Response Measurment (ARM) instrumented middleware support, GPMP provides
  - End-to-end transaction response times
  - Detailed transaction topology

ARM enabled middleware

- Web Server support provided by WebSphere plugin IHS/Apache, IIS, Domino, iPlanet
- WebSphere Application Server V6, V7, V8
- DB2 Universal Database – including z/OS DB2

Application Response Measurement

Standards based application instrumentation

Method to monitor the performance and availability across multiple servers in a distributed workload

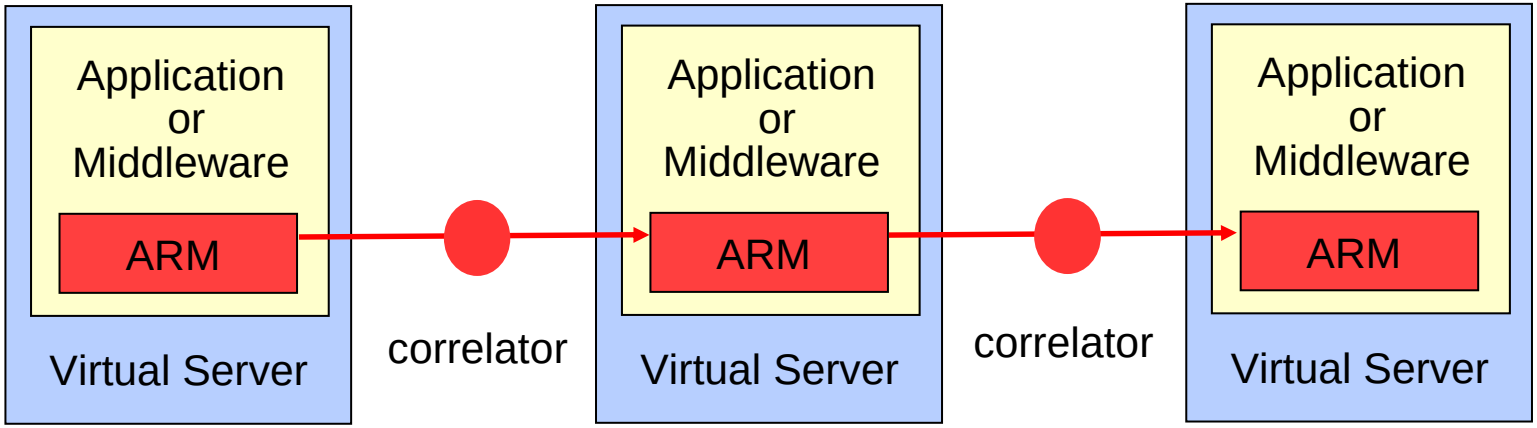

Information in the correlator used by zManager to report

- Name of the applications, middleware, and servers processing transaction
- End-to-end transaction response time
- Time spent in each “HOP”

All applications and servers processing work request must be ARM enabled

zManager PPM Workload

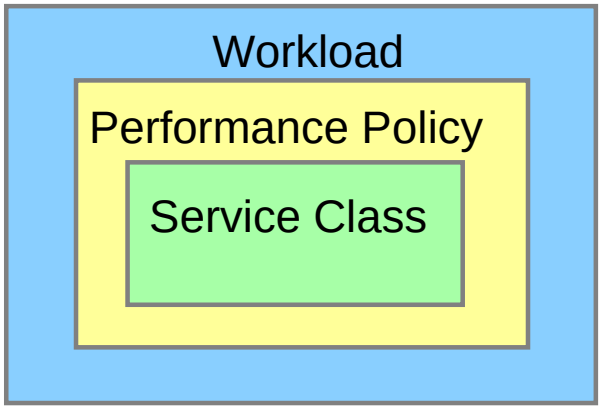

WLM and PPM Terminology

z/OS WLM and zManager PPM terms may have different meanings

- z/OS Workload Manager
  - WLM has a service definition
  - Service definition has an active policy
    - Policy contains all workloads and service classes
    - Definition may have multiple policies, only one can be active
  - Incoming work is classified into a service class
- zManager PPM
  - Workload
    - Default Workload is provided
    - Custom Workloads may be defined
  - Virtual servers are assigned to a Workload
  - Workload has an active performance policy
    - Policy contains service classes
    - Workload may have multiple policies, only one can be active
  - Virtual servers are assigned to a service class

WLM and PPM Policy

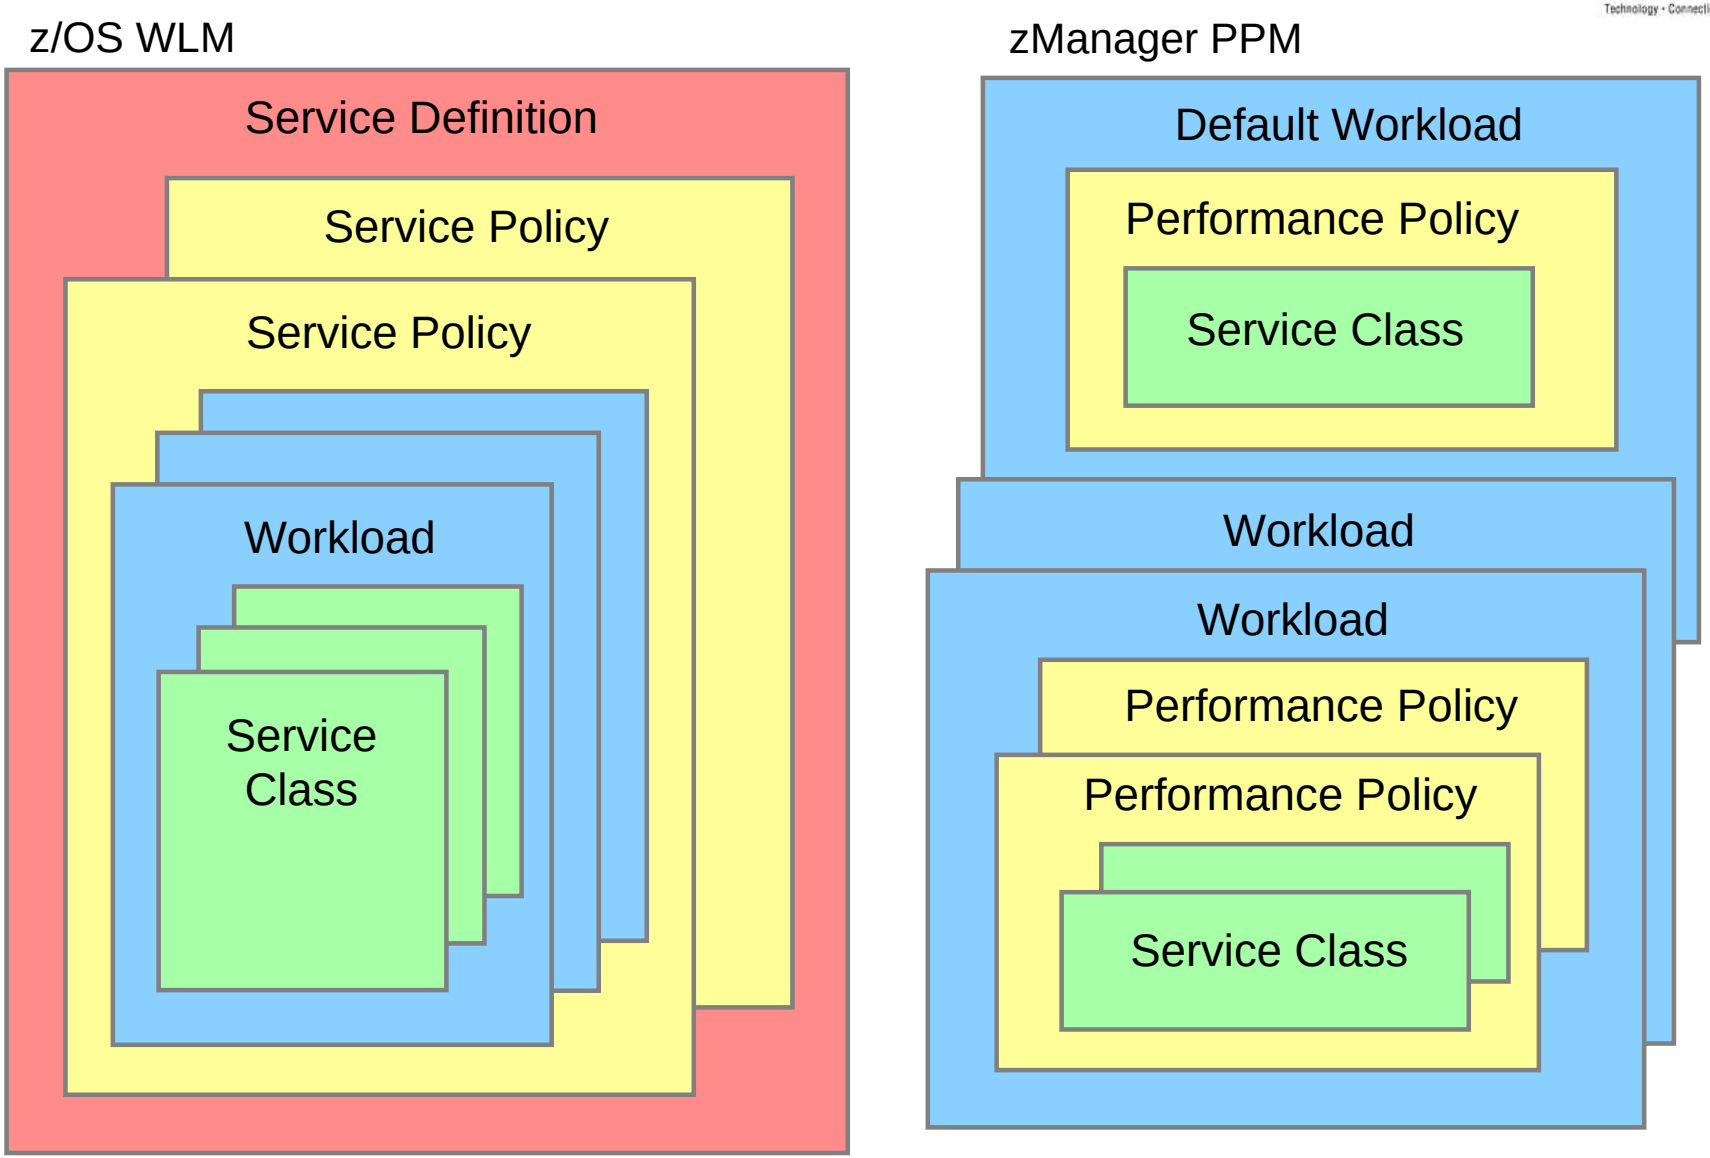

zManager Workload

A workload is a customer defined collection of virtual servers

- Provides a way to group virtual servers to manage and monitor performance for a business application
- Has one or more performance policies
- Specify workload importance and goals in performance policy

Virtual servers in a workload

- Must be in the same ensemble
- Can be on different blades or nodes
- May reside in more than one workload
- Virtual servers not assigned to a custom workload are in default workload

Workload Performance Policy

Performance policy defines performance goals for virtual servers in the workload

Each performance policy has

- Name
- Business importance:
  - Highest, High, Medium, Low, or Lowest
- One or more service classes

Multiple policies may be defined; Only one policy can be active

Active policy may be changed dynamically

- Through zManager on HMC
- With a time-based schedule

Service Classes

Each service class has

- Name
- Performance goal type
  - Discretionary
  - Velocity – Fastest, Fast, Moderate, Slow, Slowest
- Business importance for velocity goal
  - Highest, High, Medium, Low, or Lowest
- Classification rules to assign incoming work to service class

zManager will assign a PI for every service class

- PI = 1.0 – Service Class Achieving Goal
- PI > 1.0 – Service Class Missing Goal
- PI < 1.0 – Service Class Overachieving Goal

Service Class Grid

Imp/Velocity	Fastest - 1	Fast - 2	Moderate - 3	Slow - 4	Slowest - 5
Highest - 1	Group11	Group12	Group13	Group14	Group15
High - 2	Group21	Group22	Group23	Group24	Group25
Medium - 3	Group31	Group32	Group33	Group34	Group35
Low - 4	Group41	Group42	Group43	Group44	Group45
Lowest - 5	Group51	Group52	Group53	Group54	Group55

PPM Workload Planning

Steps for planning PPM Workloads

1. Determine workloads for applications
2. Assign virtual servers to each workload
3. Define a performance policy and assign importance
4. Determine number of service classes
   Velocity goal and importance for each service class
5. Assign virtual servers to each service class
6. Create the classification rules to assign virtual servers to the service classes

An Example: New Enterprise Application

A new banking application to be run in the zEnterprise ensemble

The sizing work is finished and the application will be running on the following virtual servers

Ensemble Node	BladeCenter & Blade Number	Virtual Server	Hypervisor	Function
Node1	B.1.01	WSCAIX11	PowerVM	AppServer
	B.1.01	WSCAIXT2	PowerVM	AppServer
	B.1.02	WSCAIX12	PowerVM	AppServer
	B.1.02	WSCAIXT2	PowerVM	AppServer
	B.2.01	WSCLNX21	X Hyp	HTTP Server
	B.2.01	WSCLNXT1	X Hyp	HTTP Server
	B.2.02	WSCLNX22	X Hyp	HTTP Server
	B.2.02	WSCLNXT2	X Hyp	HTTP Server
		LPAR5	PR/SM	DB2

Example: Define Workload

Step 1: Define the workload for the new banking application

The workload name will be Banking

Step 2: Assign the virtual servers to the workload

Virtual servers are
WSCAIX11, WSCAIX12, WSCAIXT1, WSCAIXT2
WSCLNX21, WSCLNX22, WSCLNXT1, WSCLNXT2
LPAR5

Step 3: Define a performance policy and assign importance

One performance policy – workload priority same all shifts
Business importance of workload is High
Performance policy name is Standard

Step 4: Determine number of service classes, names, velocity goals, and business importance

Servers running on the same blade compete for resources

Importance and velocity goals are ONLY relevant within the boundaries of one blade. This allows for limited

PPM Workload Service Class Names

Performance policy service class names

- Must start with alphabetic character
- Mixed case (case sensitive)
- Up to 64 characters in length
- If name is used by z/OS WLM to classify work, maximum length is 32 characters

Assigning Common Service Class Names

A method for assigning service class names is to incorporate the workload performance policy business importance and the service class business importance into the service class name

The service class name includes the numbers representing the policy and service class importance levels

Imp/Velocity	Fastest - 1	Fast - 2	Moderate - 3	Slow - 4	Slowest - 5
Highest - 1	Group11	Group12	Group13	Group14	Group15
High - 2	Group21	Group22	Group23	Group24	Group25
Medium - 3	Group31	Group32	Group33	Group34	Group35
Low - 4	Group41	Group42	Group43	Group44	Group45
Lowest - 5	Group51	Group52	Group53	Group54	Group55

Example: Service Classes & Virtual Servers

Step 4: Determine number of service classes, names, velocity goals, and business importance

Step 5: Define service classes and assign virtual servers

Service Class	Velocity	Importance	Virtual Servers
GROUP11Banking	Fastest	Highest	WSCAIX11, WSCAIX12 WSCLNX21, WSCLNX22
GROUP44Banking	Slow	Low	WSCAIXT1, WSCAIXT2 WSCLNXT1, WSCLNXT2

Example: Classification Rules

Step 6: Create the classification rules to assign virtual servers to the service classes

Service Class	Classification Rule	Virtual Servers
GROUP11Banking	Virtual Server Name Equals	WSCAIX11 or WSCAIX12 WSCLNX21 or WSCLNX22
GROUP44Banking	Virtual Server Name Equals	WSCAIXT* or WSCLNXT*
???	ESC / DDF	LPAR5

PPM does not classify work in z/OS WLM
Service class for z/OS work is assigned by WLM

z/OS WLM and zManager Workloads

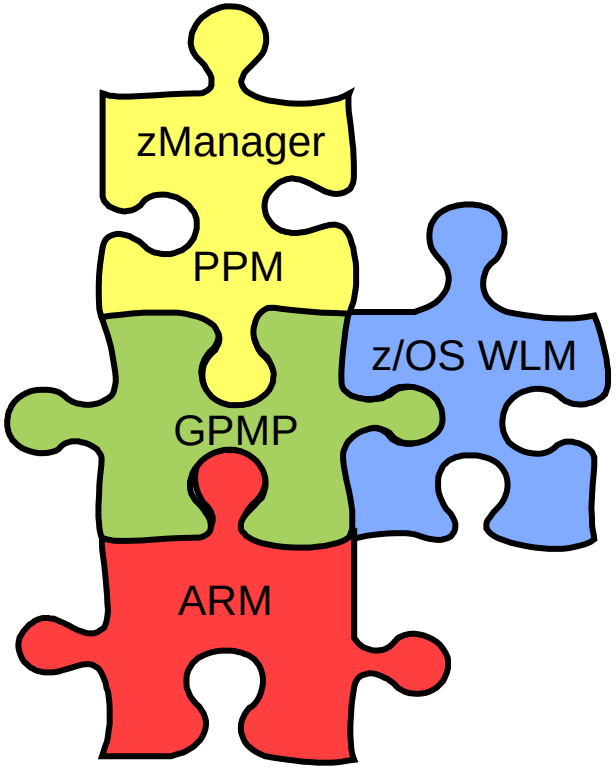

z/OS WLM Classification

PPM performance policy manages the velocity of the virtual servers, it does not manage the work running on z/OS

Work coming into z/OS from the zEnterprise can be classified using EWLM rules

- Requires GPMP and ARM to be active
- Uses name of PPM service class for classification

If ARM is not enabled, applications are not ARM enabled, or no EWLM classification rules apply

- Standard WLM classification rules will apply (DDF, JES, etc.)

