zEnterprise eXposed!

zManager and z/OS Workload Manager

Session 10907

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.

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

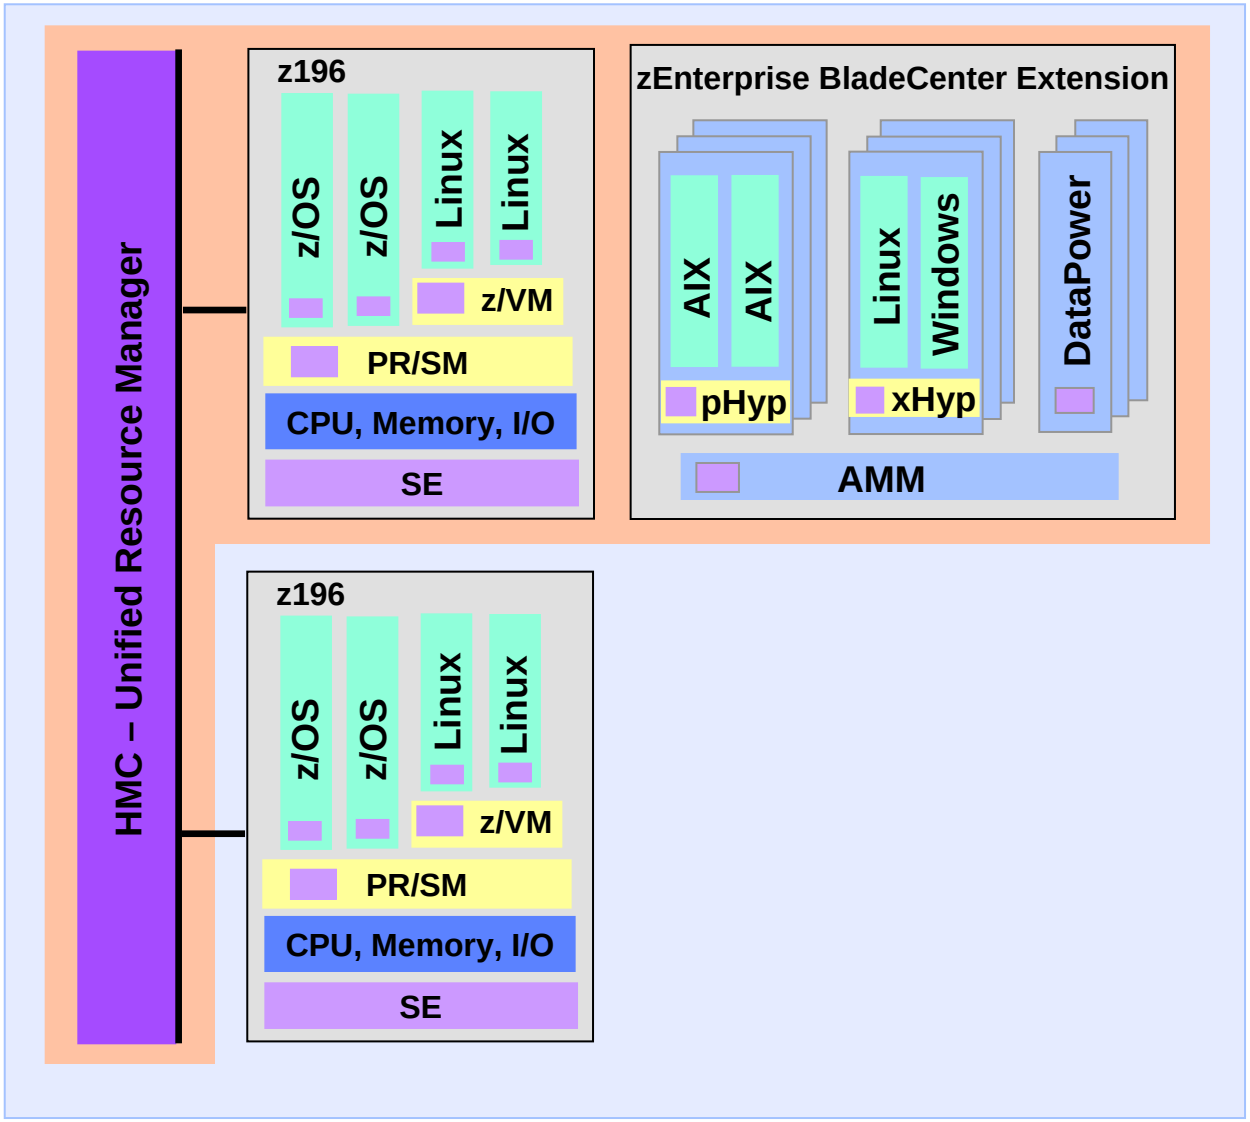

Unified Resource Manager (zManager) Platform Performance Management

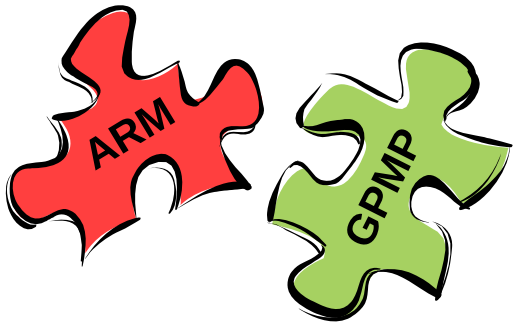

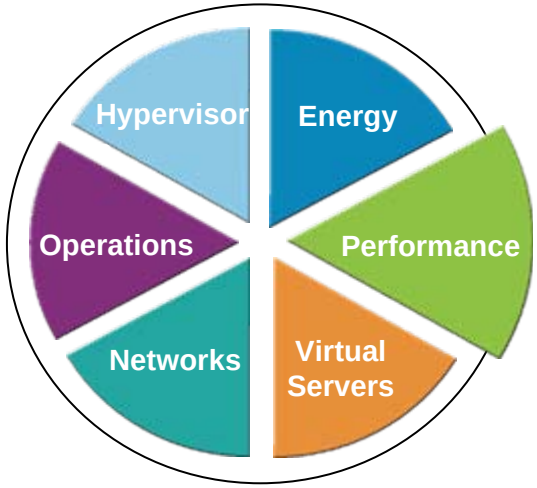

Platform Performance Management

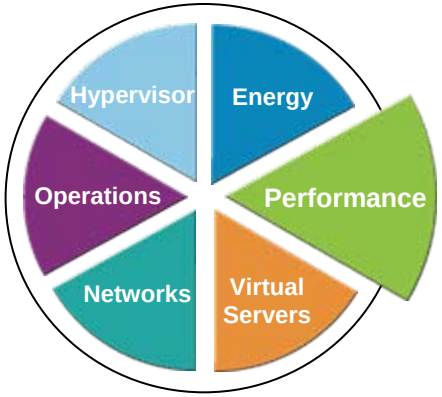

- zManager component responsible for monitoring, reporting, and management of resources used by virtual servers
- Scope is the ensemble
- User interface is the ensemble HMC
- Workload goals specified in workload performance policy
- Based on goals in workload policy, PPM may adjust processor resources 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
- Sample job in SYS1.SAMPLIB(HVEENV) to implement GPMP
  - Define RACF security environment
  - Specify run-time environment parameters
  - Verify HVEMCA procedure is in SYS1.PROCLIB
- Operator command to start GPMP => F WLM,START,GPMP
  - z/OS R12, R13 – WLM can be setup 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

