

Smarter Systems for a  
Smarter Planet

**zEnterprise eXposed!**

**zManager and z/OS Workload Manager**

**Advanced  
Technical  
Skills**

TECHNICAL SALES  
NORTH AMERICA

**Session 10004**

**Bradley Snyder**

**IBM Advanced Technical Skills**



# Trademarks

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

AlphaBlox*	GDPS*	RACF*	Tivoli*
APPN*	HiperSockets	Redbooks*	Tivoli Storage Manager
CICS*	HyperSwap	Resource Link	TotalStorage*
CICS/VSE*	IBM*	RETAIN*	VSE/ESA
Cool Blue	IBM eServer	REXX	VTAM*
DB2*	IBM logo*	RMF	WebSphere*
DFSMS	IMS	S/390*	xSeries*
DFSMSShsm	Language Environment*	Scalable Architecture for Financial Reporting	z9*
DFSMSrmm	Lotus*	Sysplex Timer*	z10
DirMaint	Large System Performance Reference™ (LSPR™)	Systems Director Active Energy Manager	z10 BC
DRDA*	Multiprise*	System/370	z10 EC
DS6000	MVS	System p*	z/Architecture*
DS8000	OMEGAMON*	System Storage	zEnterprise
ECKD	Parallel Sysplex*	System x*	z/OS*
ESCON*	Performance Toolkit for VM	System z	z/VM*
FICON*	PowerPC*	System z9*	z/VSE
FlashCopy*	PR/SM	System z10	zSeries*
	Processor Resource/Systems Manager		

\* Registered trademarks of IBM Corporation

## 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.

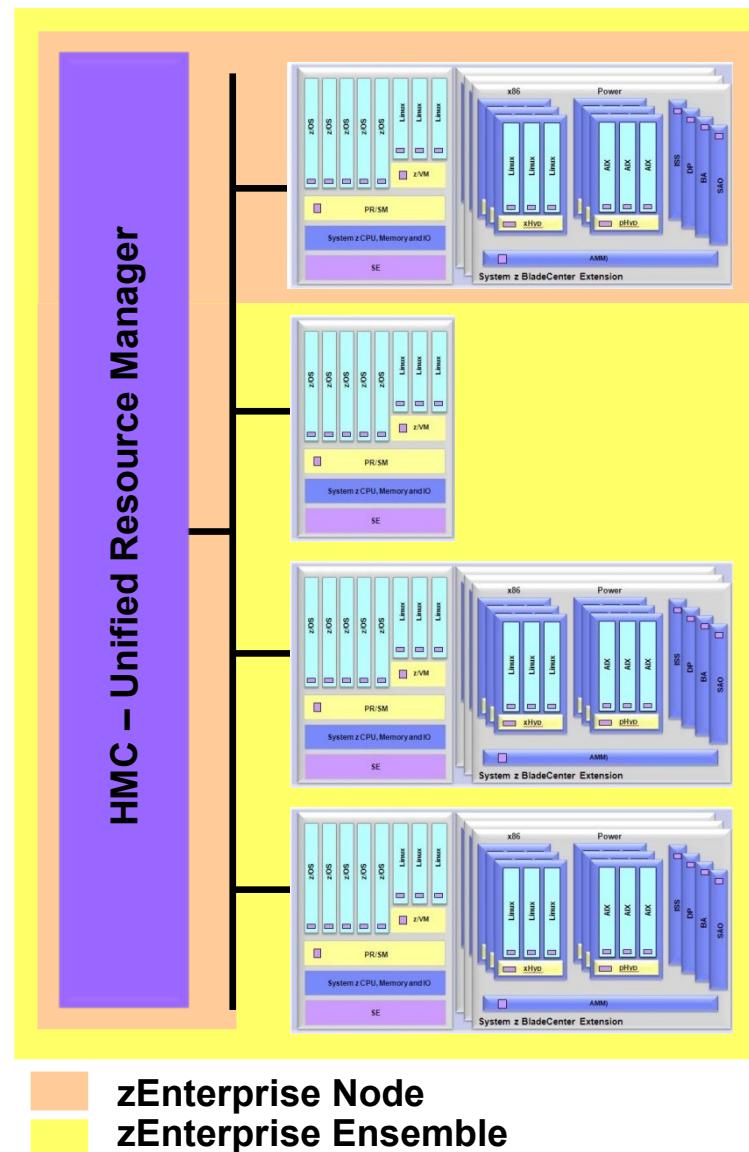
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

# Agenda

- Short zEnterprise Overview
- ARM, GPMP, and Platform Performance Manager
- ARM – Application Response Measurement
- Platform Performance Manager – Performance Policy
  - ▶ Customer example
- z/OS Workload Manager Classification
- Summary

## zEnterprise Ensemble

- **A zEnterprise ensemble is a collection of 1 to 8 nodes managed collectively by the Unified Resource Manager as a single logical virtualized system using the HMC**
- **A zEnterprise node is a z196 CPC with/without zBX.**
  - A zEnterprise node can be a member of at most one ensemble
  - zEnterprise modes are deployed within a single site
- **A zBX is 1 to 4 racks with up to 2 BladeCenters per rack**
- **Z196 CPCs are deployed within a single site**
- **Blade based fit-for-purpose Solutions**
- **Integrated Advanced Virtualization Management**
- **Implements well-defined external interface to Data Center Service Management functions**
- **Virtual Resource Management and Automation**