End-to-end goal-based performance management

EWLM Subsystem Type

EWLM subsystem type allows WLM service or report classes to be assigned to EWLM work arriving from zEnterprise

ESC (EWLM service class) is the only Qualifier type

Qualifier name is the performance policy service class name

- Sub-rules used when service class name is longer than 8 characters
- WLM allows up to 4 sub-rules
- z/OS WLM will use the first 32 characters

A default service class can be specified

EWLM service classes must be single period with response time goal

EWLM Classification

```
                Subsystem Type Selection List for Rules     Row 1 to 12
Command ===> ____________________________________________________________

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

                                                      ----- Class -----
Action  Type      Description                         Service   Report
  __    ASCH      APPC scheduled trans programs       _________ _________
  __    CB        Component Broker requests           _________ _________
  __    CICS      CICS transaction level rules        _________ _________
  __    DB2       DB2 Parallel Query transactions     _________ _________
  __    DDF       Distributed DDF work                DDFDEF    _________
  __    EWLM      EWLM Rules for PPM                  PPMDEFLT  _________
  __    IMS       IMS transaction level rules         _________ _________
  __    IWEB      Scalable WebServer Transactions     _________ _________
  __    JES       JES classification rules            BAT_MED   RBAT_MED
  __    LSFM      Lan Server for MVS rules            _________ _________
  __    MQ        MQ Series Workflow requests         _________ _________
  __    OMVS      Unix System Services requests       UNIX      _________
```

EWLM Rules


```
          Modify Rules for the Subsystem Type                Row 1 to 9
Command ===> _____________________________________________________

Subsystem Type . : EWLM           Fold qualifier names?   N  (Y or N)
Description  . . . EWLM Rules for PPM

Action codes:   A=After     C=Copy        M=Move     I=Insert rule
                B=Before    D=Delete row  R=Repeat   IS=Insert Sub-rule
                                                        More ===>
          --------Qualifier--------                 -------Class--------
Action    Type       Name     Start                  Service     Report
                                              DEFAULTS: PPMDEFLT    _________
 _____  1  ESC        GROUP11B 1                      _________    _________
 _____  2    ESC        anking   9                    PPMHGHST    _________
 _____  1  ESC        GROUP44B 1                      _________    _________
 _____  2    ESC        anking   9                    PPMTEST_    _________
 _____  1  ESC        Default  1                      PPMDF2      _________
 _____  1  ESC        GROUP21  1                      BANK2       _________
 _____  1  ESC        GROUP44  1                      GRP44TST    _________
```

Putting the Pieces Together

PPM Performance Policy is only managing the velocity of the virtual servers

GPMP reports operating system performance information to zManager

ARM provides information only for overall transaction response times

WLM manages work on z/OS – ARM token used to clasify work only

If transactions are running too long and CPU on a blade is the constraint

- Velocity goals of individual servers may need to be adjusted
- Number of virtual processors increased

Summary

- Platform Performance Manager function of zManager provides tools to manage and monitor workloads running in ensemble
- ARM provides a method for gathering transaction information and reporting application performance
- z/OS Workload Manager classifies and manages the work running on z/OS
- Combination of PPM, ARM, and z/OS WLM allow for excellent reporting of response time and performance of applications running on zEnterprise

