

# zManager: Platform Performance Manager

Hiren Shah  
IBM  
March 14, 2012  
10658



# 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](http://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.

IBM products may be trademarks or registered trademarks of their respective companies.

Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will achieve will vary depending on many 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, individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM 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 performance characteristics 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:



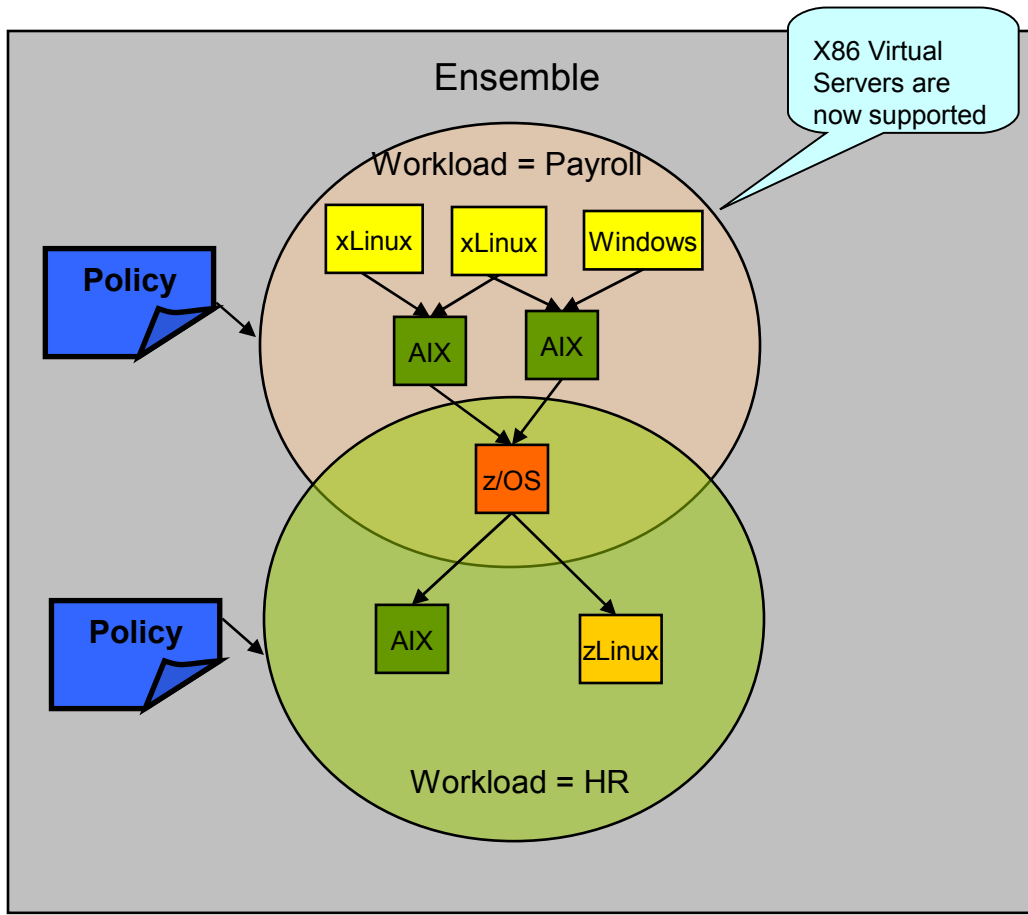
<b>1</b>	<b>Platform Performance Manager Overview</b>
<b>2</b>	<b>Workload Monitoring</b>
<b>3</b>	<b>PPM GPMP Support</b>
<b>4</b>	<b>WLM and PPM relationship</b>
<b>5</b>	<b>Dynamic Processor Management</b>
<b>6</b>	<b>SASP Load Balancing Support</b>

# zEnterprise Platform Performance Manager

- Platform management component responsible for goal-oriented resource monitoring, management, and reporting across the zEnterprise Ensemble
- 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
- Load Balancing recommendation to SASP enabled switches/routers
- Functionality integrated into the Unified Resource Manager