## ARM GPMP and PPM

- ARM – Application Response Measurement
  - ▶ Provides tracking and measurement of a transaction as it flows across multiple servers
  - ▶ Allows measurement of end-to-end response time, as well as showing how much time was spent on each server in its path
- GPMP – Guest Platform Management Provider
  - ▶ Primarily collects statistics and data provided by ARM, sends data to Support Element and HMC for reporting
- PPM – Platform Performance Management
  - ▶ Define a PPM Performance Policy to measure, monitor, and adjust performance among a group of individual virtual servers
  - ▶ Goals are set in the Performance Policy to define how much CPU resource a virtual server is entitled to, and its relative importance to other virtual servers

# ARM

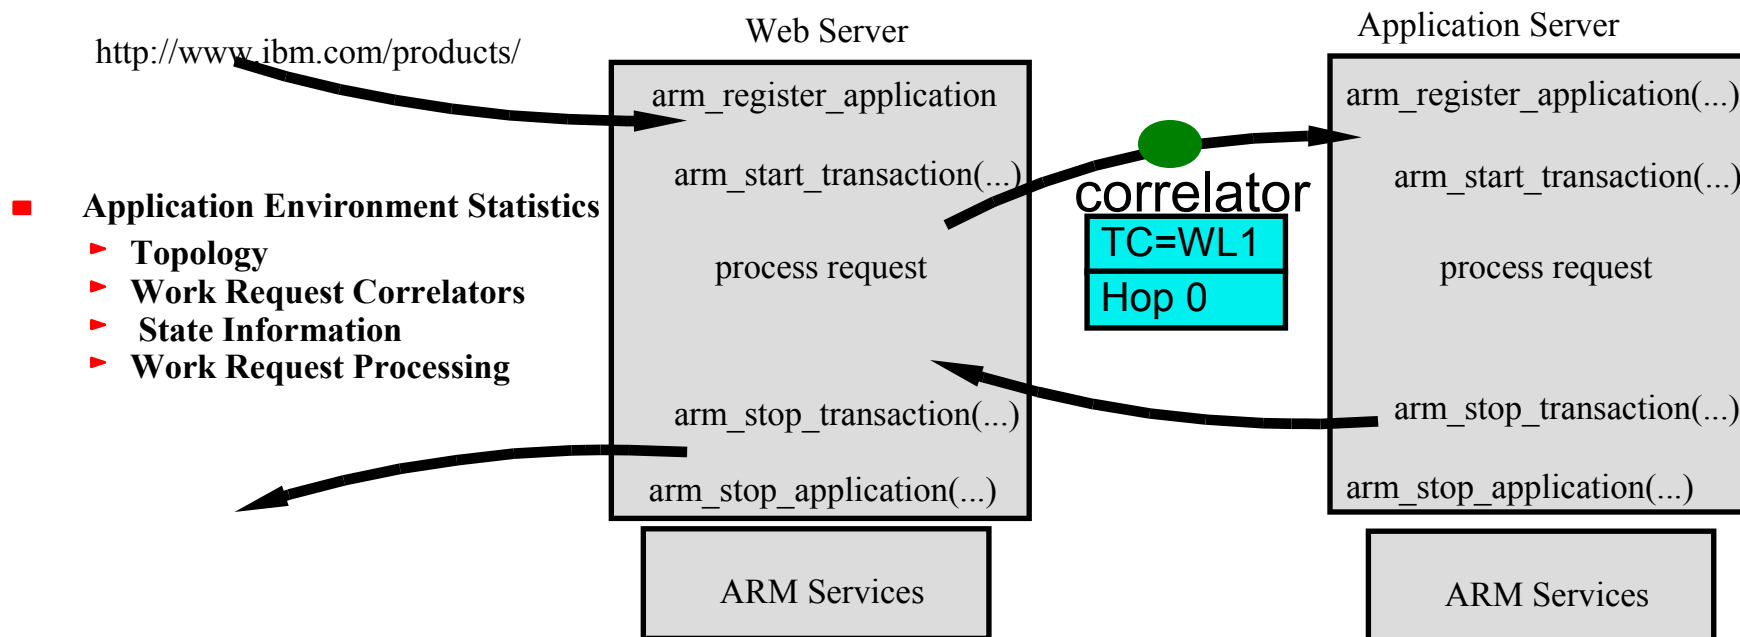
- ARM – Application Response Measurement
  - ▶ Allows measurement of end-to-end response time
    - Correlator assigned to classify work. Correlator is passed to secondary applications and other managed servers that process the application
    - Similar to Performance Blocks in z/OS Workload Manager
- zManager uses information to determine
  - ▶ The amount of time that each application or server used to process the transaction
  - ▶ The name of the application or server that processed the transaction
  - ▶ The end-to-end transaction flow as it moves from one application or server to the next
- Movement of work request from one application OR server to another is considered a 'HOP'
  - ▶ Data viewed from the HOPS Report
- For accurate data, all applications and servers that will process a work request must be ARM enabled

## 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

# Workflow Example with ARM

## Standards Based Application Instrumentation



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




## ARM Instrumented Middleware

- Web Server support:
  - ▶ WebSphere provided plugin
    - IHS/Apache
    - IIS
    - Domino
    - iPlanet
  