Second Half Agenda

Defining a Workload

Enabling GPMP

Reports

- Workloads report
- Virtual servers report
- Hops report
- Topology report

Defining a Workload in an Ensemble

New Workload

Define using the New Workload Wizard on the ensemble HMC

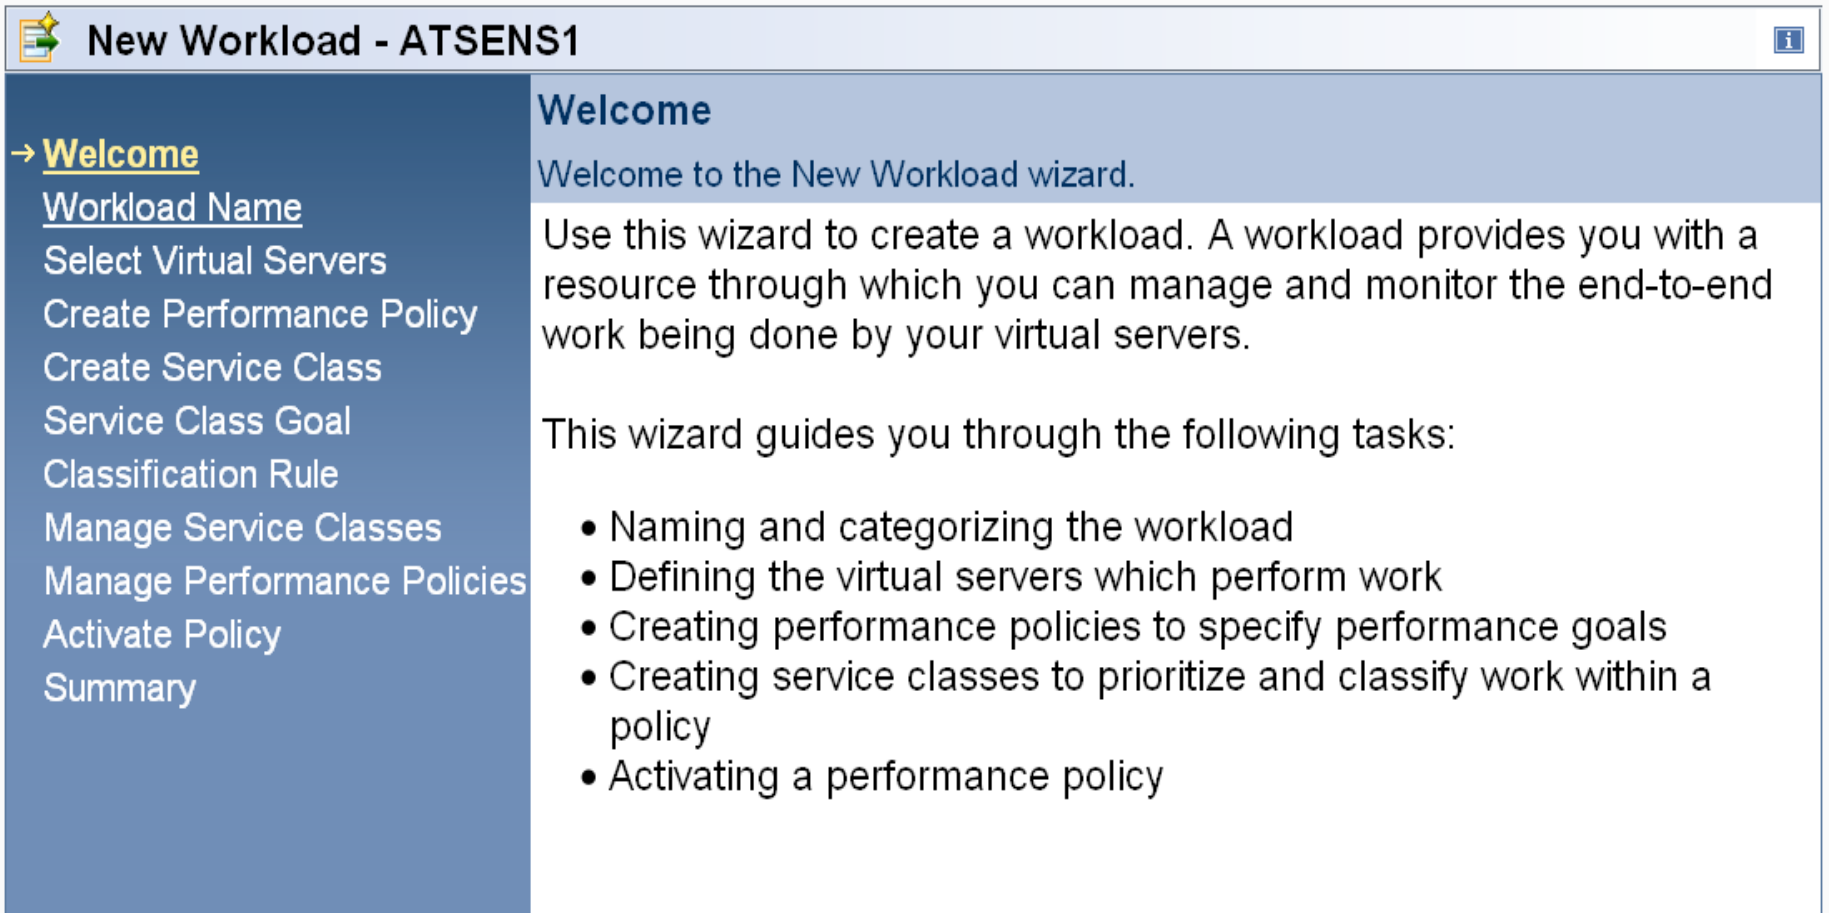

Workload Name

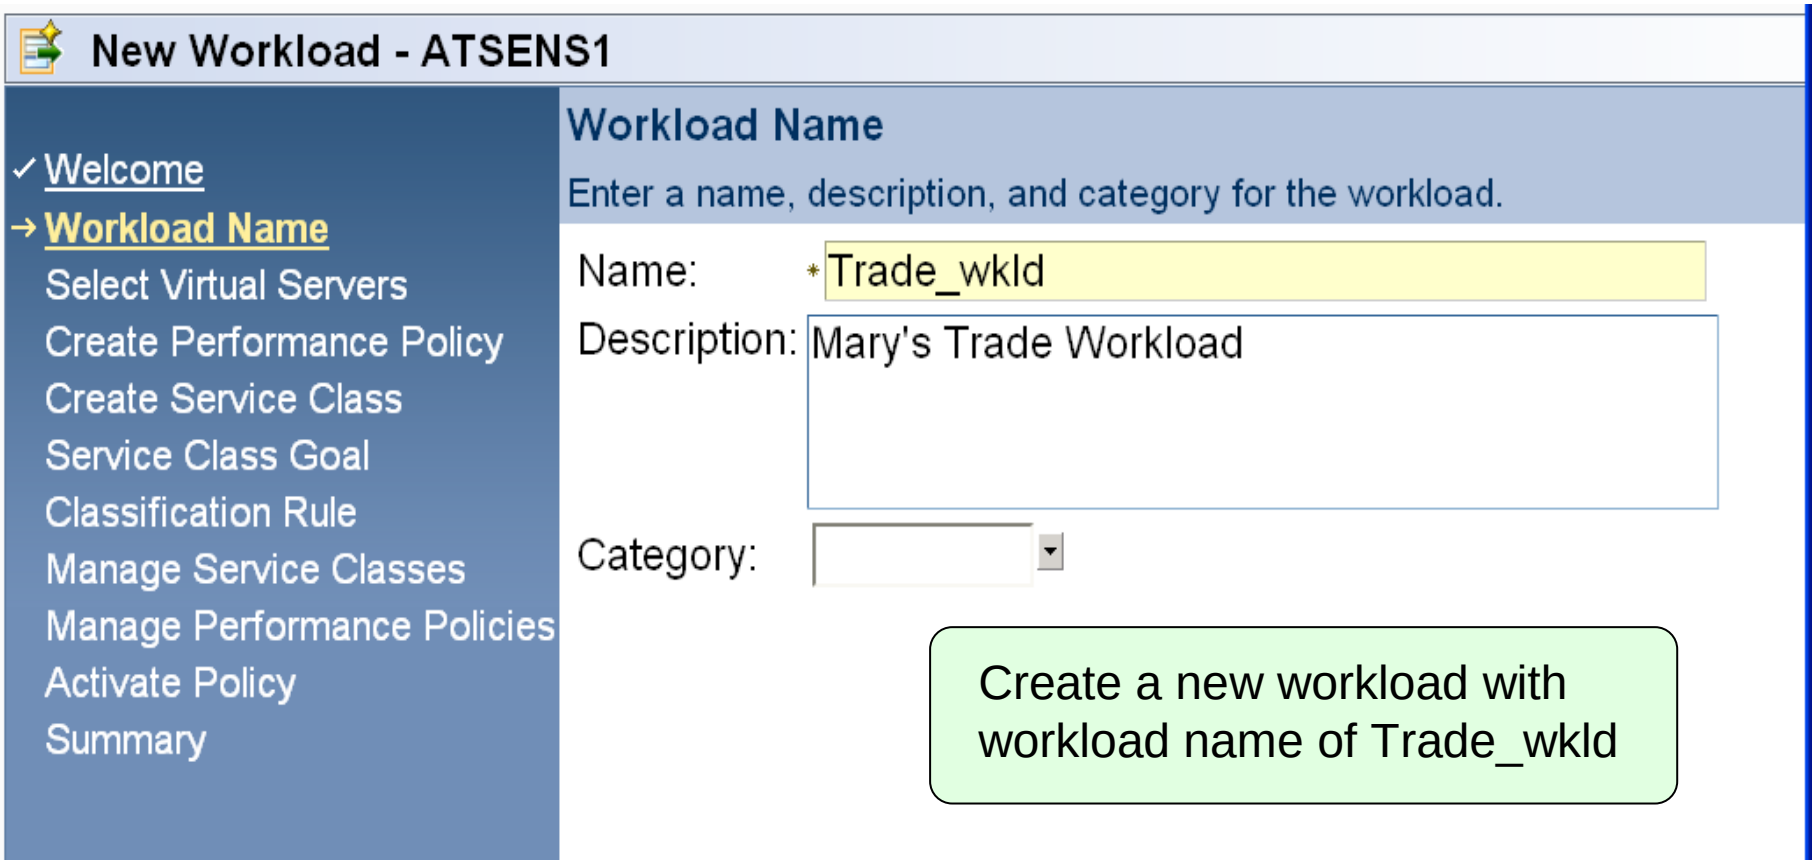

Workload Virtual Servers

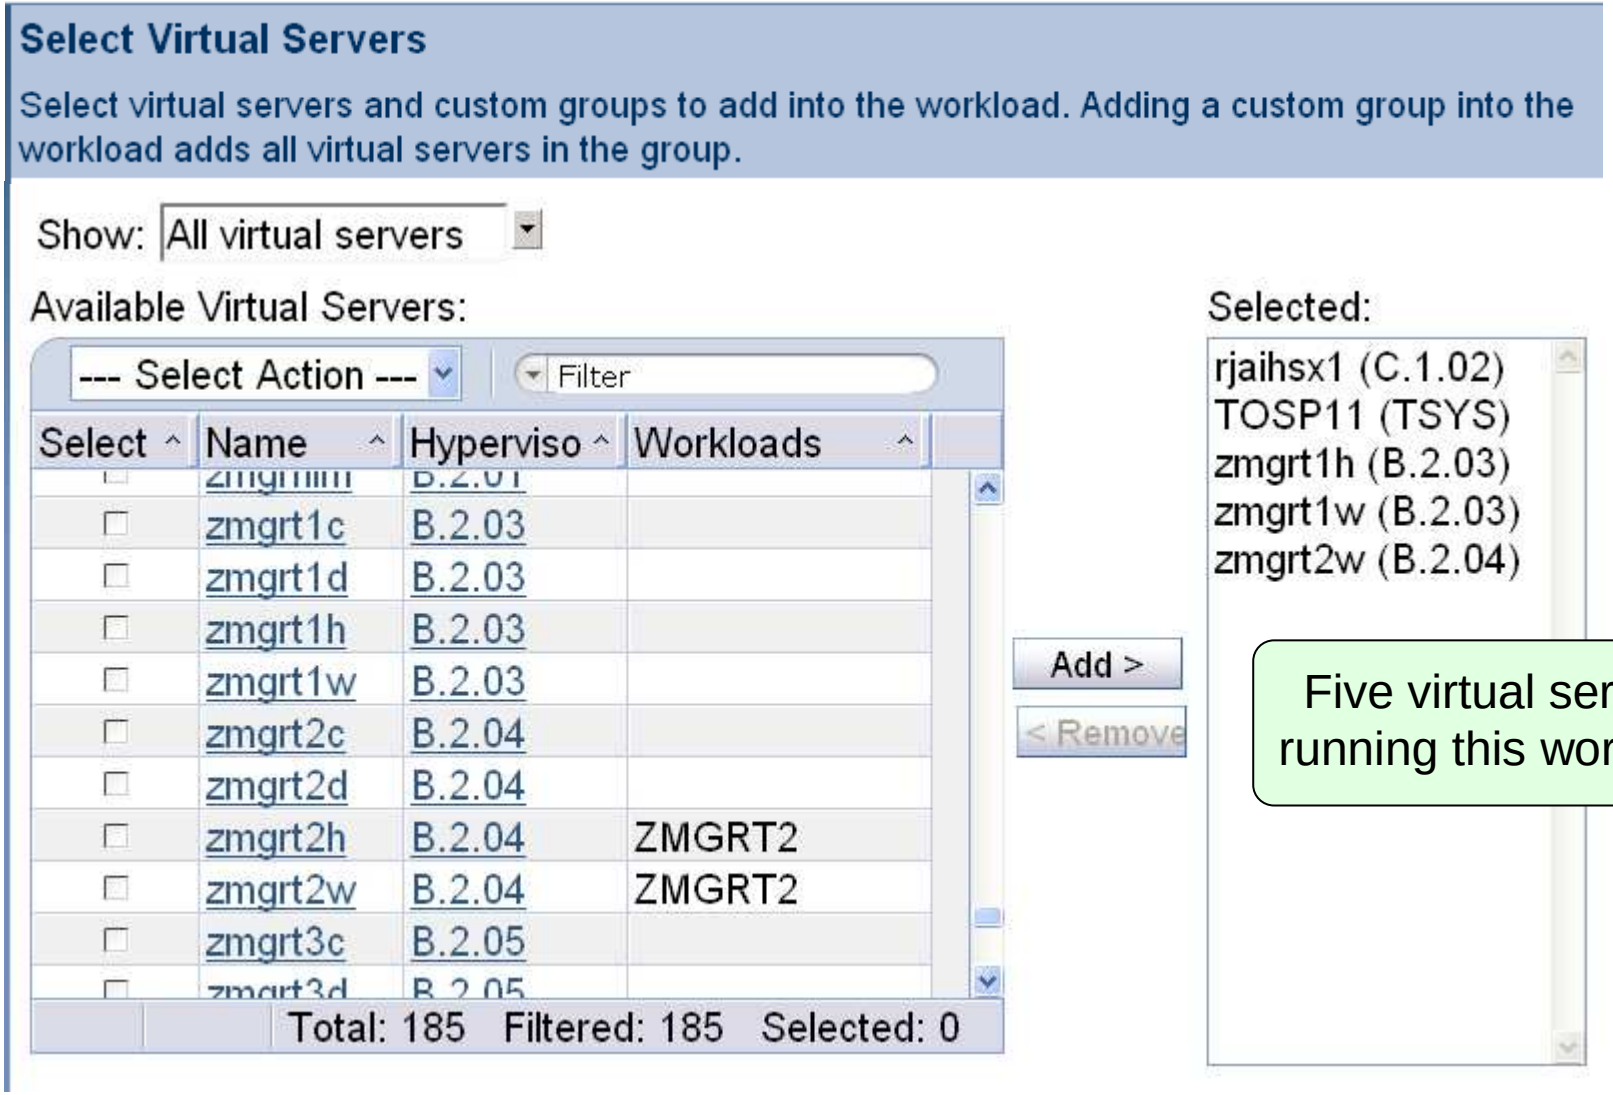

Workload Performance Policy

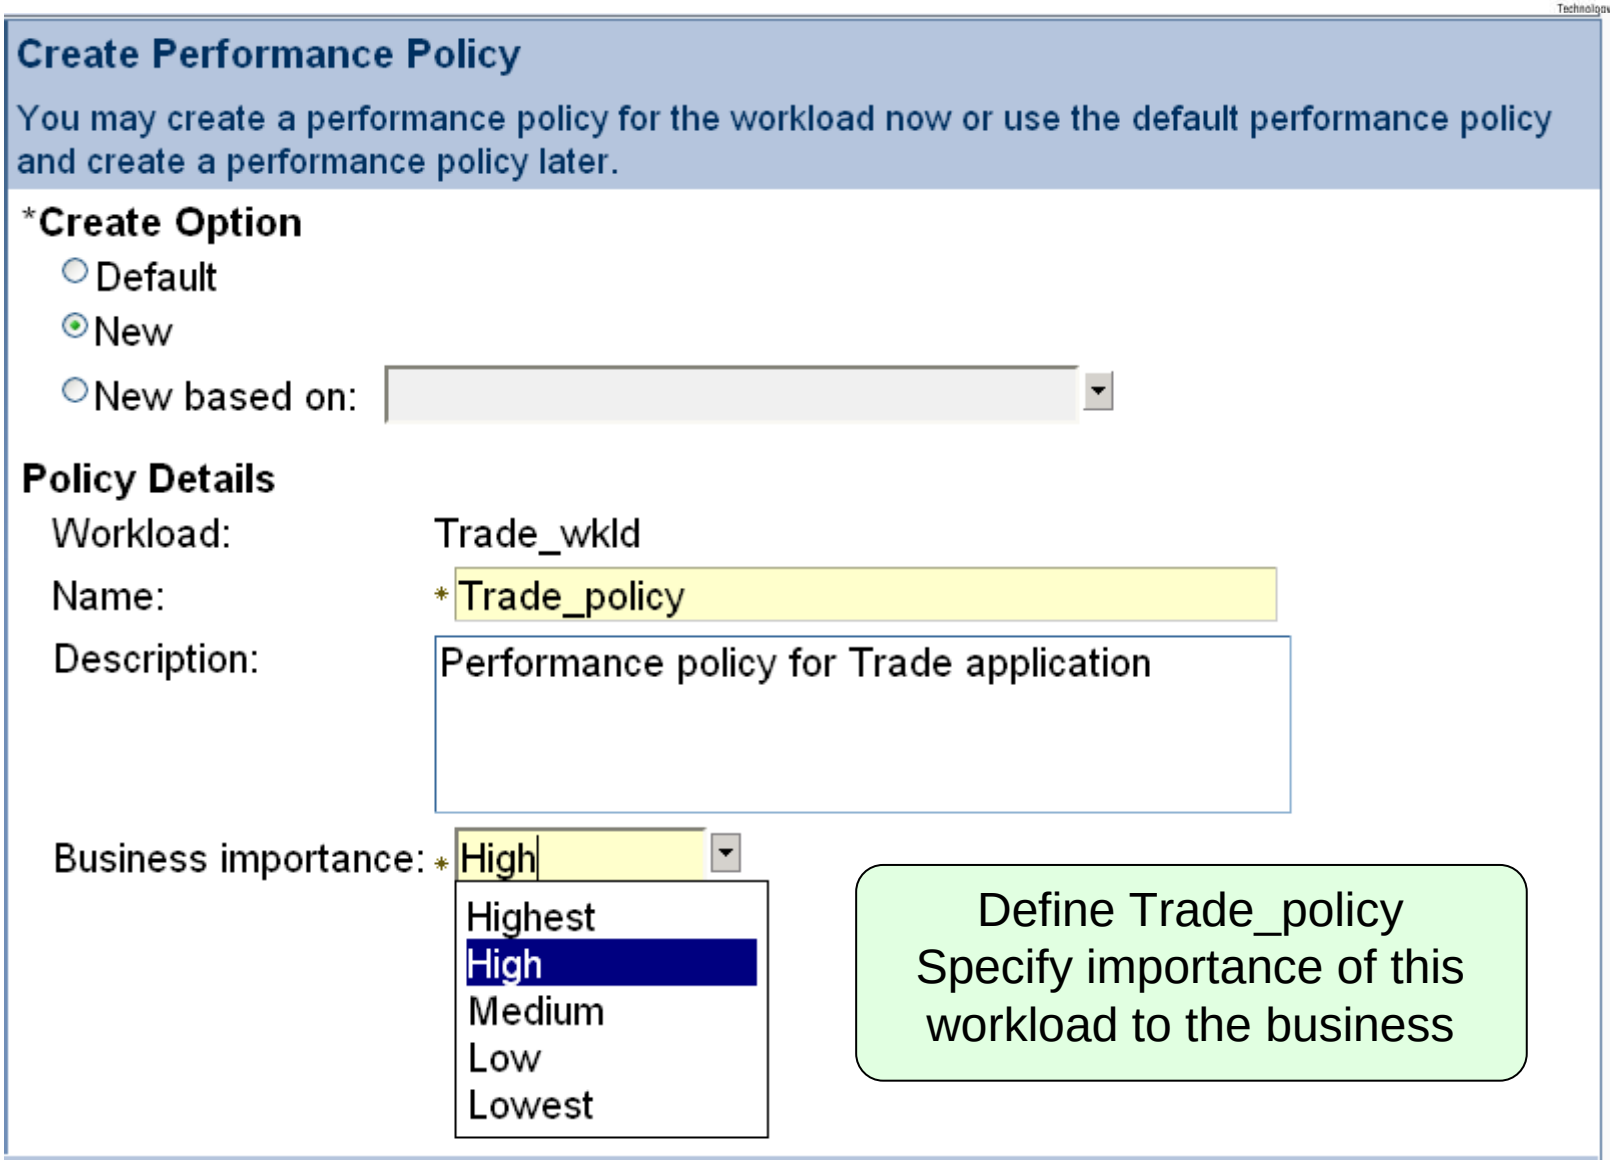

Workload Service Class