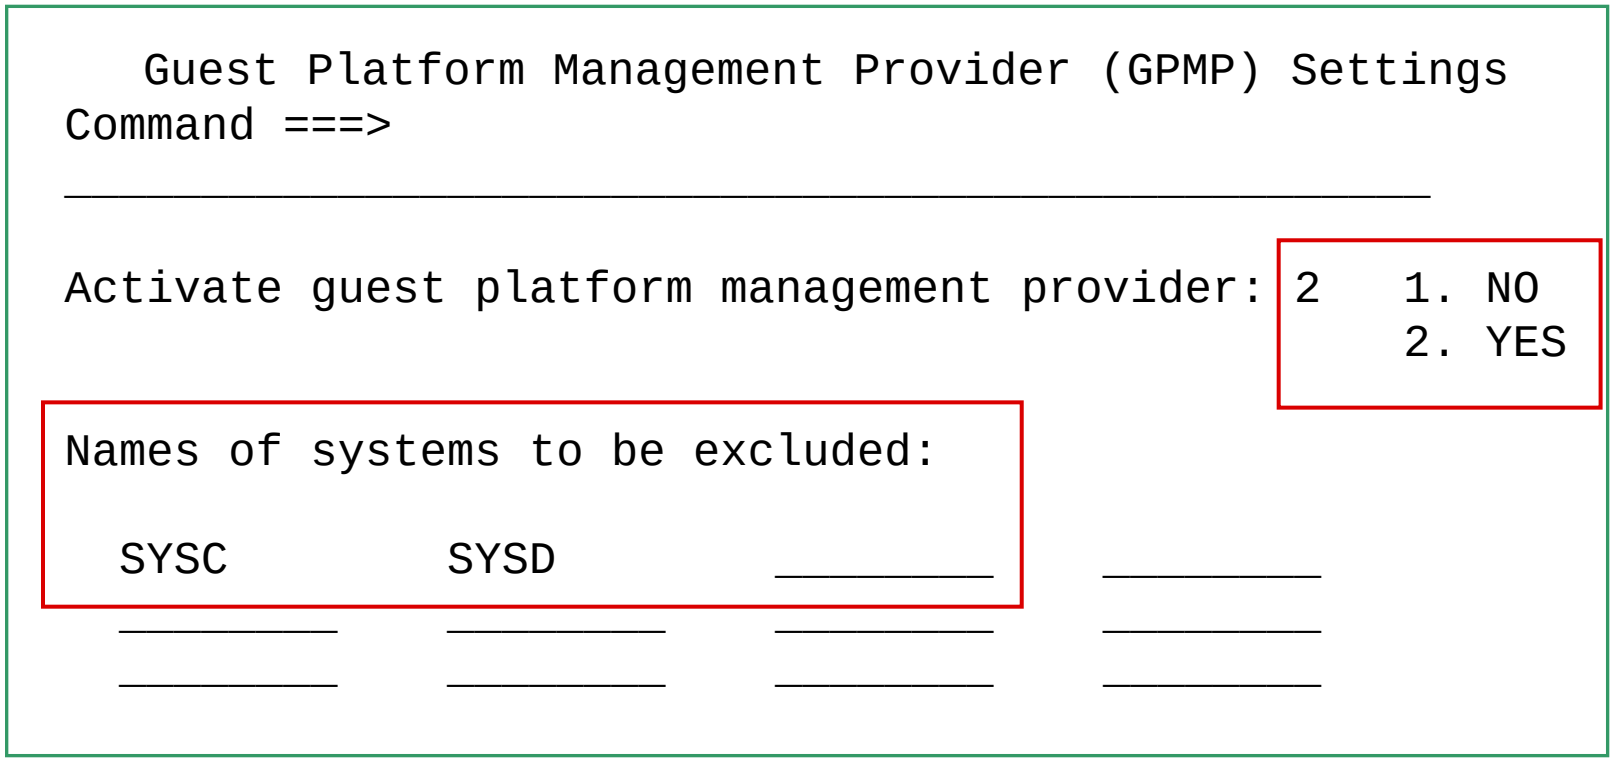

Benefits of GPMP

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

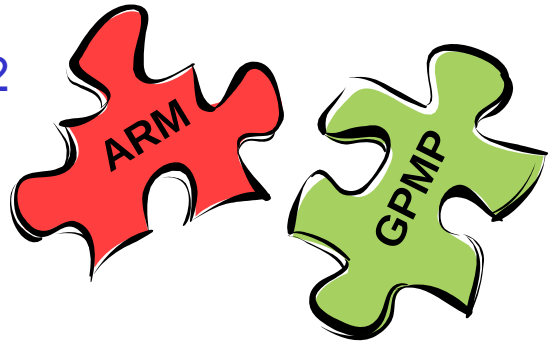

Application Response Measurement

Standards based application instrumentation

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

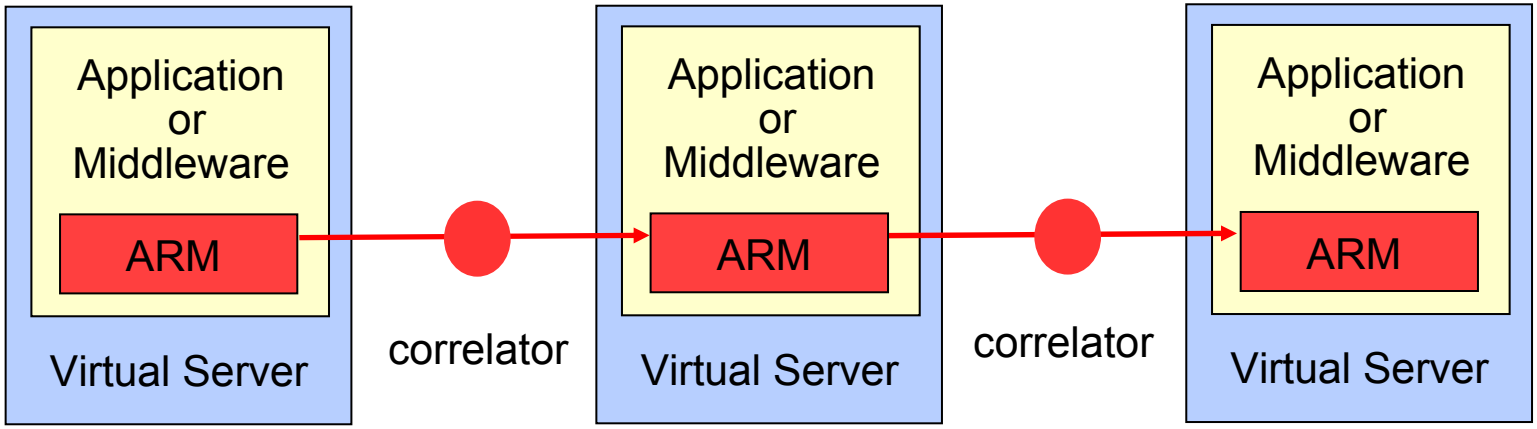

Information in the correlator used by zManager to report

- 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

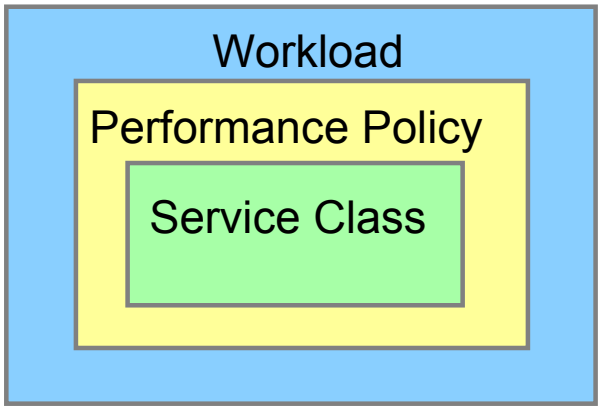

WLM and PPM Terminology

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

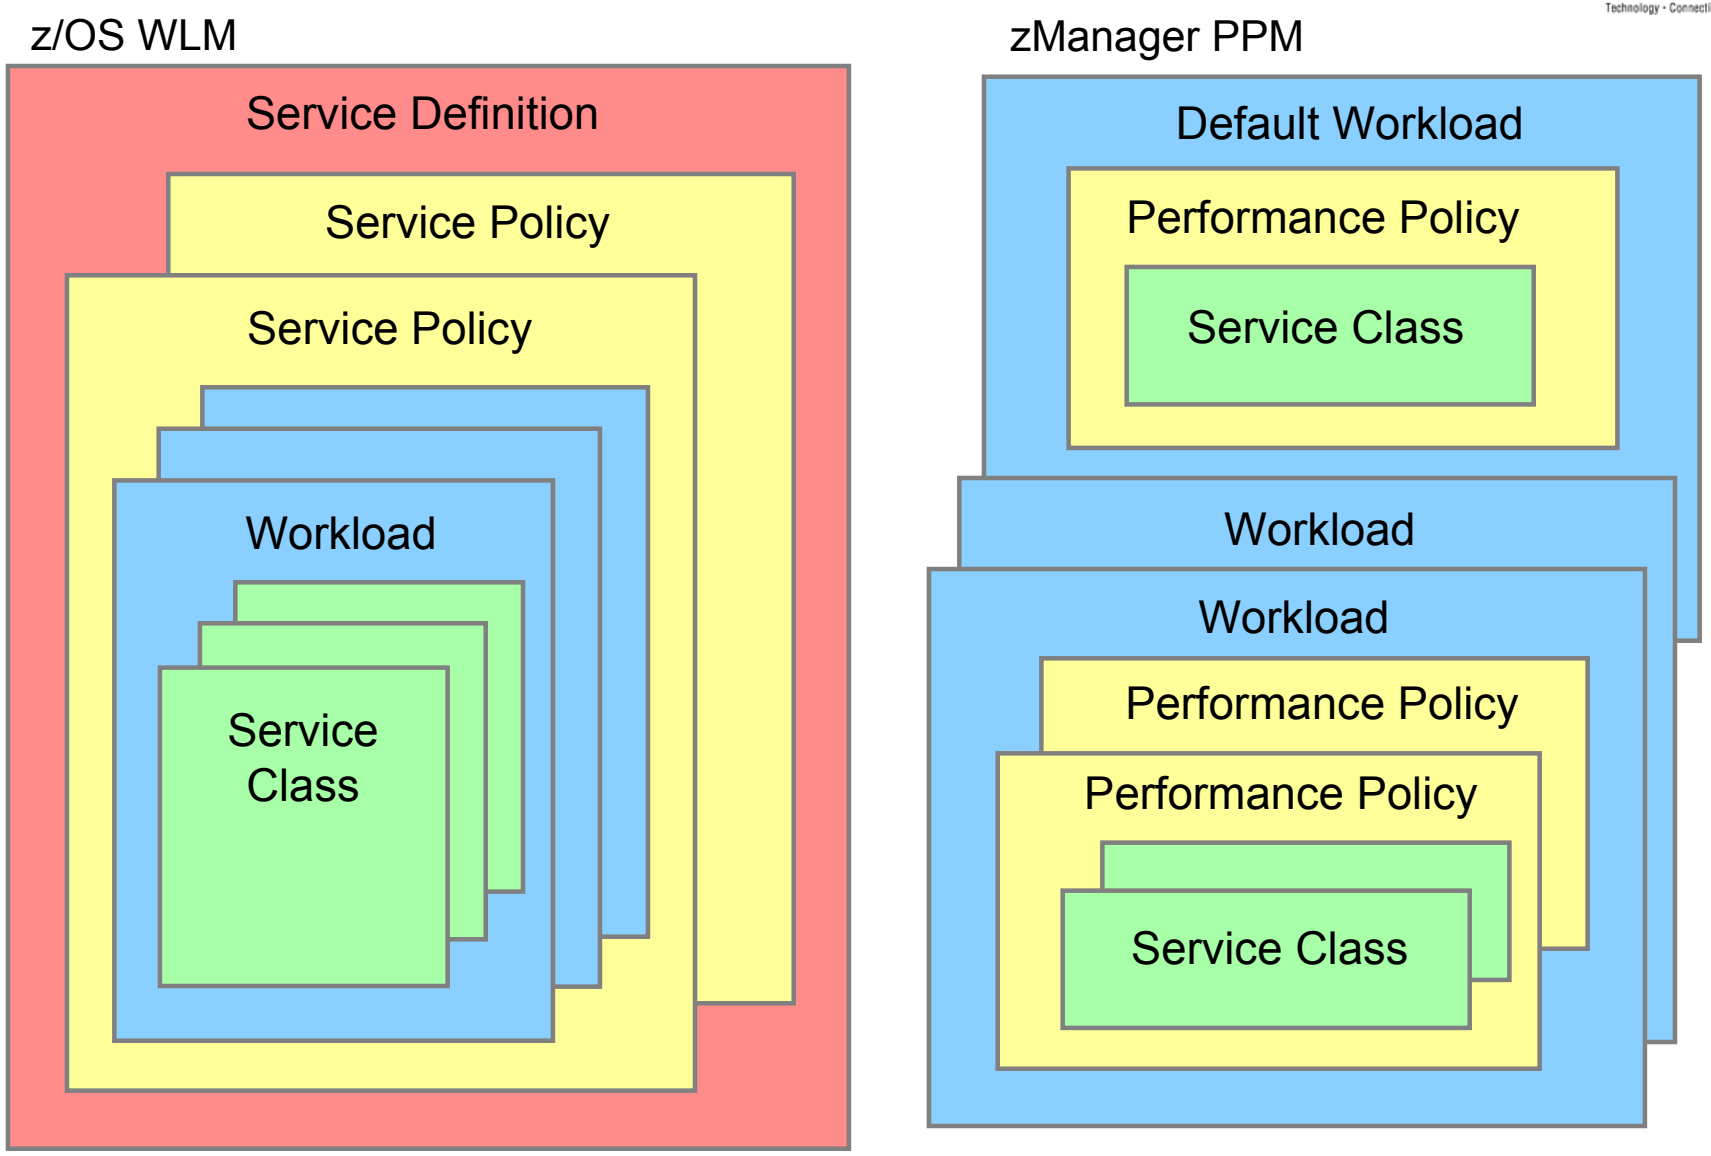

zManager Workload