- WebSphere Application Server
  - ▶ WAS 6.0, WAS 7.0
  
- DB2
  - ▶ z/OS DB2
  - ▶ DB2 Universal Database for Linux/Unix/Windows

## Platform Performance Manager

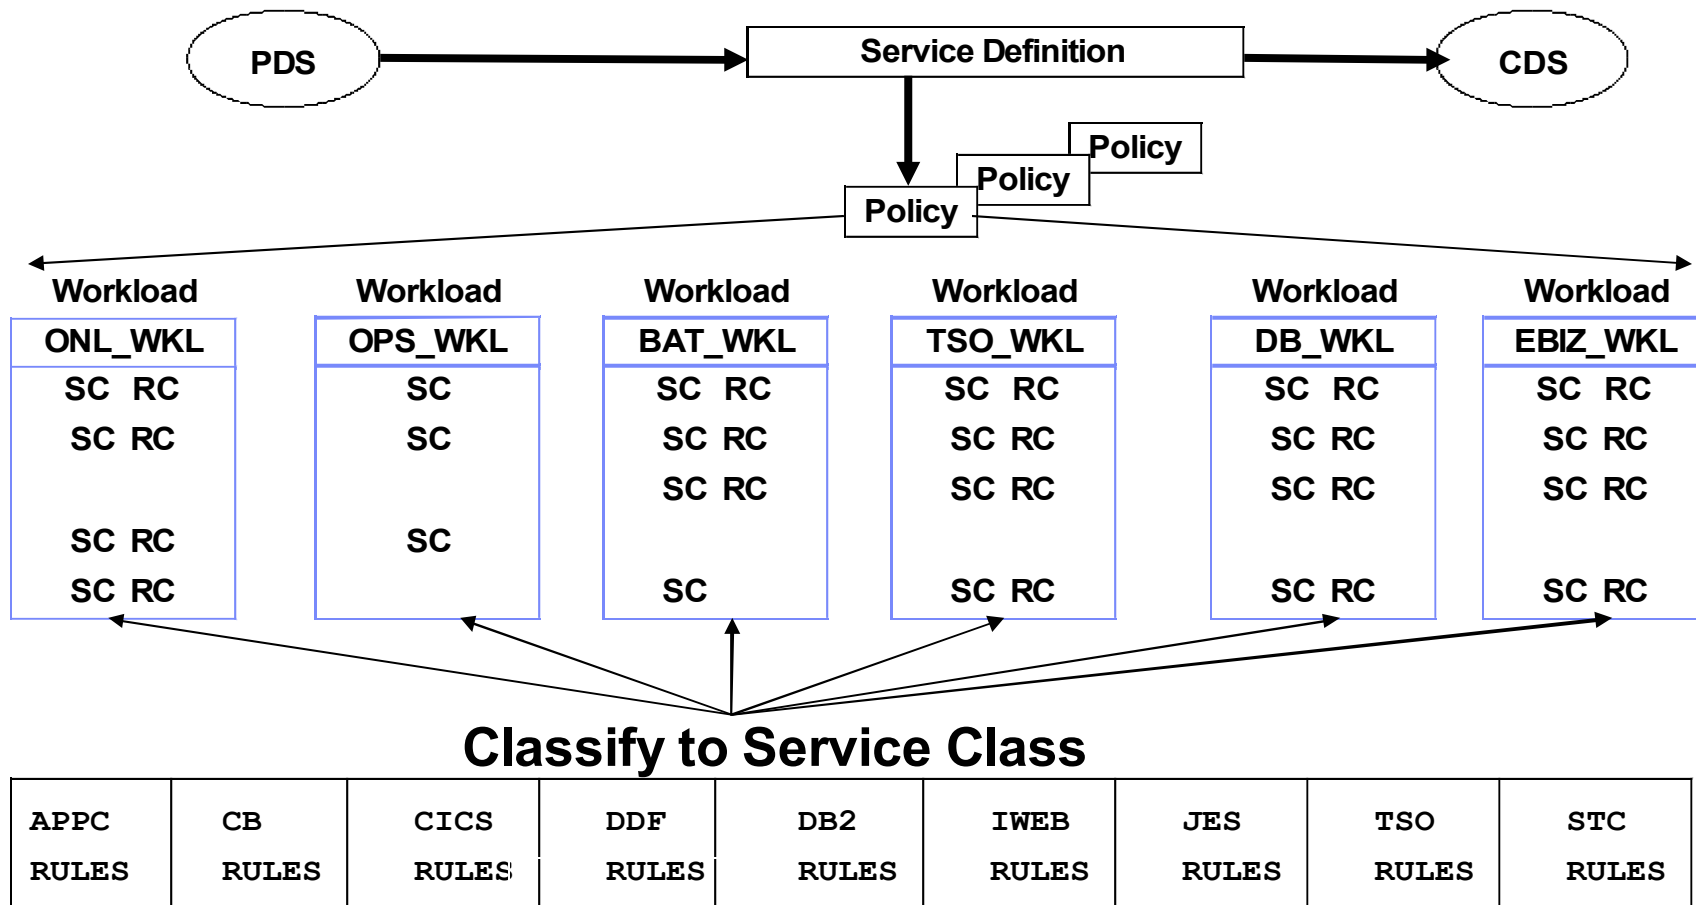
- Component of zManager
- Monitors and manages performance of virtual servers in the zEnterprise
- Based on a workload performance policy, PPM will adjust CPU resources between those virtual servers that share resources



## z/OS Workload Manager and PPM Terminology

- Both z/OS use the same terms, but will mean different things
- z/OS
  - ▶ Workload Manager has a Service Definition
  - ▶ Service Definition has an active POLICY, which is the container for all Workloads and Service Classes
    - Definition can have multiple Policies, only one active
  - ▶ As work comes into the system, it is classified to a Service Class
  - ▶ Service Classes can be grouped together in a Workload for reporting purposes
- PPM
  - ▶ Workload Performance Policy
  - ▶ First definition is to define one or more virtual servers to a workload
  - ▶ Within the workload, define one or more service classes and assign the virtual servers to the appropriate service class