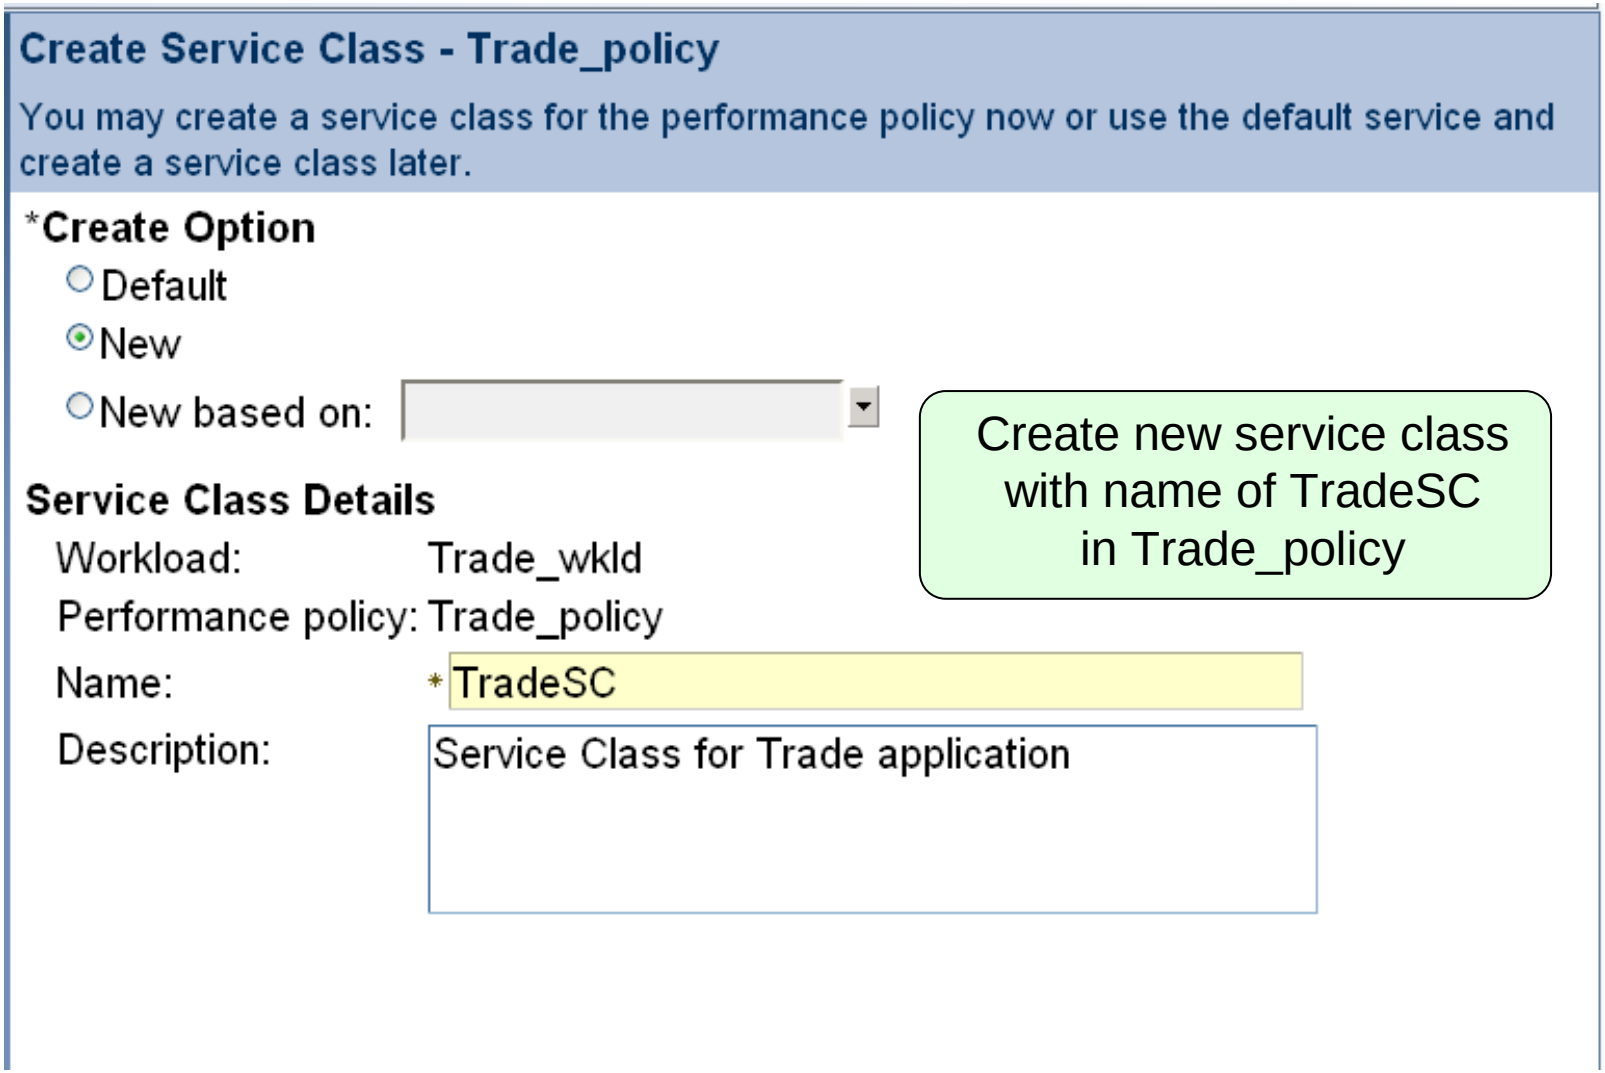

Workload SC Performance Goal

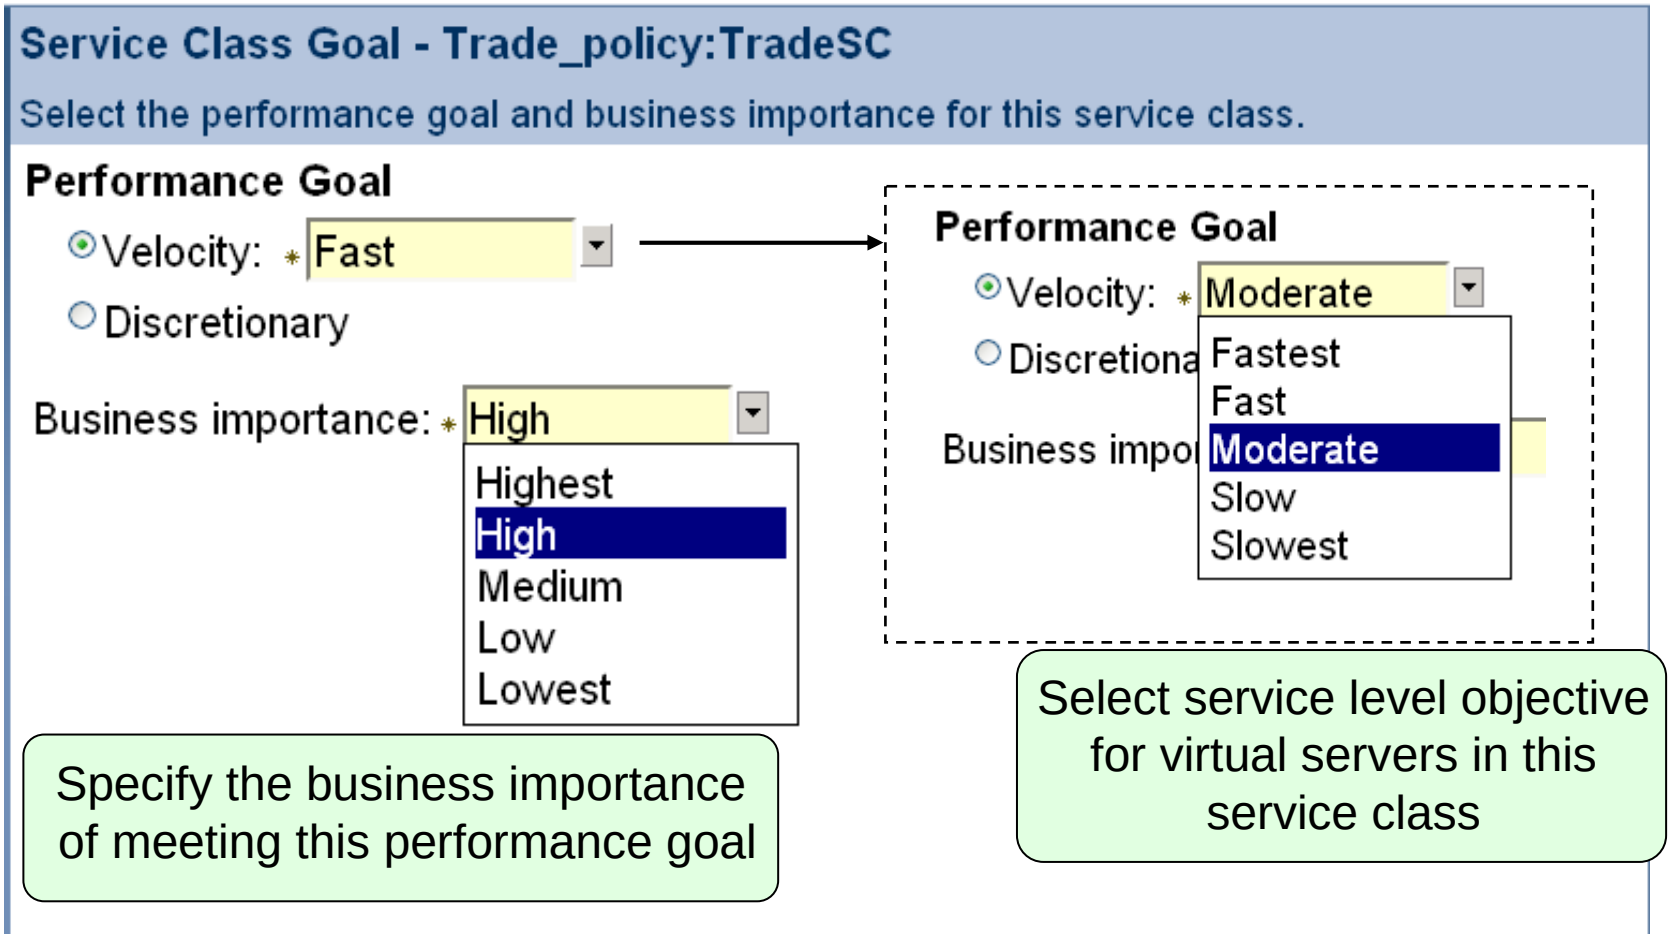

Workload Classification Rule

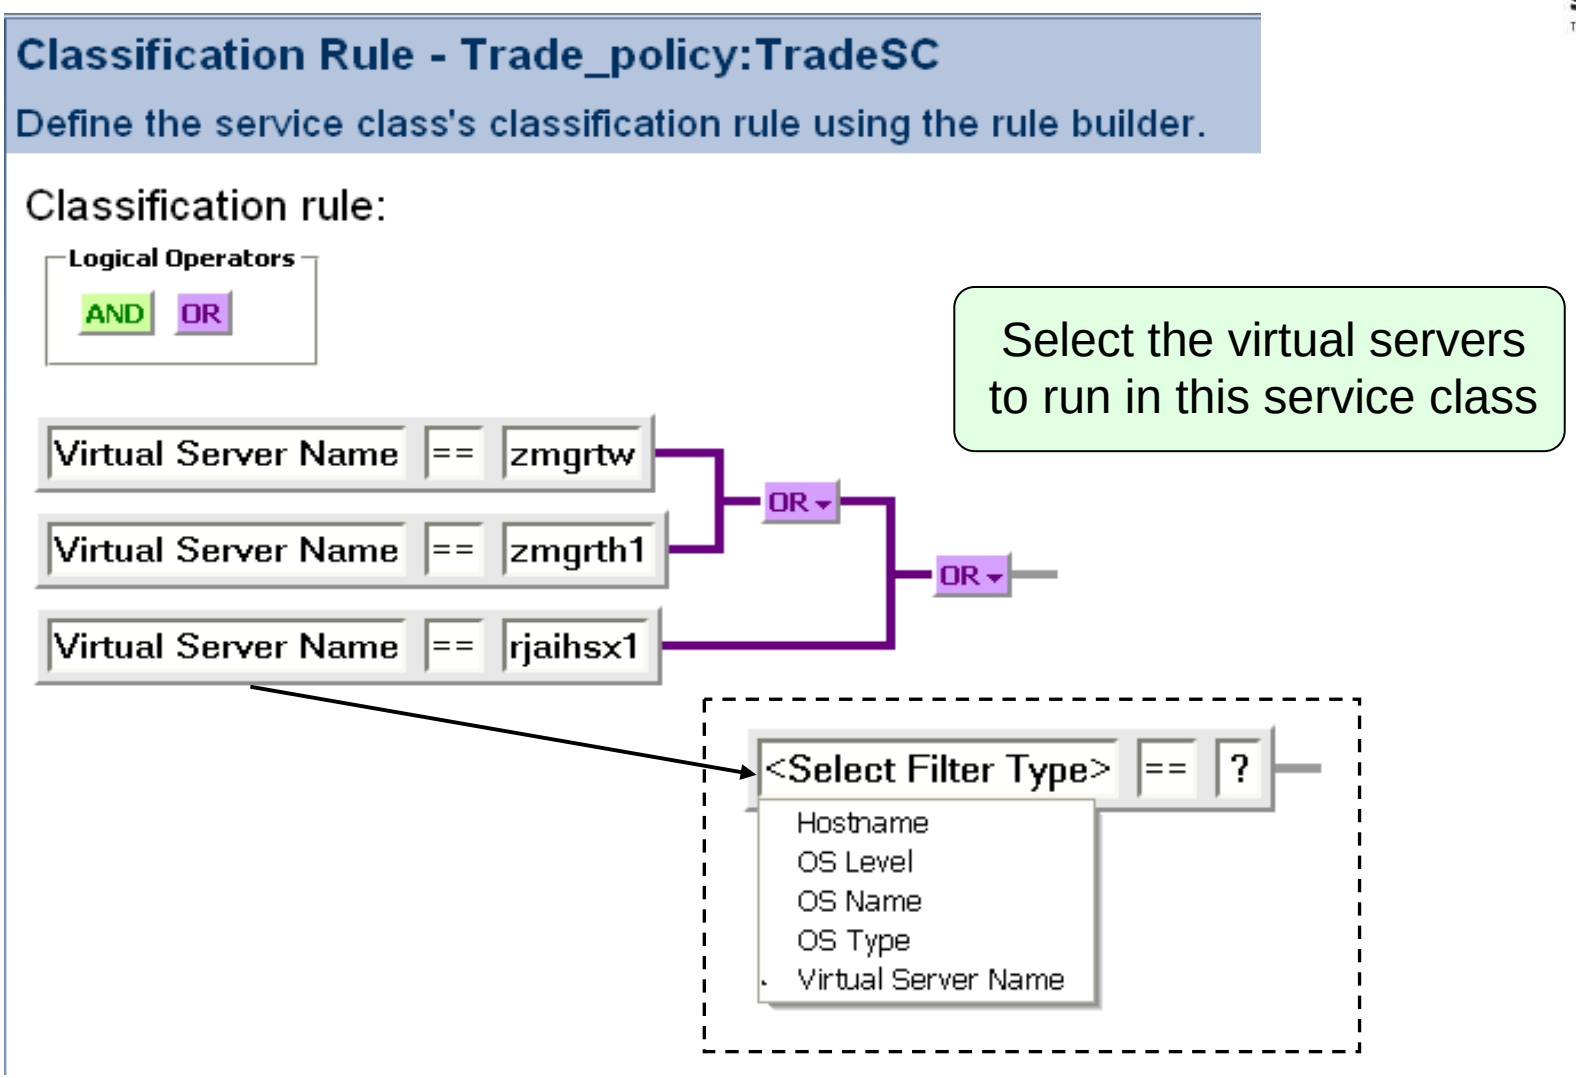

Workload Summary

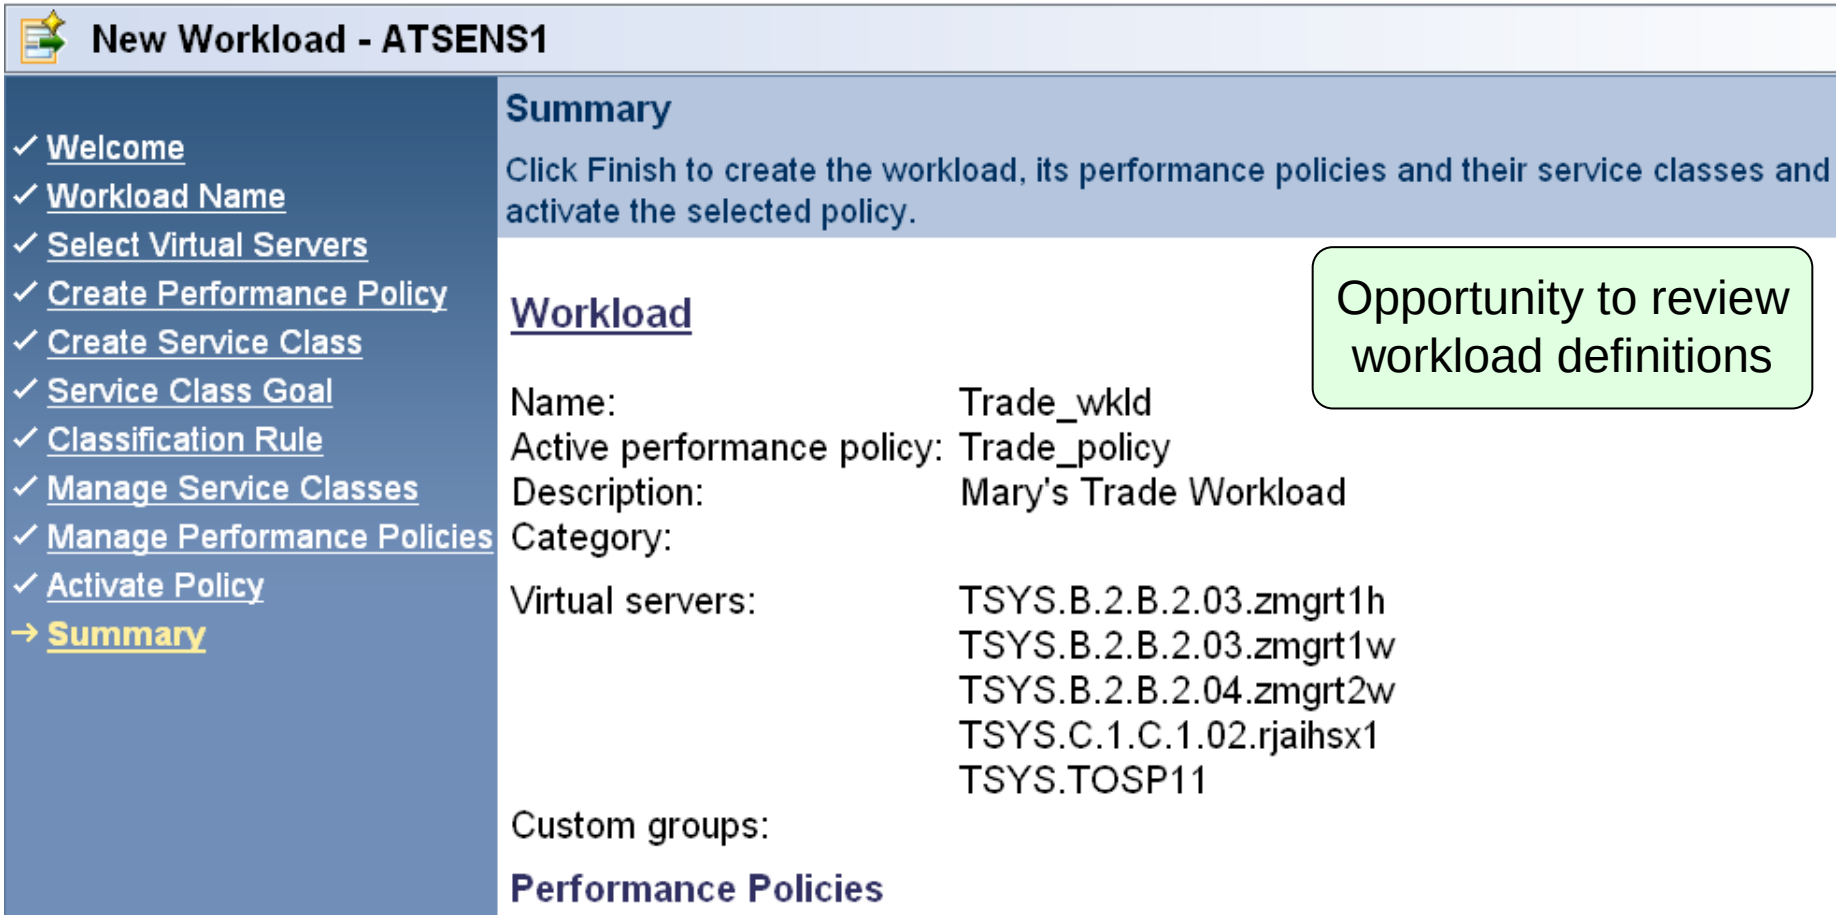

Opportunity to review workload definitions

Workload Summary – Default Policy

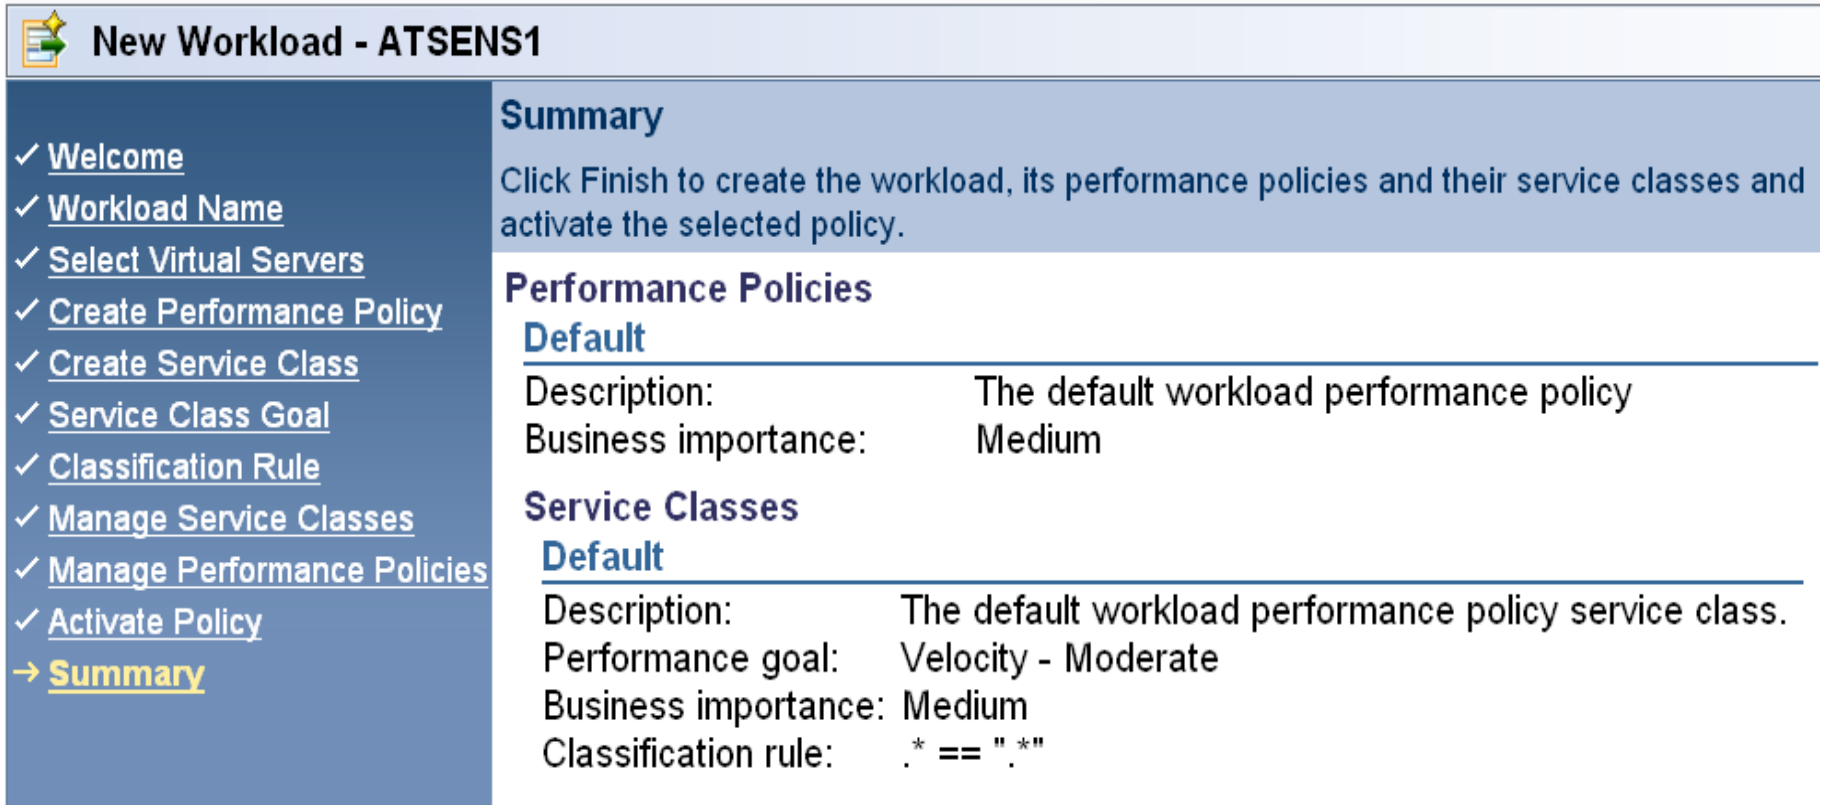

Workload Summary – Trade_Policy

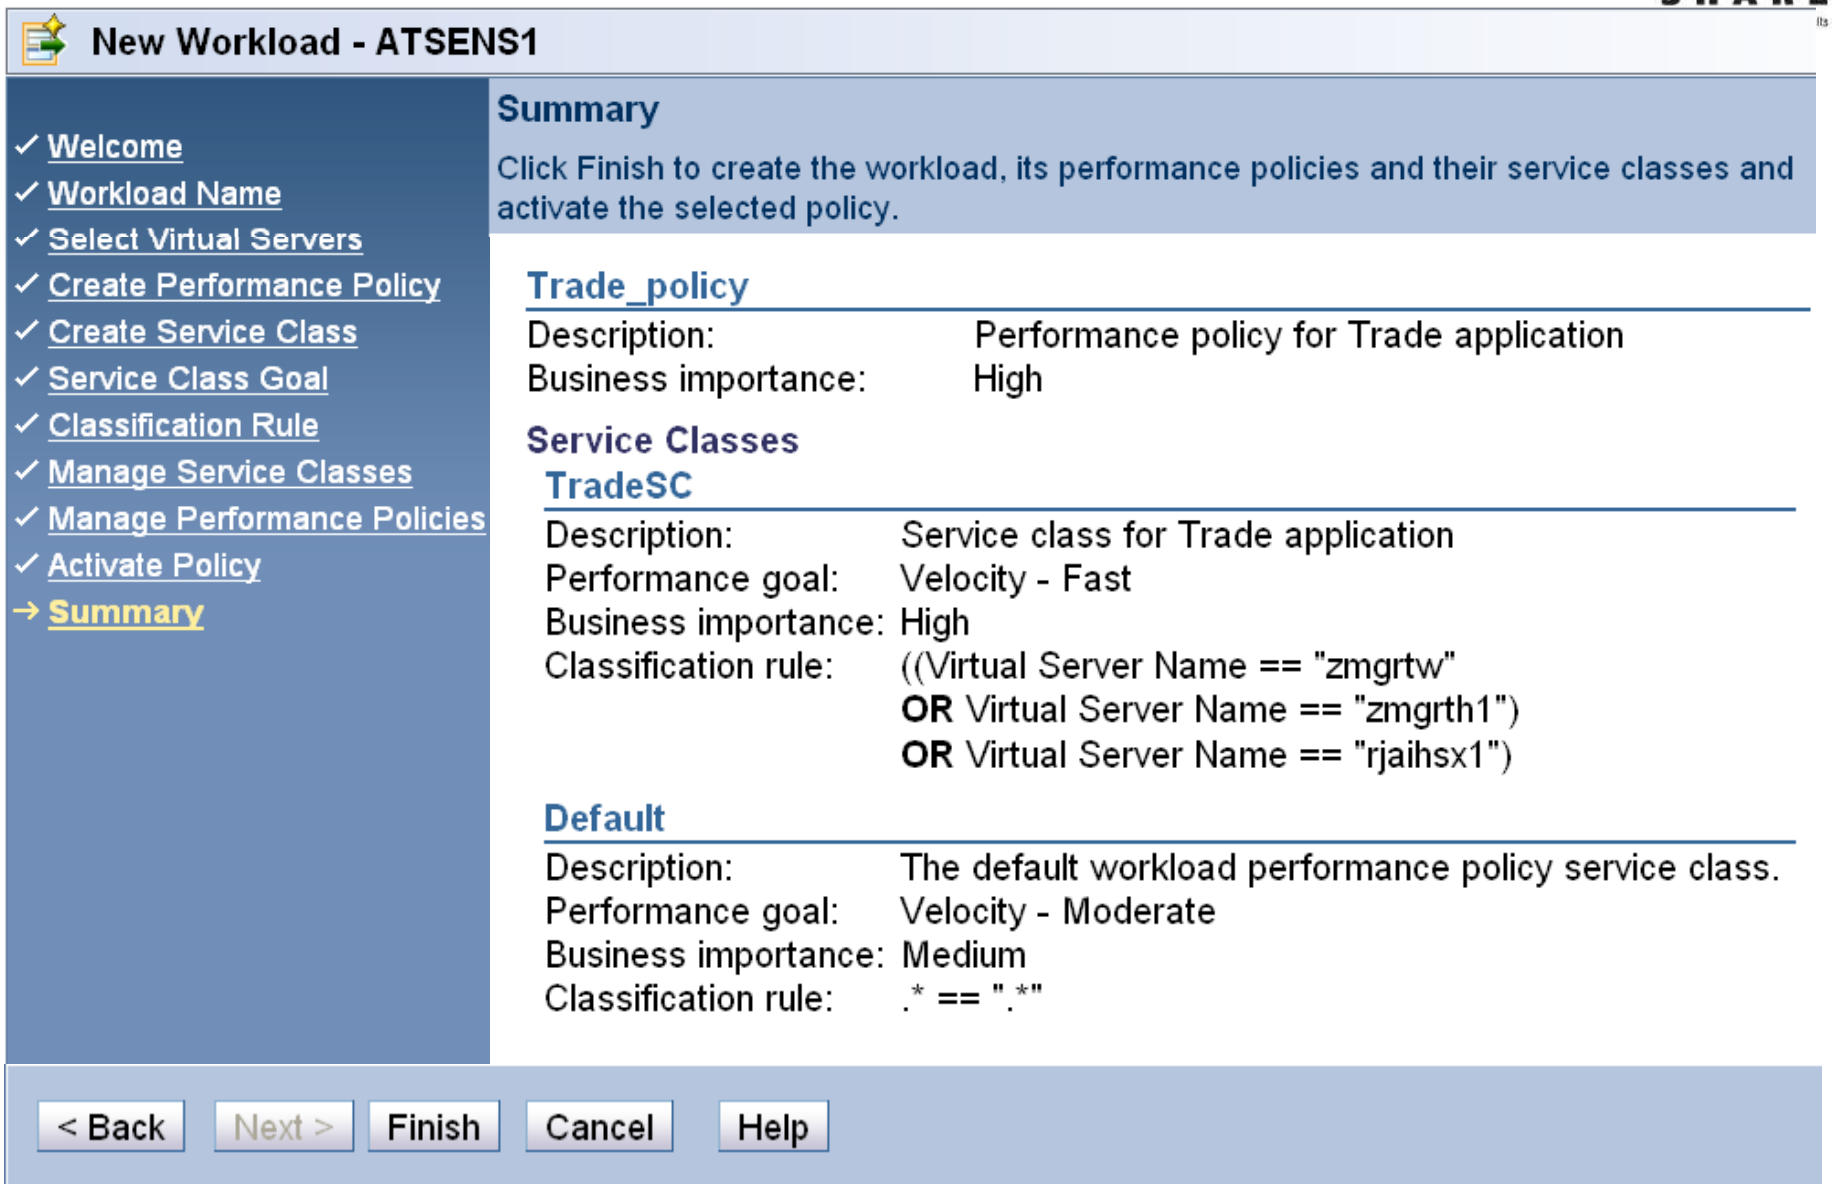

Trade Workload Performance Policy

Workload contains:

- Virtual servers
- Performance policies
- Each performance policy has service classes and classification rules

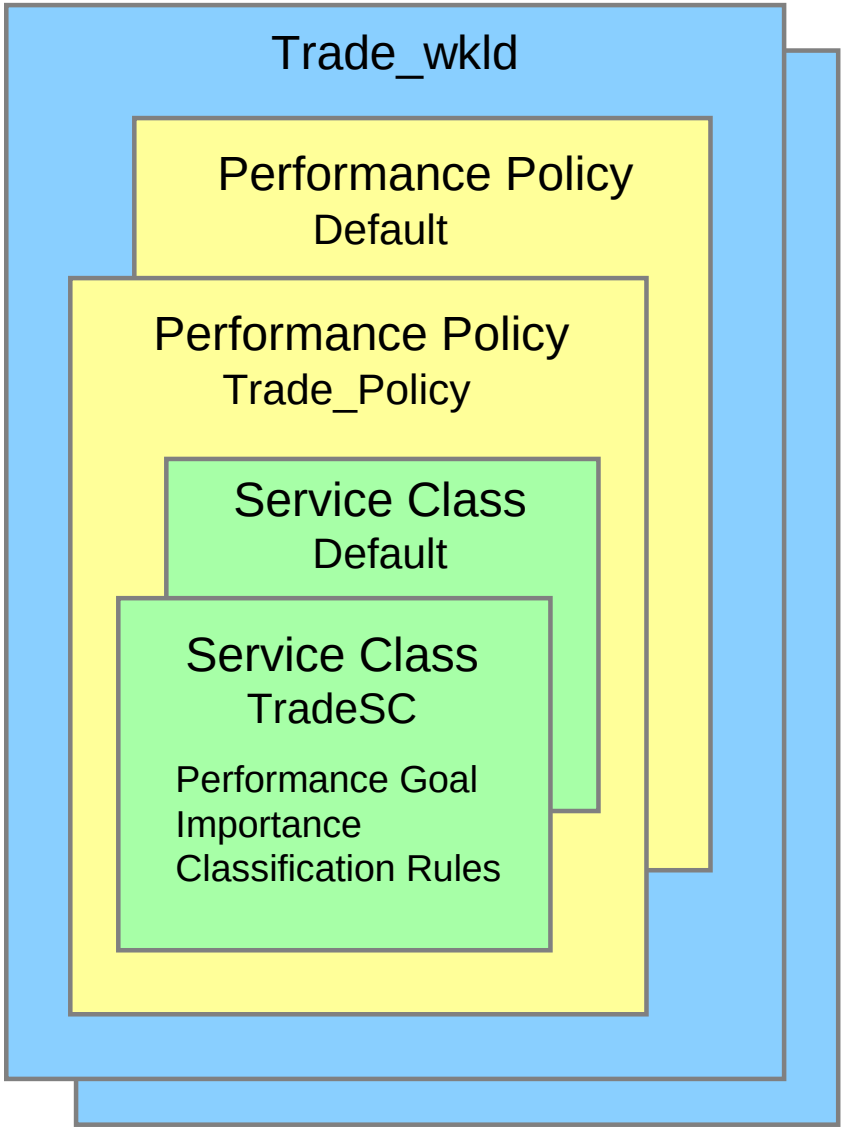

Enabling GPMP

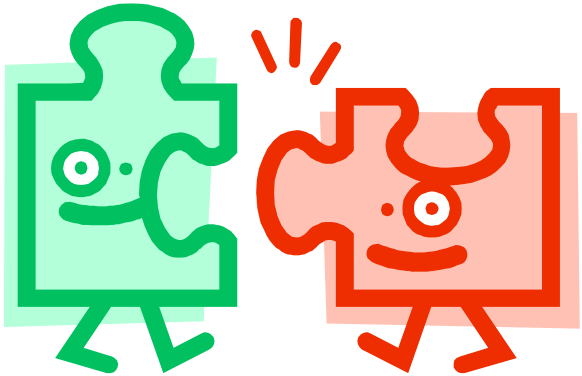

GPMP - Big Picture

GPMP – Operating system monitor information
Hypervisor - Monitoring, Resource optimization
Node – Data collection and aggregation
Ensemble HMC - Reporting, workload management

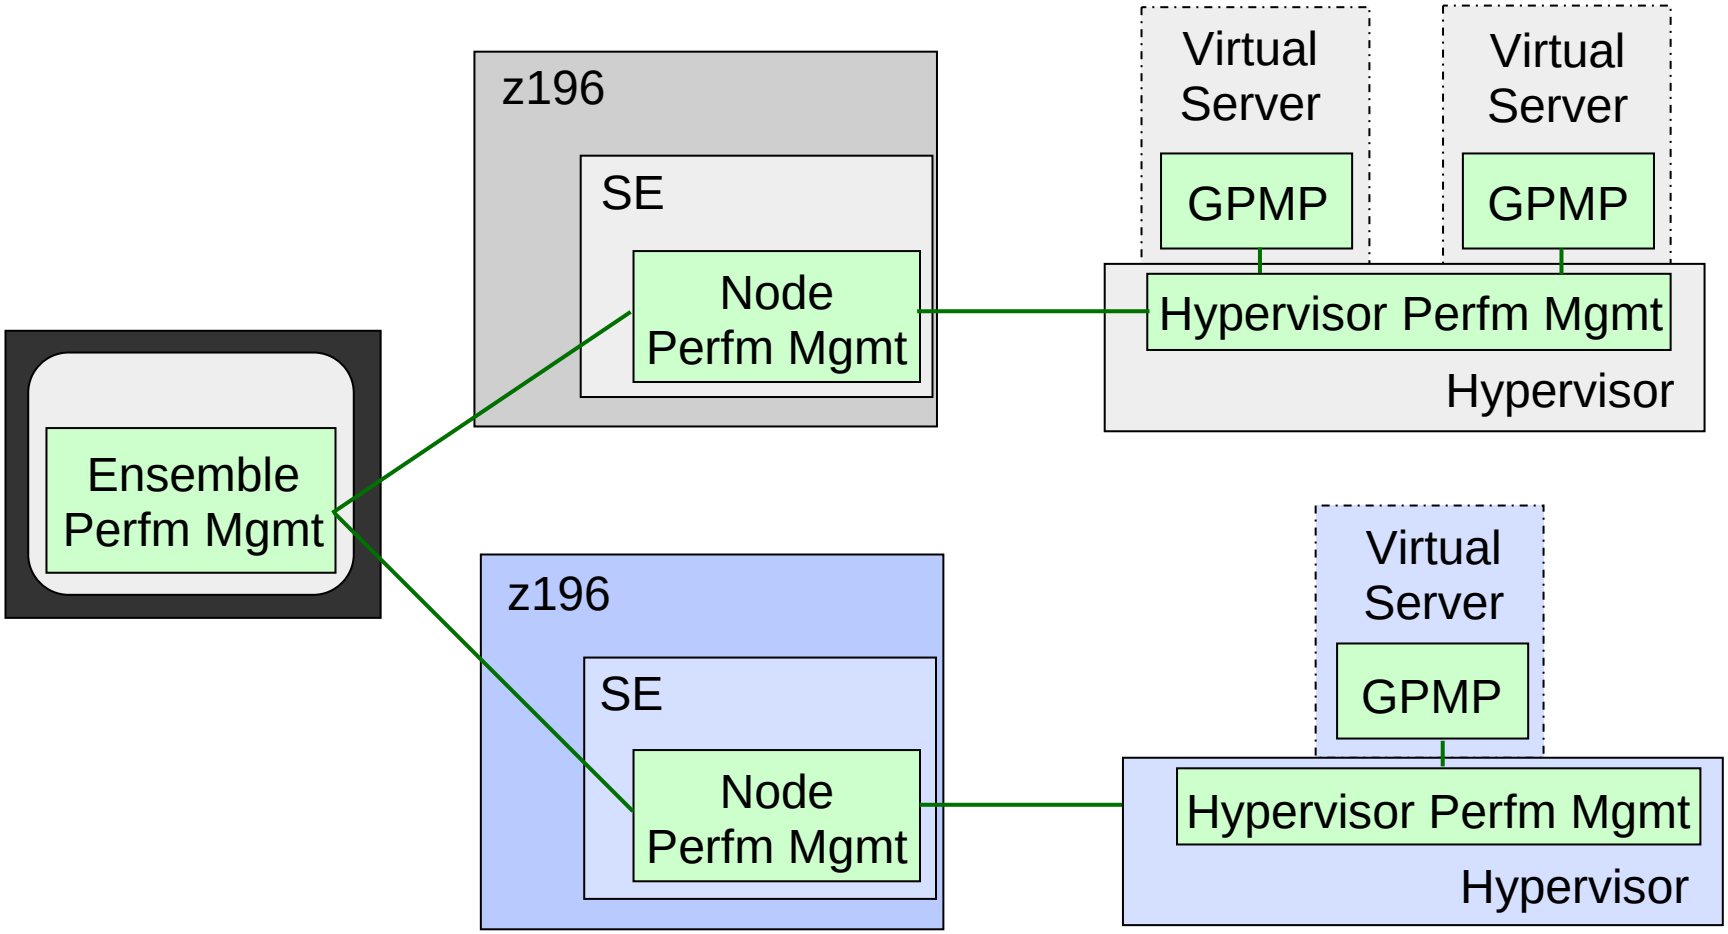

Enable GPMP on z/OS

Sample job in SYS1.SAMPLIB(HVEENV)