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

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 relevant within the boundaries of one blade

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

Thoughts for assigning zManager service class names

- Names should be meaningful for installation
- Names based on application
  Banking, WebSales
- Names with service class business importance and velocity
  SrvClsForFastestHighest, SCforModerateLow
- Common names that indicate combinations of performance policy importance and service class importance or velocity

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

Policy Importance	SC Highest -1	SC High -2	SC Medium - 3	SC Low - 4	SC Lowest - 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
BankingProdHighFastestSC	Fastest	Highest	WSCAIX11, WSCAIX12 WSCLNX21, WSCLNX22
BankingTestLowSlowSC	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
BankingProdHighFastestSC	Virtual Server Name Equals	WSCAIX11 or WSCAIX12 WSCLNX21 or WSCLNX22
BankingTestLowSlowSC	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

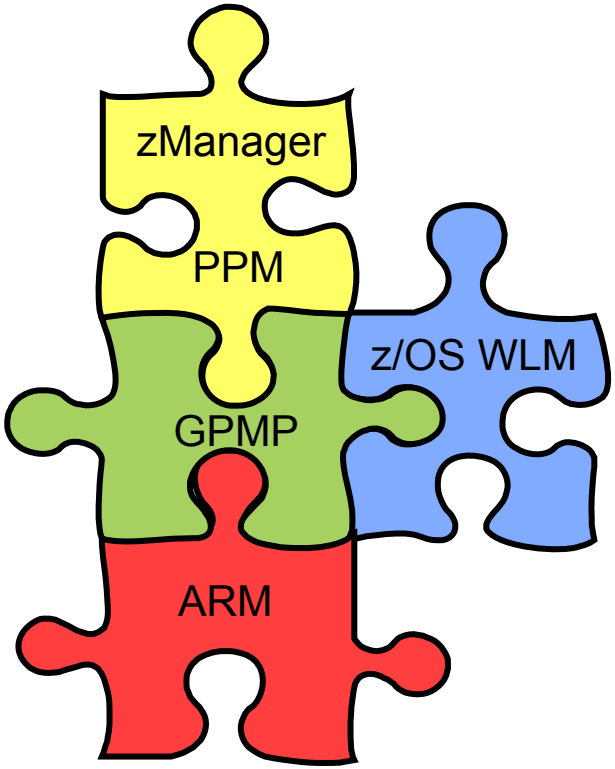

z/OS WLM Classification

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        BankingP 1                      _________    _________
 _____  2    ESC       rodHighF 9                     _________    _________
 _____  3      ESC       astestSC 17                  PPMHGHST     _________
 _____  1  ESC        BankingT 1                      _________    _________
 _____  2    ESC       estLowSl 9                     _________    _________
 _____  3      ESC       owSC     17                  PPMTEST      _________
 _____  1  ESC        Default  1                      PPMDF2       _________
 _____  1  ESC        GROUP21  1                      BANKPROD     _________
 _____  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

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