# z/OS Workload Manager



**Work Enters the System**

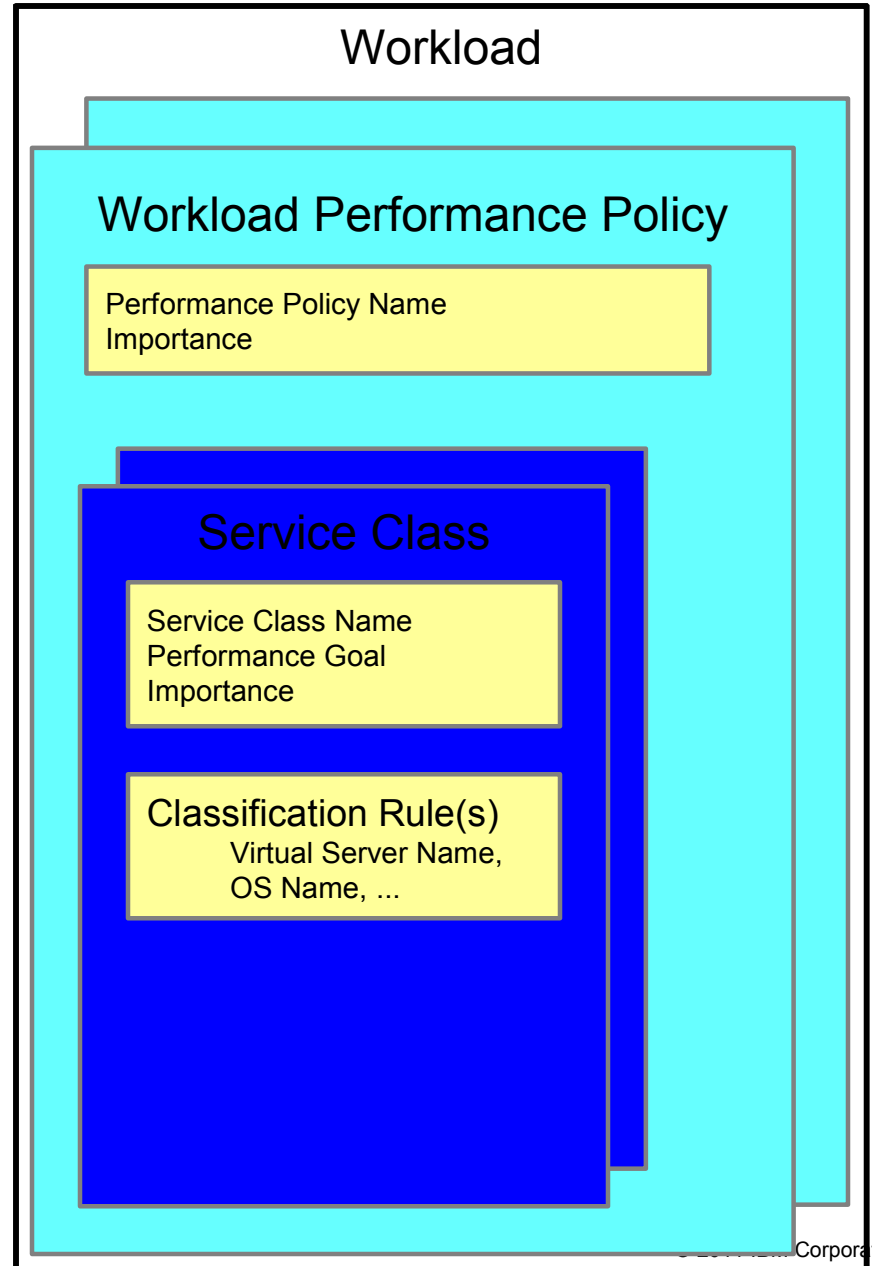
## Workload Performance Policy

- 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:
  - ▶ Multiple performance policies associated with a workload
  - ▶ A single policy is active at a given time
  - ▶ Can dynamically change the policy that is active
    - Through the UI
    - Through a time-based schedule
    - Example: Day shift / night shift policy



# zManager Performance Policy

- Define one or more virtual servers to a workload
- A workload can have one or more performance policies to describe its business importance and objectives
- Only one Policy is active at any time
  - ▶ Every new workload has at least the 'Default' performance Policy
  - ▶ Additional policies may be defined as needed
- Each performance policy has one or more service classes that set the priority of and classify the virtual servers
  - ▶ zManager uses the active performance policy and its service classes to determine how physical resources are applied to the virtual servers



## PPM Workload Service Class Goals

- Every Workload has a business importance level (Highest to Lowest) to manage resources among different workloads running on same physical server
- Every Service Class has a business importance level (Highest to Lowest) to help manage resource among different Service Classes running on the same physical Server
  - ▶ Therefore, at most can only have 25 combinations
  - ▶ For simplification, use common names for a specific combination