- Define the RACF security environment
  - Create user and group ids - group HVEMCA, user HVEMCA1
  - Authorize access to INMN
    - RDEFINE SERVAUTH EZB.OSM.sysname.tcpname
- Specifying parameters for run-time environment
  - Provide location of Java 2 1.5 or 1.6 runtime environment
  - Create UNIX file system directories
  - Specify parameters for GPMP
- Verifying HVEMCA procedure is in SYS1.PROCLIB
- Starting the GPMP address space
  - z/OS R12, R13 – WLM set to automatically start GPMP
  - z/OS R10, R11 – GPMP must be started manually

WLM Start of GPMP

WLM option to automatically start the GPMP address space

```
Definition name  . . . . WLM       (Required)
Description  . . . . . . WSC SAOPLEX Service Definition

Select one of the
following options. . . . ___  1.   Policies
                             2.   Workloads
                             3.   Resource Groups
                             4.   Service Classes
                             5.   Classification Groups
                             6.   Classification Rules
                             7.   Report Classes
                             8.   Service Coefficients/Options
                             9.   Application Environments
                             10.  Scheduling Environments
                             11.  Guest Platform Management Provider
```

WLM Start of GPMP - Systems

Automatic activate of GPMP, ability to exclude systems in sysplex

```
   Guest Platform Management Provider (GPMP) Settings
Command ===>
_____________________________________________________________

Activate guest platform management provider: 2   1. NO
                                                 2. YES

Names of systems to be excluded:

  SYSC        SYSD        _________   _________
  _________   _________   _________   _________
  _________   _________   _________   _________
```

GPMP Operator Commands

Verify ARM is enabled - D WLM,AM if not - F WLM,AM=ENABLE

Start GPMP - F WLM,GPMP,START

Display GPMP Status - D WLM,AM,ALL

```
IWM075I  17.49.20  WLM DISPLAY 433
  ARM SERVICES ARE ENABLED
  GUEST PLATFORM MANAGEMENT PROVIDER JOBNAME=HVEMCA ASID=0018
  GPMP POLICY IS ACTIVE
  NUMBER OF REGISTERED PROCESSES=1, APPLICATIONS=1
  ADDRESS SPACES CURRENTLY REGISTERED WITH ARM:
    JOBNAME=DSNADIST ASID=004D
      APPLICATION=DDF
        IDENTITY PROPERTIES=0 CONTEXT NAMES=0
        STARTED APPLICATION INSTANCES:
          DSN9
            TRAN=0 GROUP=DSN9WSC
        REGISTERED TRANSACTIONS:
          SYS_DefaultZWLMTransactionName
```

Enable GPMP on Virtual Server

Enable GPMP for virtual server

Virtual Server Details -> Options -> Enable GPMP support

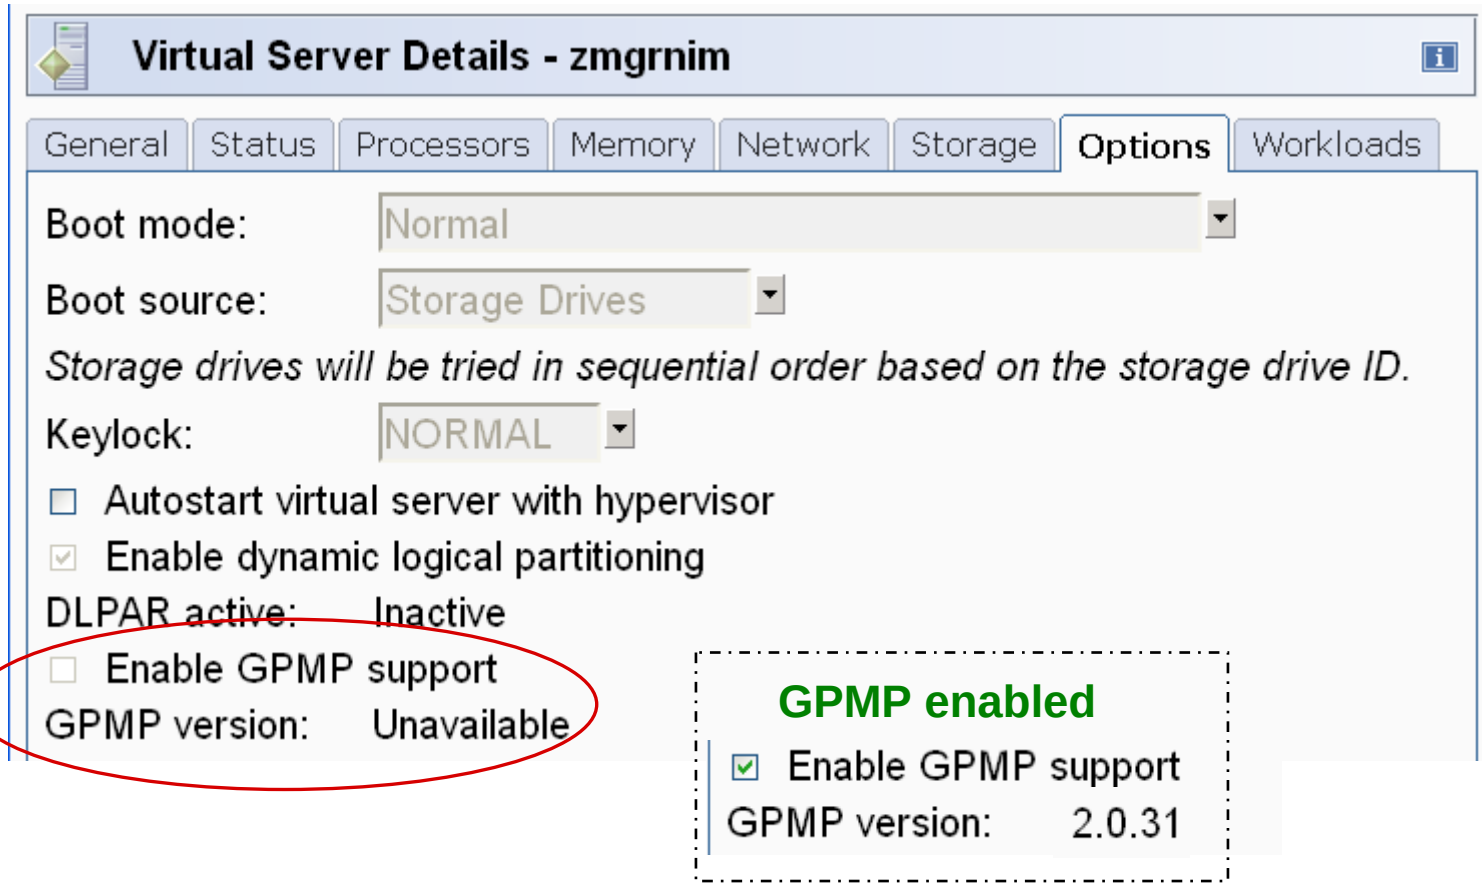

GPMP Installation Image

GPMP is delivered in a software package

For AIX, Linux on system x, and Windows

upload the image with the HMC Mount Virtual Media task

mount GPMP Installation Image on virtual server's CDROM drive

Virtual Server -> Configuration -> Mount Virtual Media

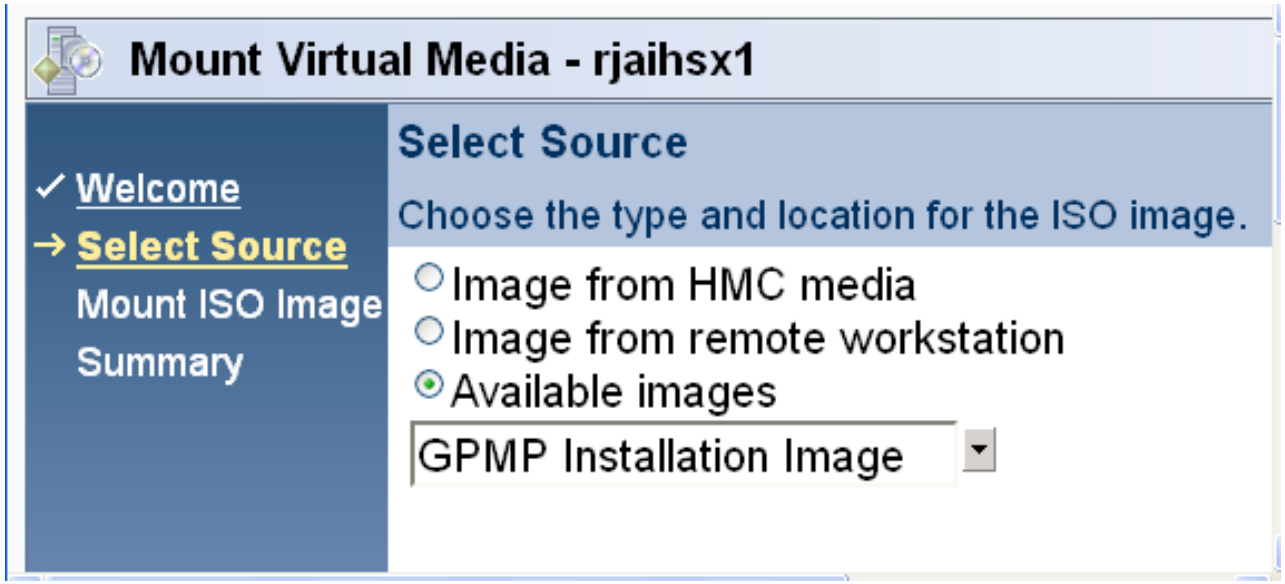

Install GPMP on Operating System

Enable IPV6 for access to INMN

Enable EWLM services on AIX virtual server

- With smitty ewlm or command ewlmcfg –c
- Confirm EWLM services enabled with command: ewlmcfg –q

Install GPMP image

- GPMP image accessible on a directory mount point
- Install using method appropriate for operating system (AIX, Linux on system x, Windows, or Linux on system z)
- Default group is ibmgpmp
- Default user is ibmgpmp

Start GPMP on Virtual Servers

Start GPMP

- AIX, Linux on system x, zVM Linux guest
  su ibmgpmp -c "/opt/ibm/gpmp/gpmp start“
- Windows
  Select Launch the IBM GPMP option on installation panel
  c:\Program Files\IBM\gpmp\gpmp start

Set GPMP to start automatically

- AIX, Linux on system x, zVM Linux guest
  su ibmgpmp -c "/opt/ibm/gpmp/gpmp autostart on“
- Windows
  c:\Program Files\IBM\gpmp\gpmp autostart on

Display GPMP status

- AIX, Linux on system x, zVM Linux guest
  su ibmgpmp -c "/opt/ibm/gpmp/gpmp”“
- Windows
  c:\Program Files\IBM\gpmp\gpmp

GPMP Status

```
su ibmgpmp -c "/opt/ibm/gpmp/gpmp"
FEW6030I Persistent storage settings for the guest platform management provider:
FEW6034I        Autostart flag is on
FEW6036I        Shared memory ID is 4
FEW6037I        The guest platform management provider is not running
```

```
 su ibmgpmp -c "/opt/ibm/gpmp/gpmp start“
FEW6101I The guest platform management provider is starting.
```

```
su ibmgpmp -c "/opt/ibm/gpmp/gpmp"
FEW6030I Persistent storage settings for the guest platform management provider:
FEW6034I        Autostart flag is on
FEW6036I        Shared memory ID is 4
FEW6038I        Main process ID is 7995420
```

Performance Monitoring and Reporting

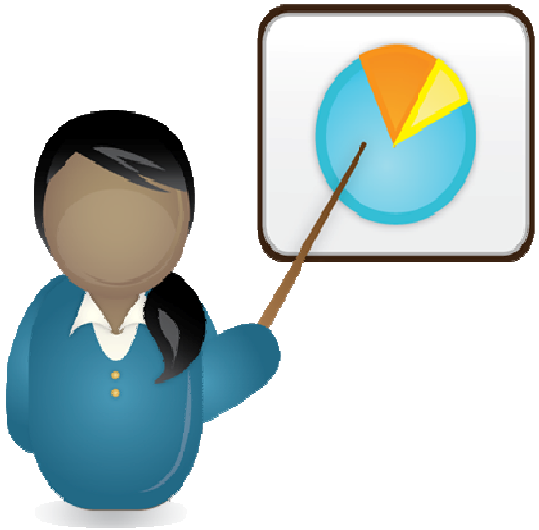

Workload Monitoring and Reporting

Report virtual server resource usage in a Workload

User interface for reports is the ensemble HMC

Reports current data and fairly recent history

- Interval of data displayed is user selectable
- Granularity of data kept changes over time
  - 1 minute granularity kept for most recent hour
  - 15 minute interval data kept after first hour
- History of 36 hours

Report data can be downloaded to local workstation

- Uses CSV format
- Can only download data currently represented on screen

Navigating the Reports

Workloads report lists all workloads

These reports are for a specific workload

- Service Classes Report
- Virtual Servers Report
- Resource Adjustment Report
- Virtual Server Topology Report
- Hops Report

Workloads Report

List of workloads

High level view of “performance health” of each workload

- Indication if workload service class is missing goals
- Locate worst performing service class / performance index (PI)
- Details for a specific workload

  Bar graph of virtual server utilization distribution

  Graph of service class PI

Workloads Report Example

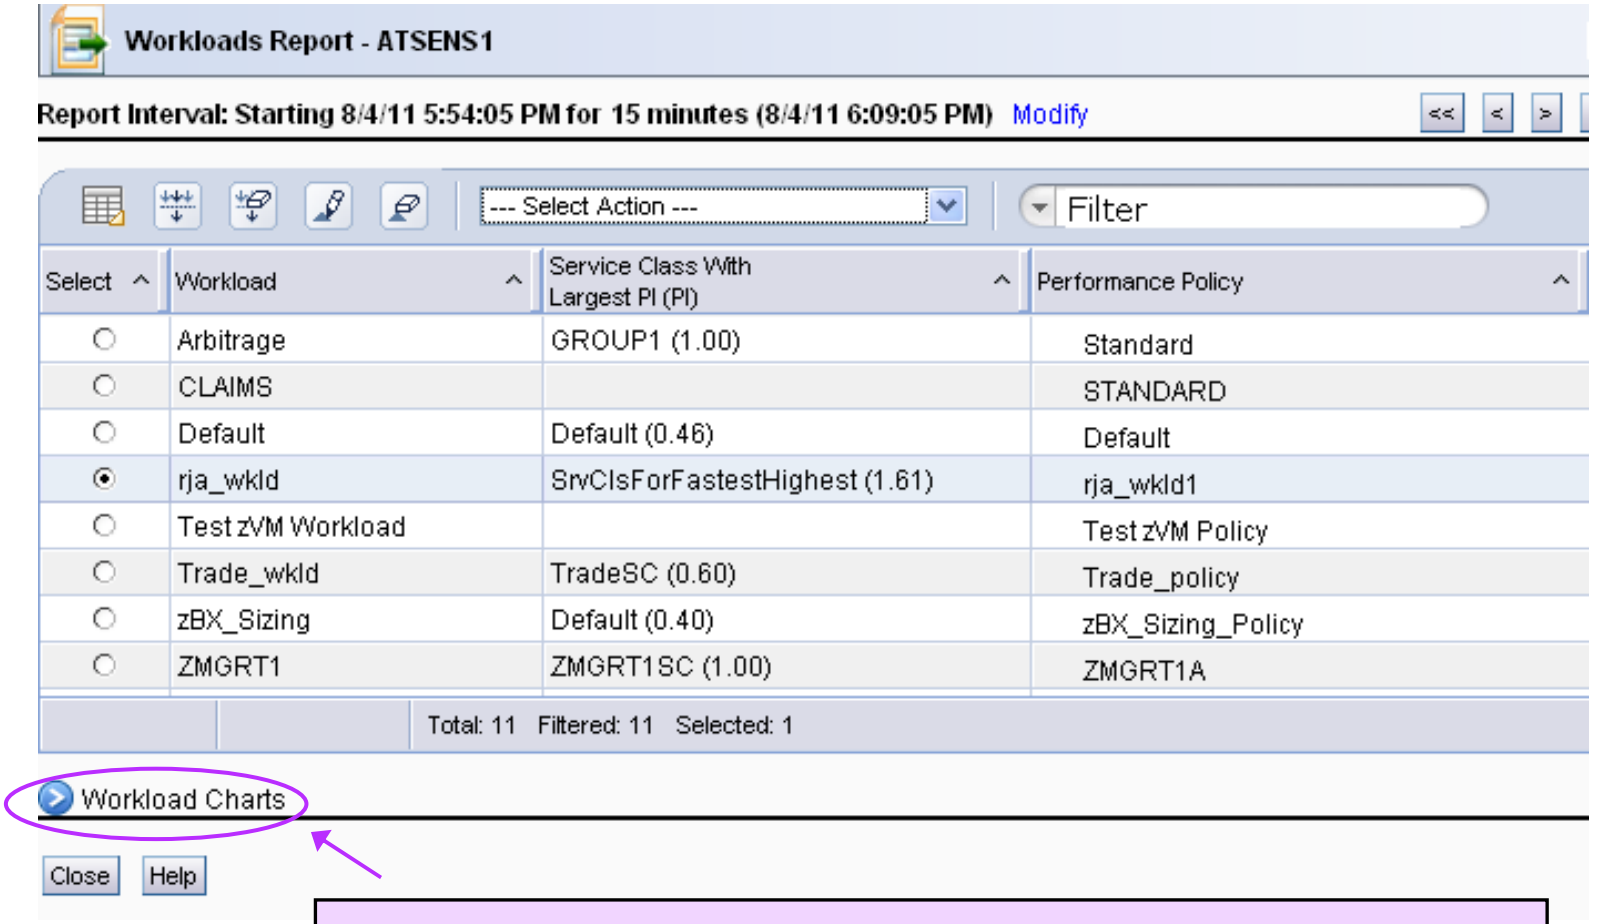

Service class with largest PI is rja_wkld with 1.61
View workload information for rja_wkld