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

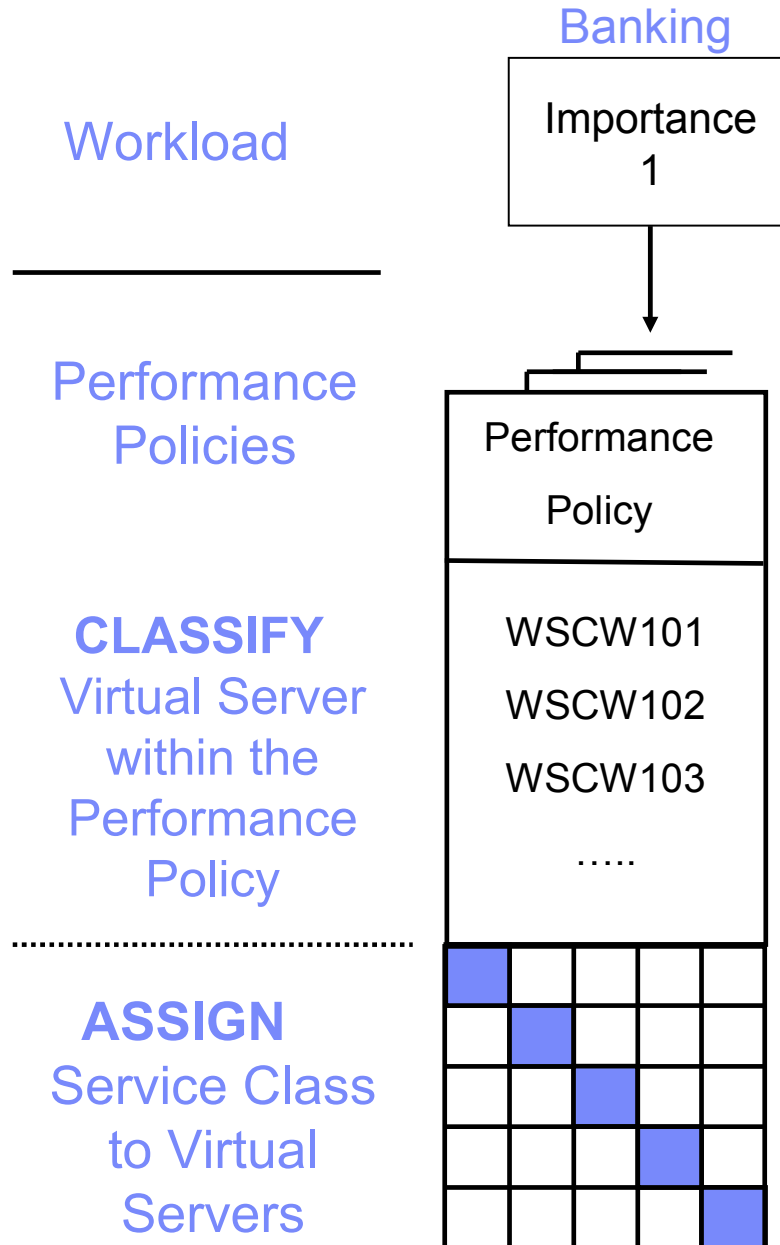
## PPM Workflow

The following steps will be taken in order when defining work with PPM

1. Define a workload with a name, ie. Banking
2. Define the virtual servers in zEnterprise and assign to the workload
3. Define a performance policy to manager the virtual servers, or edit the default performance policy
4. Determine which virtual servers will belong to which service class
5. Define necessary service classes and assign virtual servers to appropriate service class



# Graphical Example



## Customer Example

- Customer has a new banking workload to be implemented on a zEnterprise zBX system
  - ▶ Step 1 define workload called “BANKING” on HMC
- Sizing has already been done to determine best placement of each virtual server
  - ▶ Step 2 define virtual servers below on the HMC and assign them to the BANKING workload

zBX Node	zBX Chassis	Blade Number	Function	Virtual Server
Node1	B.1	1	HTTP	WSCW101
		2	HTTP	WSCW102
		3	HTTP	WSCW103
		4	HTTP	WSCW104
		5	WAS	ATAIX100
			WAS	ATTAIX10
			WAS	ATPAIX20
		6	WAS	ATAIX200
			WAS	ARPTX100
			WAS	ATMIX304

## Performance Policy Service Class Classifications

- Step 3 – Define new policy in BANKING workload called ‘STANDARD’
- Step 4 – Determine which virtual servers belong to which service class
  - ▶ Importance and Velocity Goals are only relevant within the boundaries of one blade. No need to differentiate GROUP22 on different blades
  - ▶ Part 1 – Blades with a single virtual server. Assign virtual server to Group11
    - Protects performance of virtual server in the event additional servers are added to blade
    - Can change group if necessary
  - ▶ Part 2 – Set relative importance/velocity for each virtual server on each blade
    - May need to reassign service class if virtual server is moved from one blade to another

## Service Class Assignments

- Using previous Rules, Service Class Assignments are below
  - ▶ Internal business discussions will need to be completed in order to determine relative importance of virtual servers that are sharing resources

zBX Node	zBX Chassis	Blade Number	Virtual Server	Service Class
Node1	B.1	1	WSCW101	GROUP11
		2	WSCW102	GROUP11
		3	WSCW103	GROUP11
		4	WSCW104	GROUP11
		5	ATAIX100	GROUP11
			ATTAIX10	GROUP55
			ATPAIX20	GROUP33
		6	ATAIX200	GROUP11
			ARPTX100	GROUP33
			ATMIX304	GROUP55