Workloads Report – CPU Utilization

Charts: CPU Utilization | Performance Index

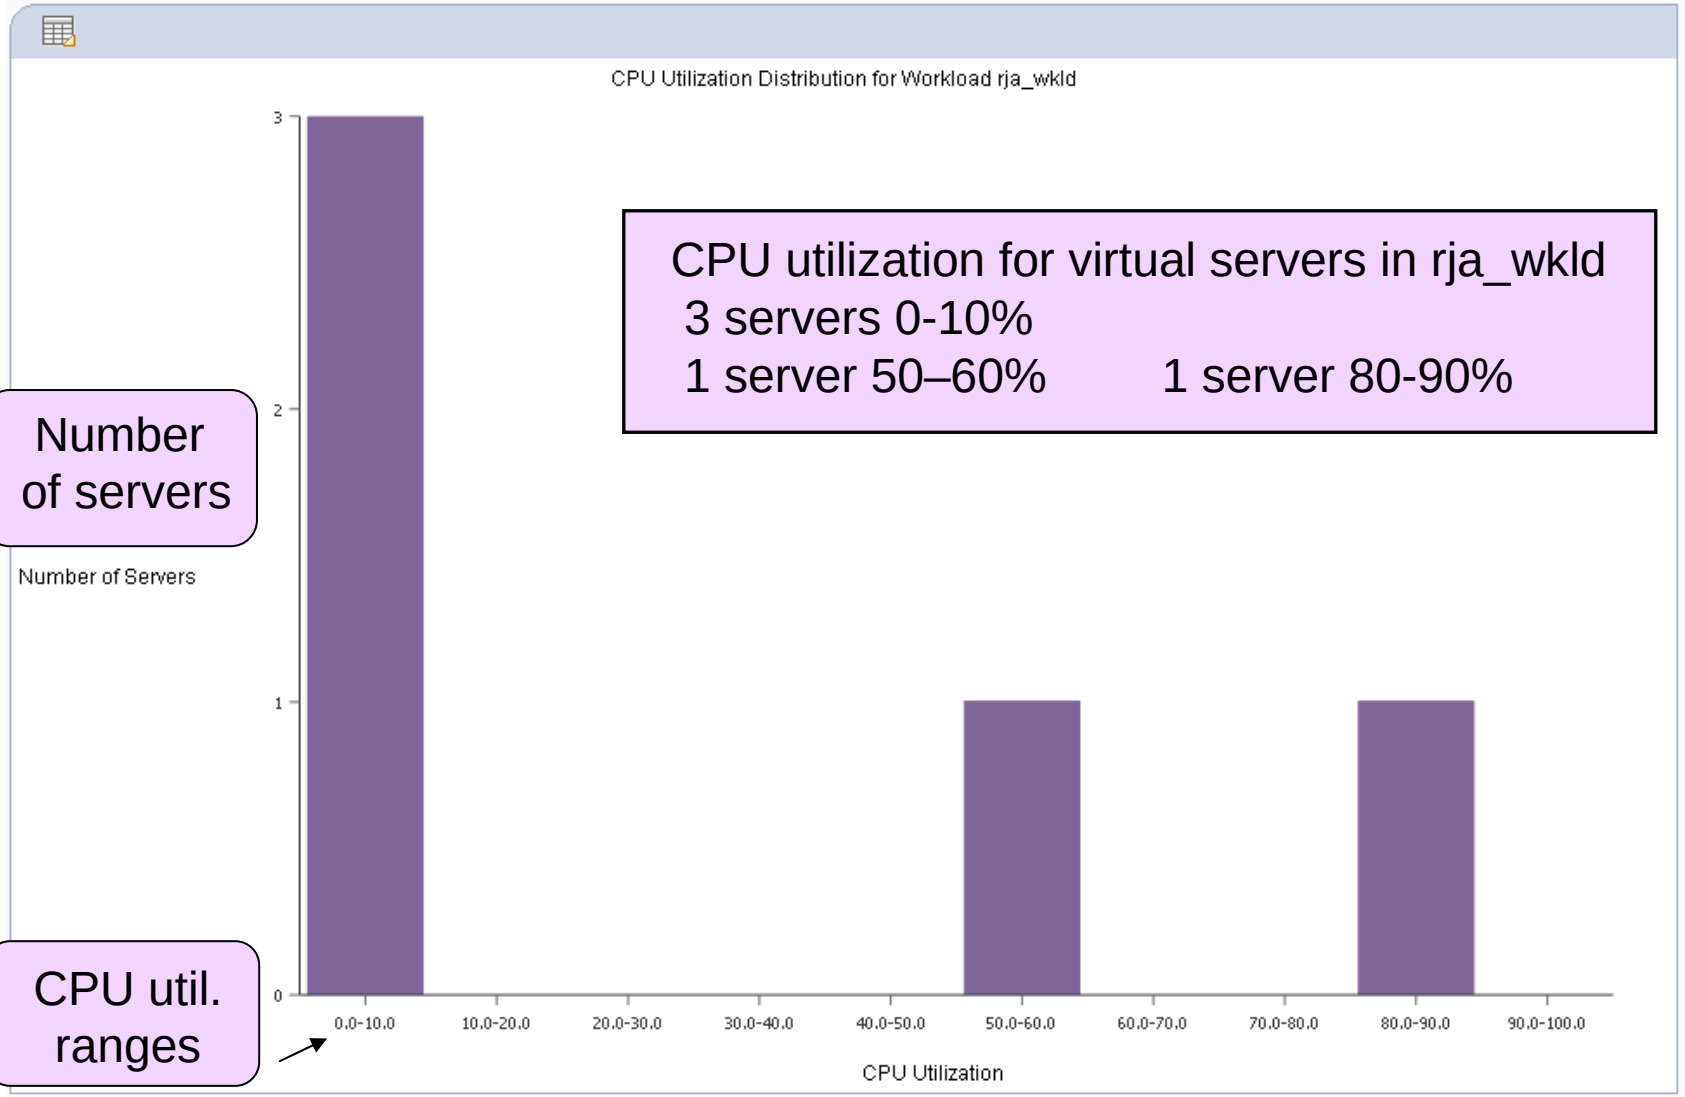

CPU utilization for virtual servers in rja_wkld
3 servers 0-10%
1 server 50–60% 1 server 80-90%

Number of servers

CPU util. ranges

Workloads Report – PI

Charts: CPU Utilization | Performance Index

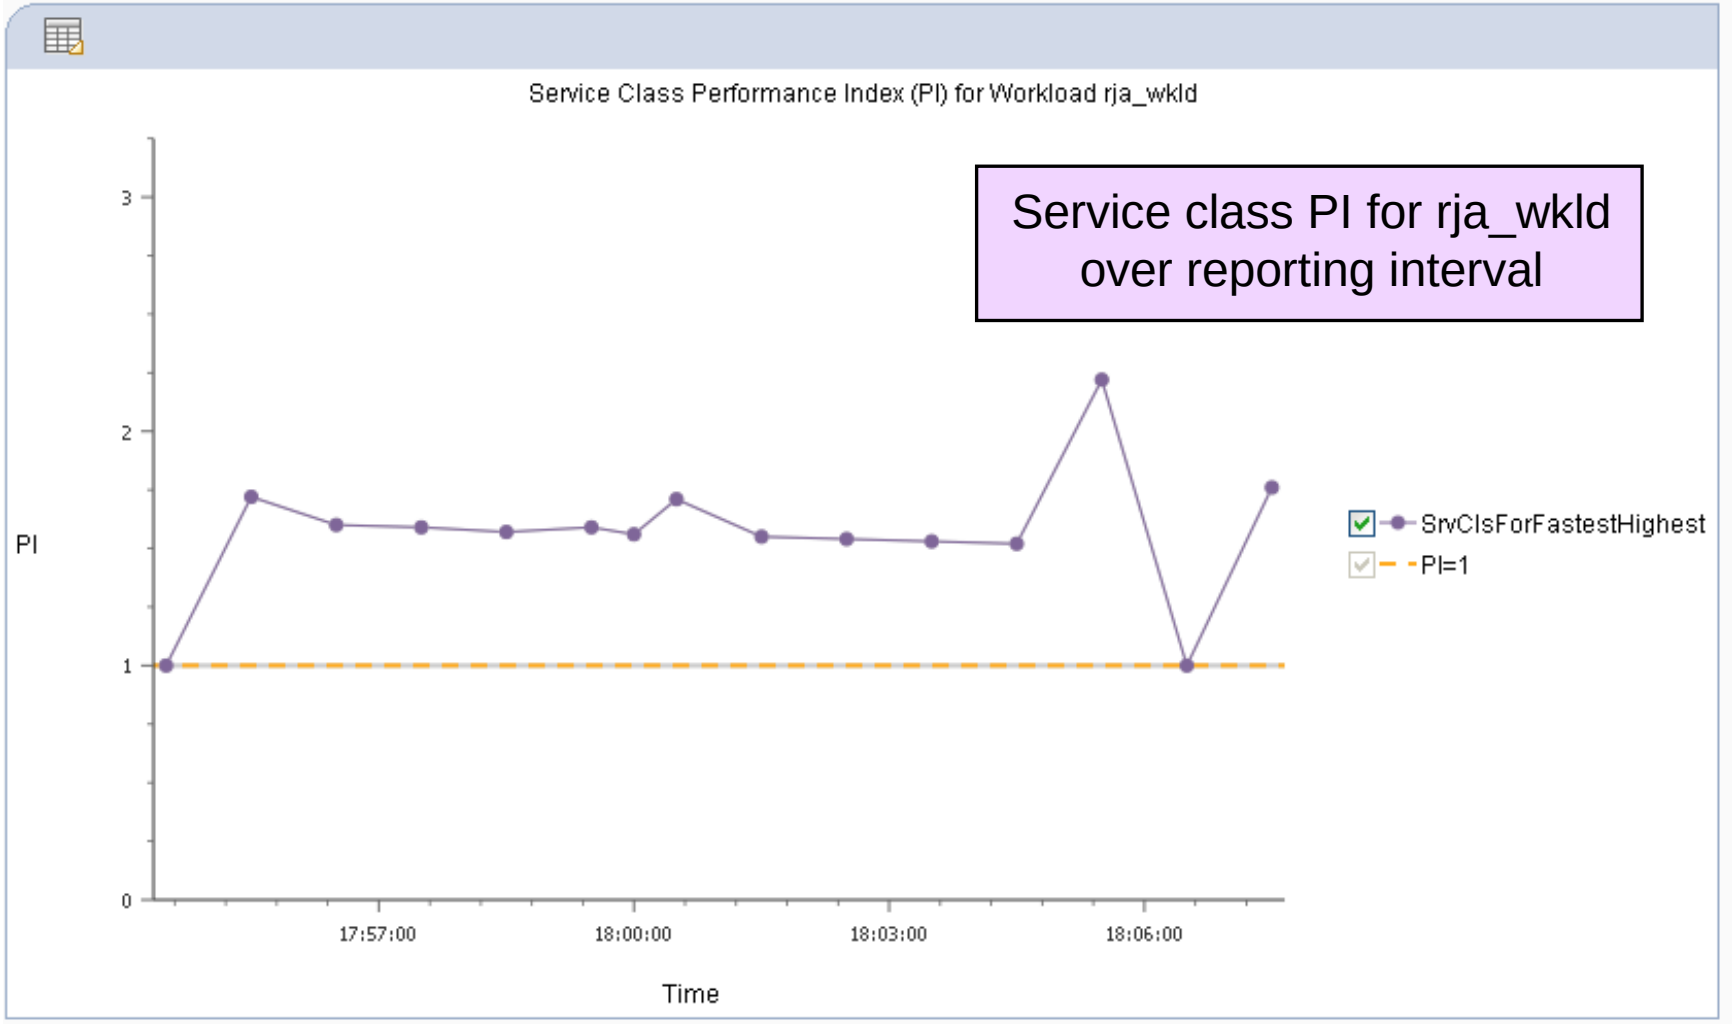

Virtual Servers

Virtual Servers report

- List of virtual servers in the workload
- Resource information for each virtual server
  - Hypervisor
  - Virtual processors
  - Service class and PI
  - Allocated memory
  - Physical CPU utilization
  - OS view of CPU utilization

Resource adjustment report

- Resource adjustment actions taken over report interval
- Donor and receiver virtual servers

Virtual Servers Report

Virtual Servers Report - rja_wkld

Report Interval: Starting 2/2/12 1:30:00 AM for 15 minutes (2/2/12 1:45:00 AM) Modify

--- Select Action --- Filter

Select	Virtual Server	Hypervisor	Hypervisor Type	Hostname	Virtual Processors	Allocated Memory (MB)	Physical CPU Utilization (%)	Hypervisor CPU Delay (%)	Service Class (PI)
◉	rjaihs1	B.2.14	PowerVM	rja_ihs1.dmz	2	8,192	60.7	39.2	Default (0.48)
○	rjaihsx1	C.1.02	x Hyp	zbx-rja-ihsx1	4	4,096	0.0	0.1	SrvClsForFastestHighest (1.59)
○	rjawas1	B.2.14	PowerVM	rjawas1.dmz	5	8,192	0.5	0.1	SrvClsForFastestHighest (1.59)
○	rjawas2	B.2.14	PowerVM	rja_was2.dmz	5	8,192	93.6	6.0	Default (0.48)
○	TOSP11	TSYS	PR/SM	ZBLC.dmz	4	4,096	0.0		SrvClsForFastestHighest (1.59)

Page 1 of 1 Total: 5 Filtered: 5 Displayed: 5 Selected: 1

Five virtual servers in rja_wkld

Physical CPU Utilization

Service Class

Resource Adjustment Report

Service Class Resource Adjustments Report - SrvClsForFastestHighest

Report Interval: Starting 2/2/12 2:13:08 AM for 15 minutes (2/2/12 2:28:08 AM) Modify

Successful Adjustments:

--- Select Action --- | Filter

Virtual Server	Type	Workload	Service Class	Processing Units Before	Processing Units After	Time
rjawas1	Receiver	rja_wkld	SrvClsForFastestHighest	0.50	0.55	Feb 2, 2012 2:22:54 AM EST
rjawas2	Donor	rja_wkld		4.00	3.95	
rjawas1	Receiver	rja_wkld	SrvClsForFastestHighest	0.55	0.57	Feb 2, 2012 2:24:26 AM EST
rjawas2	Donor	rja_wkld		3.95	3.93	
rjawas1	Receiver	rja_wkld	SrvClsForFastestHighest	0.57	0.59	Feb 2, 2012 2:25:57 AM EST
rjawas2	Donor	rja_wkld		3.93	3.91	
rjawas1	Receiver	rja_wkld	SrvClsForFastestHighest	0.59	0.79	Feb 2, 2012 2:27:28 AM EST
rjawas2	Donor	rja_wkld		3.91	3.71	

Page 1 of 1 Total: 8 Filtered: 8 Displayed: 8

Receiver rjawas1 before 0.50 after 0.55 processing units
Donor rjawas2 before 4.00 after 3.95 processing units

Hops and Topology Reports

Information from ARM and GPMP used to create reports

Hops report

- Shows each hop for application in a specific service class
- For each hop provides
  - Name, hop number
  - Transaction information
  - Average response times

Virtual Server Topology Report

- Relationship of virtual servers running a workload
- Graphical representation of virtual servers