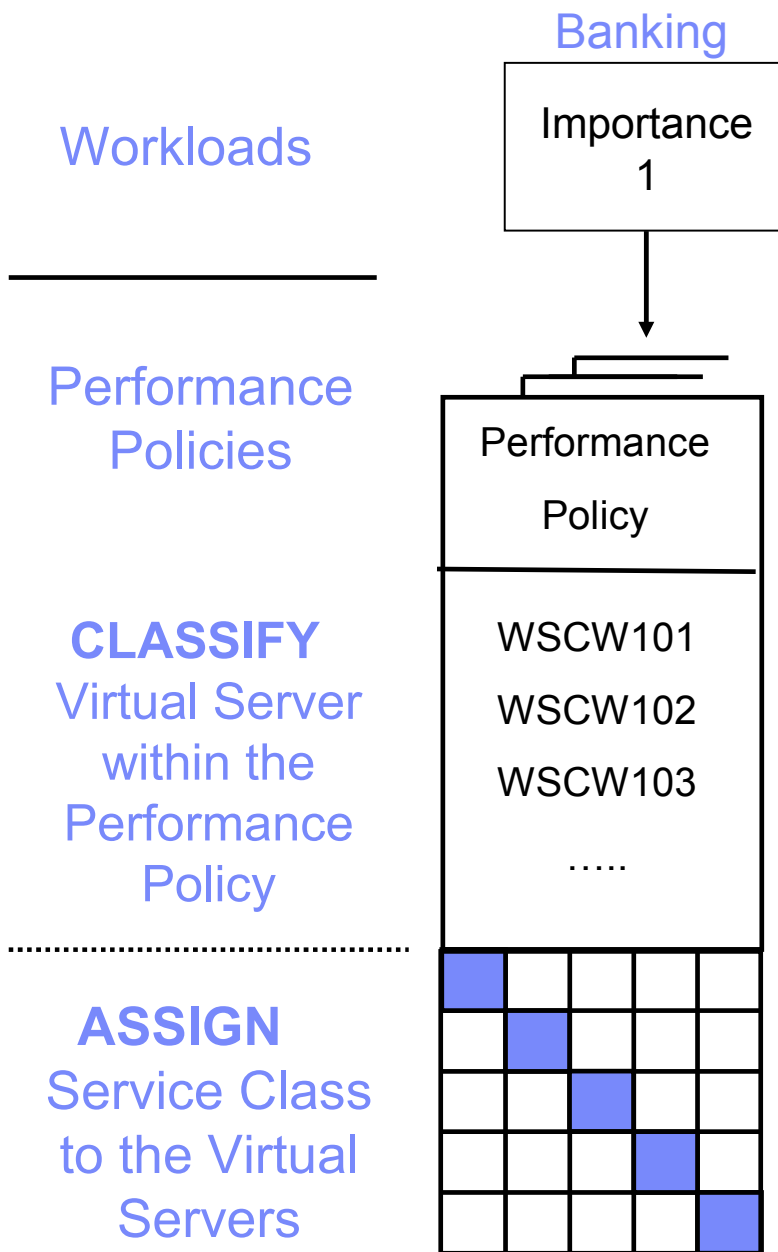
## Grouping Servers by Service Class

- Step 5 - Define necessary service classes and assign virtual servers to appropriate service class

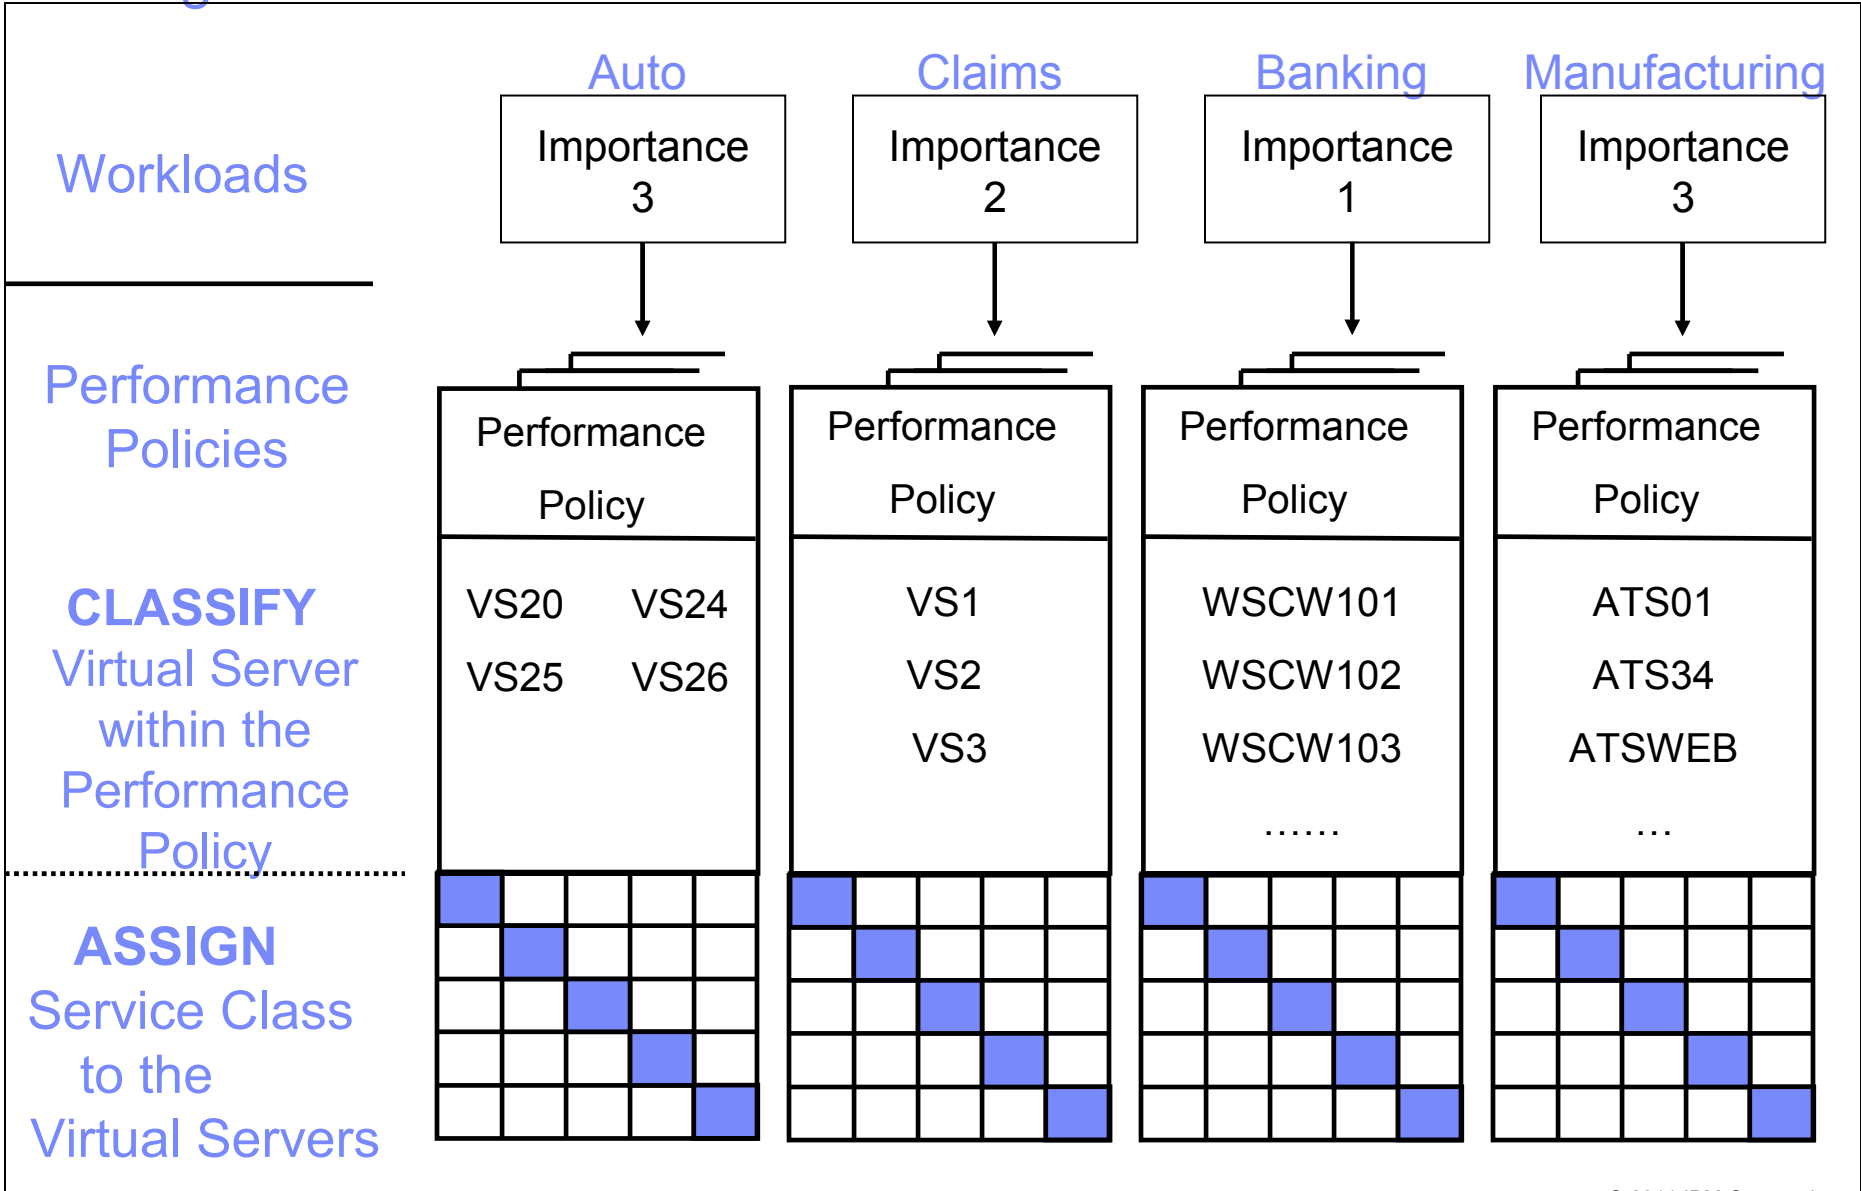
Workload	Policy	Service Class	Virtual Server
Banking	Standard	Group11	WSCW101, WSCW102, WSCW103, WSCW104, ATAIX100, ATAIX200
		Group33	ATPAIX20, ARPTX100
		Group55	ATTAIX10, ATMIX304

- Optional Additional Step – Repeat steps 3 through 5 for additional policies if different performance goals are needed at different times
  - ▶ I.e. define policy 'Weekend' if test servers can get high importance and velocity goals then