Virtual Server Topology Report

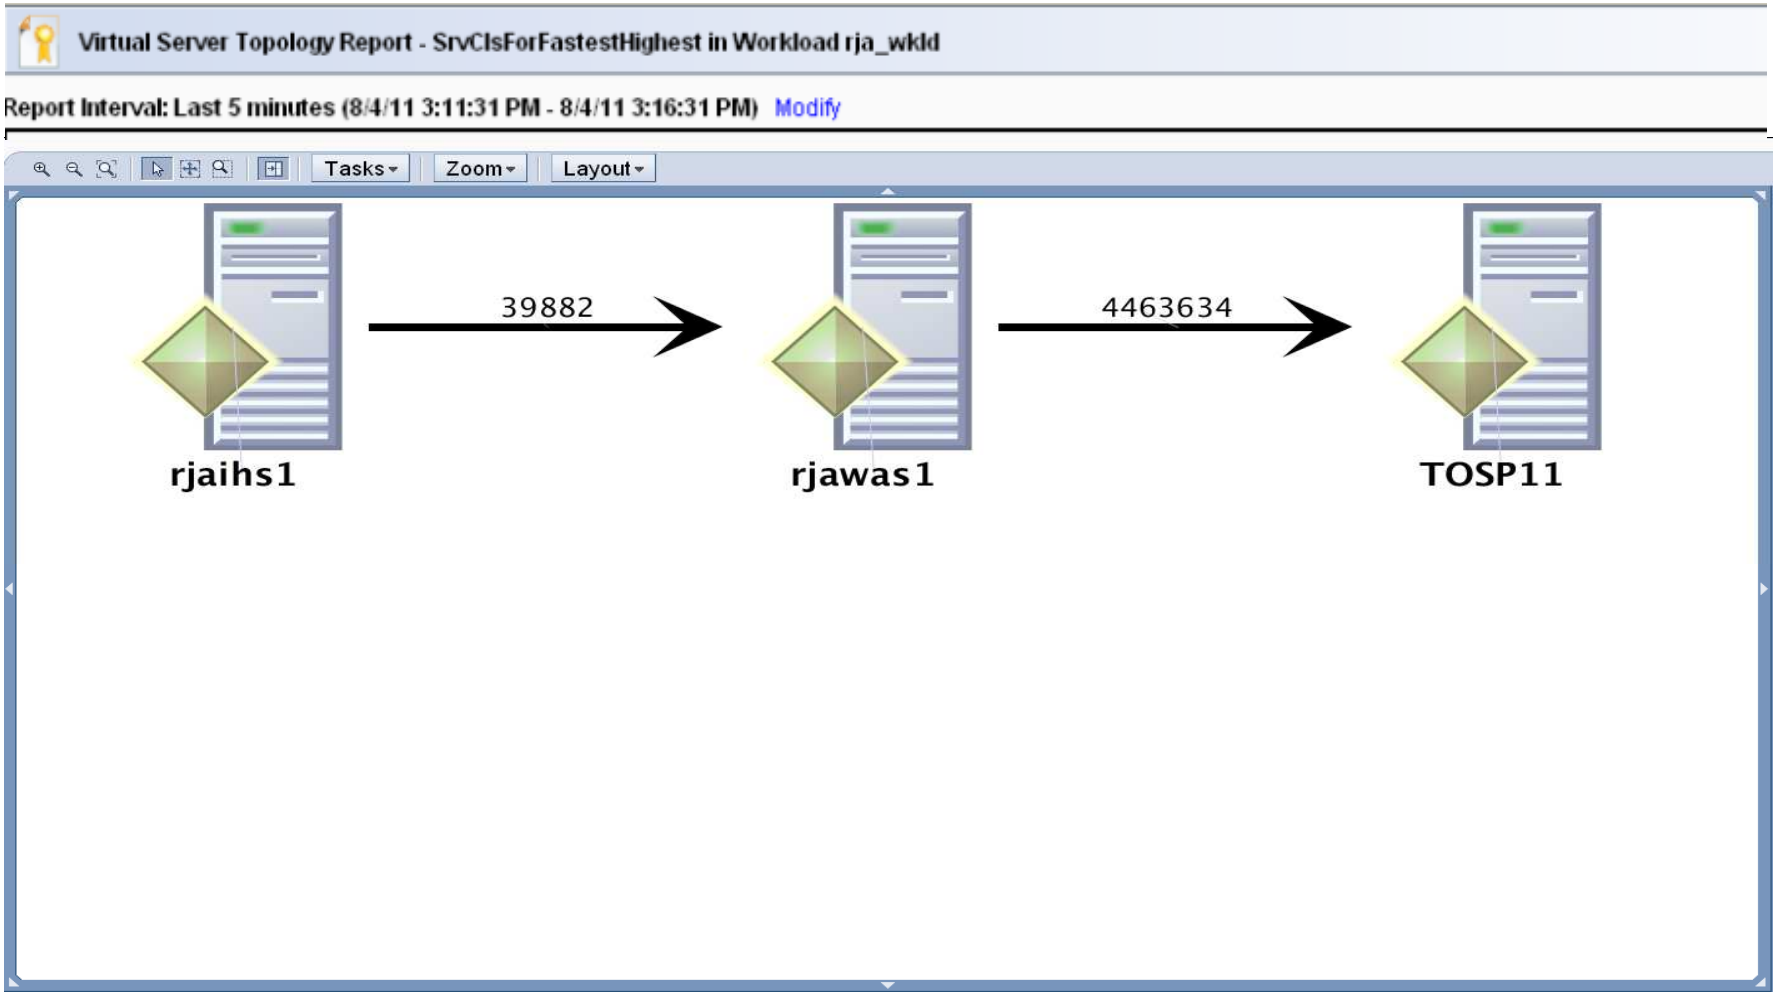

Hops Report

Hops Report - SrvClsForFastestHighest in Workload rja_wkld

Report Interval: Last 5 minutes (8/4/11 3:08:36 PM - 8/4/11 3:13:36 PM) Modify

Details for SrvClsForFastestHighest:

Workload:	rja_wkld	Performance goal:	Velocity - Fastest	PI:	1.00
Performance policy:	rja_wkld1	Business importance:	Highest	Performance:	Fastest

--- Select Action --- Filter

Name	Hop Num	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		22,281	0	0	202	0.000000	0.000273	0.165120	0.094535
IBM Webserving Plugin	0	IBM_HTTP_Server	22,281	0	0	202	0.000000	0.000273	0.165120	0.094535
rjaihs1	0		22,281	0	0	202	0.000000	0.000273	0.165120	0.094535
Hop 1	1		37,486	0	0	90	0.000000	0.001384	0.019575	0.015683
WebSphere:APPLICATION	1	server1	37,486	0	0	90	0.000000	0.001384	0.019575	0.015683
rjawas1	1		37,486	0	0	90	0.000000	0.001384	0.019575	0.015683
Hop 2	2		4,363,388	0	0	3	0.000000	0.000265	0.000265	0.000673
DDF	2	DSN9WSC	4,363,388	0	0	3	0.000000	0.000265	0.000265	0.000673
TOSP11	2		4,363,388	0	0	3	0.000000	0.000265	0.000265	0.000673

Transaction average
- Execution time
- Response time

GPMP Stopped

Test with GPMP stopped on AIX virtual servers and z/OS

Reports affected

- Hops report
- Virtual server topology report
- Virtual server report

Topology Report – GPMP Stopped

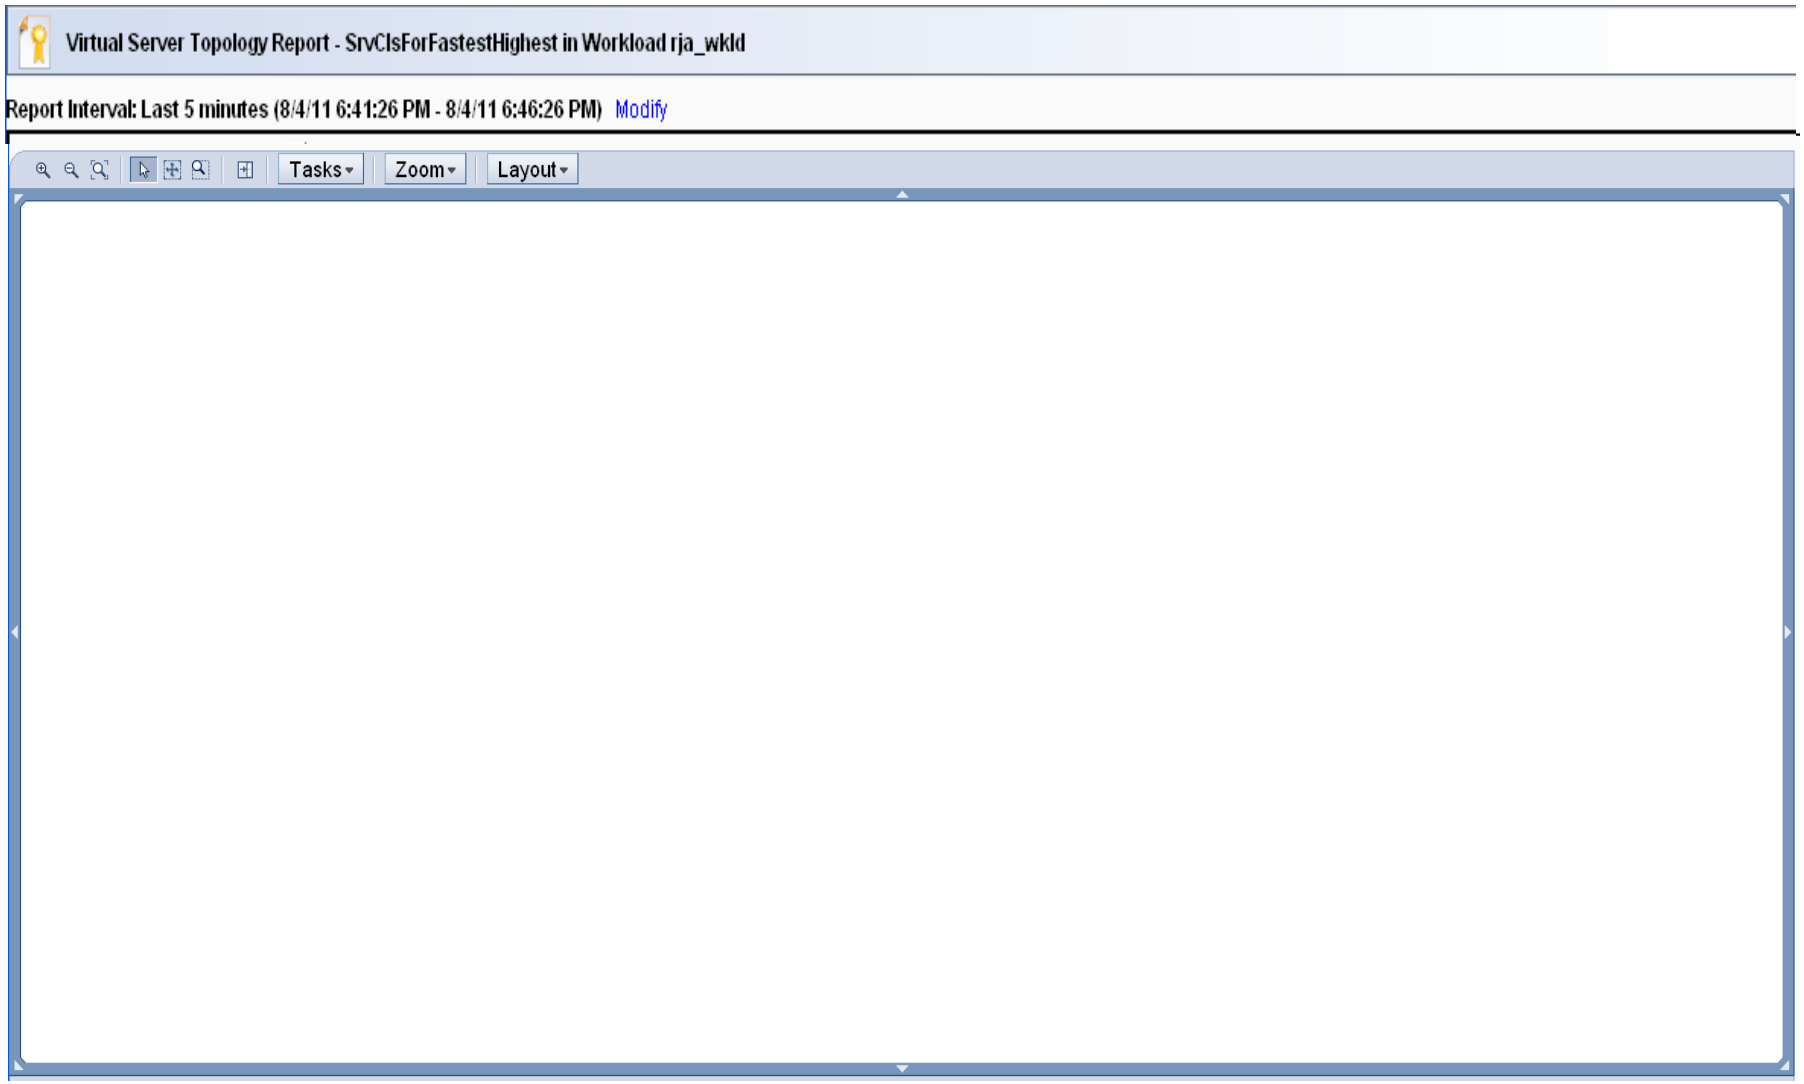

Hops Report – GPMP Stopped

Hops Report - SrvClsForFastestHighest in Workload rja_wkld

Report Interval: Last 5 minutes (8/4/11 6:39:39 PM - 8/4/11 6:44:39 PM) Modify

Details for SrvClsForFastestHighest:

Workload: rja_wkld Performance goal: Velocity - Fastest PI: 1.00

Performance policy: rja_wkld1 Business importance: Highest Performance: Fastest

--- 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		0	0	0	4	0.000000	0.000000	0.000000	374.435577
IBM Webserving Plugin	0	IBM_HTTP_Server	0	0	0	0	0.000000	0.000000	0.000000	0.000000
rjaihs1	0		0	0	0	0	0.000000	0.000000	0.000000	0.000000
WebSphere:APPLICATION_SERVER	0	server1	0	0	0	4	0.000000	0.000000	0.000000	374.435577
rjawas1	0		0	0	0	4	0.000000	0.000000	0.000000	374.435577
Hop 1	1		0	0	0	0	0.000000	0.000000	0.000000	0.000000
WebSphere:APPLICATION_SERVER	1	server1	0	0	0	0	0.000000	0.000000	0.000000	0.000000
rjawas1	1		0	0	0	0	0.000000	0.000000	0.000000	0.000000
Hop 2	2		0	0	0	0	0.000000	0.000000	0.000000	0.000000
DDF	2	DSN9WSC	0	0	0	0	0.000000	0.000000	0.000000	0.000000
TOSP11	2		0	0	0	0	0.000000	0.000000	0.000000	0.000000

No data from systems

Virtual Server Report – GPMP Started

Virtual Servers Report - rja_wkld

Report Interval: Starting 8/4/11 2:54:07 PM for 15 minutes (8/4/11 3:09:07 PM) Modify

<< < > >>

--- Select Action ---

Select	Virtual Server	Virtual Processors	Allocated Memory (MB)	Physical CPU Utilization (%)	Hypervisor CPU Delay (%)	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 (%)
◉	rjaihs1	2	8,192	23.4	10.8	SrvClsForFastestHighest (1.00)	13.5	20.9	0.0	65.6
○	rjawas1	4	8,192	71.1	3.1	SrvClsForFastestHighest (1.00)	3.4	1.8	0.0	94.8
○	rjawas2	5	8,192	45.5	46.4	Default (0.53)	2.9	0.3	0.0	96.7
○	TOSP11	2	4,096	6.0		SrvClsForFastestHighest (1.00)	1.0	40.3	0.4	0.0

Virtual Servers

OS Processes columns contain data provided by GPMP

Virtual Server Report – GPMP Stopped

Virtual Servers Report - rja_wkld

Report Interval: Last 5 minutes (8/4/11 6:34:34 PM - 8/4/11 6:39:34 PM) Modify

--- Select Action ---

Select	Virtual Server	Virtual Processors	Allocated Memory (MB)	Physical CPU Utilization (%)	Hypervisor CPU Delay (%)	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 (%)
◉	rjaihs1	2	8,192	22.7	8.3	SrvClsForFastestHighest (1.00)				
○	rjawas1	4	8,192	63.3	15.2	SrvClsForFastestHighest (1.00)				
○	TOSP11	2	4,096	5.9		SrvClsForFastestHighest (1.00)				

OS Processes columns are blank

Summary

Platform Performance Manager function of zManager provides tools to manage workloads running in the ensemble

Resources are directed to virtual servers based on the goals and importance levels of the workload

HMC is user interface to create workloads and view reports

ARM enabled middleware and GPMP allow for end-to-end monitoring of application performance

Manage workloads in multi-tier application environment