# Example



# Adding Additional Workloads



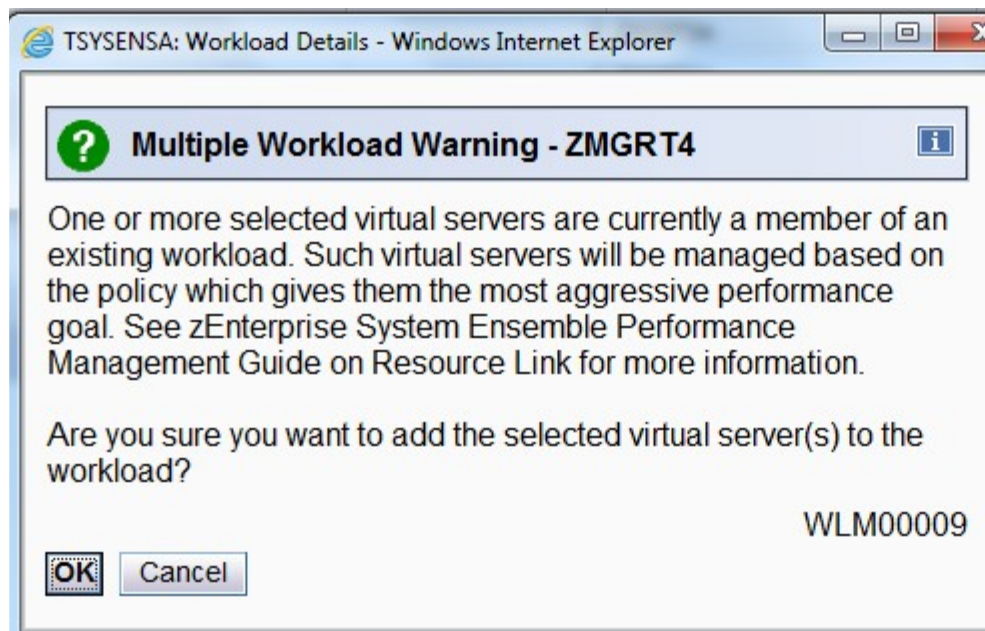
## ARM and the Performance Policy

- The PPM Performance Policy is only managing the velocity of the virtual servers
- ARM provides information only for overall transaction response times
  - ▶ Cannot specify any response time goal for transactions
- If transactions are running too long, and CPU is the constraint, then velocity goals of individual servers will need to be adjusted
  
- Similarity in z/OS to having response time goals for CICS/IMS transactions, but managing the regions to the region velocity goal



## Assigning Virtual Servers to Multiple Workloads

- It is possible to assign a virtual server to multiple workloads
- Server will be managed to most aggressive goal
- Example – Web server acts as front end to multiple applications



## z/OS Workload Manager Classification

- Work coming into z/OS from within the zEnterprise CAN be classified using EWLM rules
  - ▶ Requires ARM and GPMP to be implemented since the correlator contains information needed for classification
- Uses name of PPM service class for classification
  - ▶ PPM workload and service class names can be up to 64 characters long
    - Is case sensitive
  - ▶ z/OS WLM will use first 32 characters only for classification
    - Will need to use sub-rules when service class names are longer than 8 characters
    - Be sure FOLD Qualifier names is set to NO (makes naming case sensitive)
    - EWLM Service Class must be a single period response time goal
- PPM cannot make any changes to z/OS environment. z/OS WLM Service Class goal will need to be set appropriately
  - ▶ ARM and GPMP allow reports to show how much time is being spent in z/OS as well as rest of zEnterprise

## z/OS WLM Classification Example

- Done either through WLM ISPF application, or through z/OS MF
- From WLM Classification rules, edit EWLM rules

```

-----
                          Subsystem Type Selection List for Rules          Row 1 to 15 of 18
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  _
  3_    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    _

```

## z/OS Classification Rules Cont.

- ESC is the only Qualifier Type
- Qualifier name is name of performance policy Service Class
  - ▶ Sub-rules needed for Service Classes longer than eight characters
- As with all classification rules, a default can be used
  - ▶ May not want to use if need to let incoming work use other (ie. DDF) classification rules

```

                Modify Rules for the Subsystem Type                Row 1 to 8 of 14
Command ==> _____ Scroll ==> CSR

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-----
Action  Type      Name      Start      Service      Report
-----
_____ 1  ESC      GROUP11  1          PPMDEFLT   _____
                              DDFPROD   _____
    
```

## z/OS Classification Rules with Multiple Workloads

- Possible for multiple workloads to have applications that access z/OS but have different performance requirements on z/OS
- For example Banking and Manufacturing both access DB2 on z/OS, and transactions are coming from servers that are assigned to GROUP33
- Need to differentiate between GROUP33 for Banking and GROUP33 for Manufacturing
- In this case, add suffix to group names to allow for differentiation between workloads
  - ▶ GROUP33-BANKING
  - ▶ GROUP33-MANUFACTURING

# Classification Rules

- Classification Rules with Suffixes

```

Modify Rules for the Subsystem Type          Row 1 to 8 of 16
Command ==> _____ Scroll ==> CSR

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-----
Action  Type      Name      Start
-----
_____ 1  ESC      GROUP11  1
_____ 1  ESC      GROUP33- 1
_____ 2   ESC      BANKING  9
_____ 1  ESC      GROUP33- 1
_____ 2   ESC      MANUFACT 9

-----Class-----
Service      Report
-----
DEFAULTS: PPMDEFLT _____
           DDFPROD  _____
           DDFPROD  _____
           _____ _____
           GRP33MAN  _____
    
```

## Summary

- ARM allows for the reporting of application transaction performance
- Platform Performance Management uses a workload policy to monitor and manage the performance of virtual servers
- Combination of the two allow for excellent reporting of response time and performance of applications running on zEnterprise
  
- z/OS Workload Manager can use performance policy classifications to classify incoming work in z/OS
  - ▶ If no EWLM classification rules match, then will default to standard classification rules

# Thank You