# Platform Workload

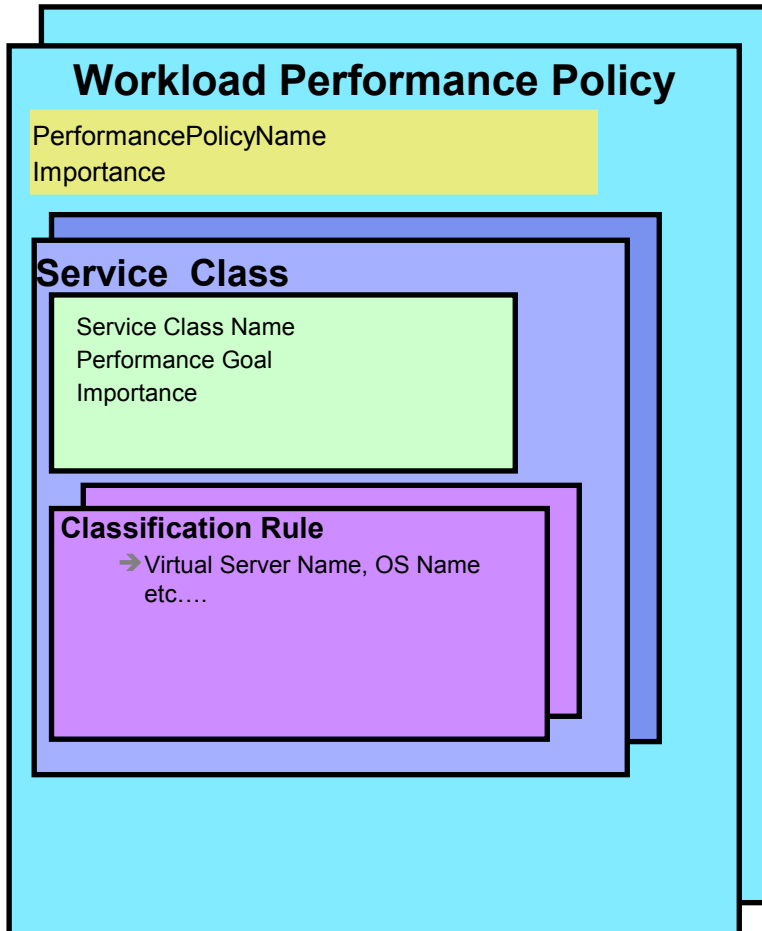
- **A Platform Workload is a grouping mechanism and “management view” of virtual servers supporting a business function**
- **Provides the context within which associated platform resources are presented, monitored, reported, and managed**
- **Management policies are associated to Platform Workload**
  - 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



**SHARE**  
Technology • Connections • Results

# 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 Monitor Dashboard 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
- Reporting data can be extracted by zManager APIs, for example
  - `POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-groups-report`

# Workload Based Monitoring and Reporting



- Display of current data and fairly recent history
  - Up to 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

[Welcome](#)  
[Systems Management](#)  
[Ensemble Management](#)  
   [Ravenclaw](#)  
     [Members](#)  
     [Workloads](#)  
[HMC Management](#)  
[Service Management](#)  
[Tasks Index](#)

Ensemble Management > Ravenclaw > Workloads

Workloads | Topology

Filter:  Tasks Views

Select	Name	Virtual Servers	Performance Policy	Performance Policy Status	Performance Policy Business Importance
<input checked="" type="checkbox"/>	Bookstore Workload	37	Bookstore Policy	Active	Highest
<input type="checkbox"/>	Default	104	Default	Active	Medium
<input type="checkbox"/>	GPMLinuxStress	47	Default	Active	Medium
<input type="checkbox"/>	GSSH17Cluster	6	GSSH17Cluster-Policy	Active	Highest
<input type="checkbox"/>	GSSH24	1	POLGSSH24	Active	High
<input type="checkbox"/>	GSSP15 and Friends	1	GSSP15-Policy	Active	High
<input type="checkbox"/>	GSSPlow	2	GSSPlow-Policy	Active	Low
<input type="checkbox"/>	LinuxTrade	6	OnlineTrades	Active	Medium
<input type="checkbox"/>	mark0005	3	test	Active	High
<input type="checkbox"/>	mark0007	2	suselinux	Active	Medium
<input type="checkbox"/>	mark0009	1	Default	Active	Medium
<input type="checkbox"/>	MIXOS	9	POMIXOS	Active	High

Max Page Size: 500 Total: 30 Filtered: 30 Selected: 1

Tasks: Bookstore Workload

Workload Details

- Daily
- Operational Customization
- Configuration
  - Delete Workload
  - New Performance Policy
  - New Workload
- Monitor
  - Service Classes Report
  - Virtual Servers Report
  - Workload Resource Adjustments Report
  - Workloads Report**

# 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

Report Interval: Starting 3/2/11 11:05:51 PM for 15 minutes (3/2/11 11:20:51 PM) [Modify](#)

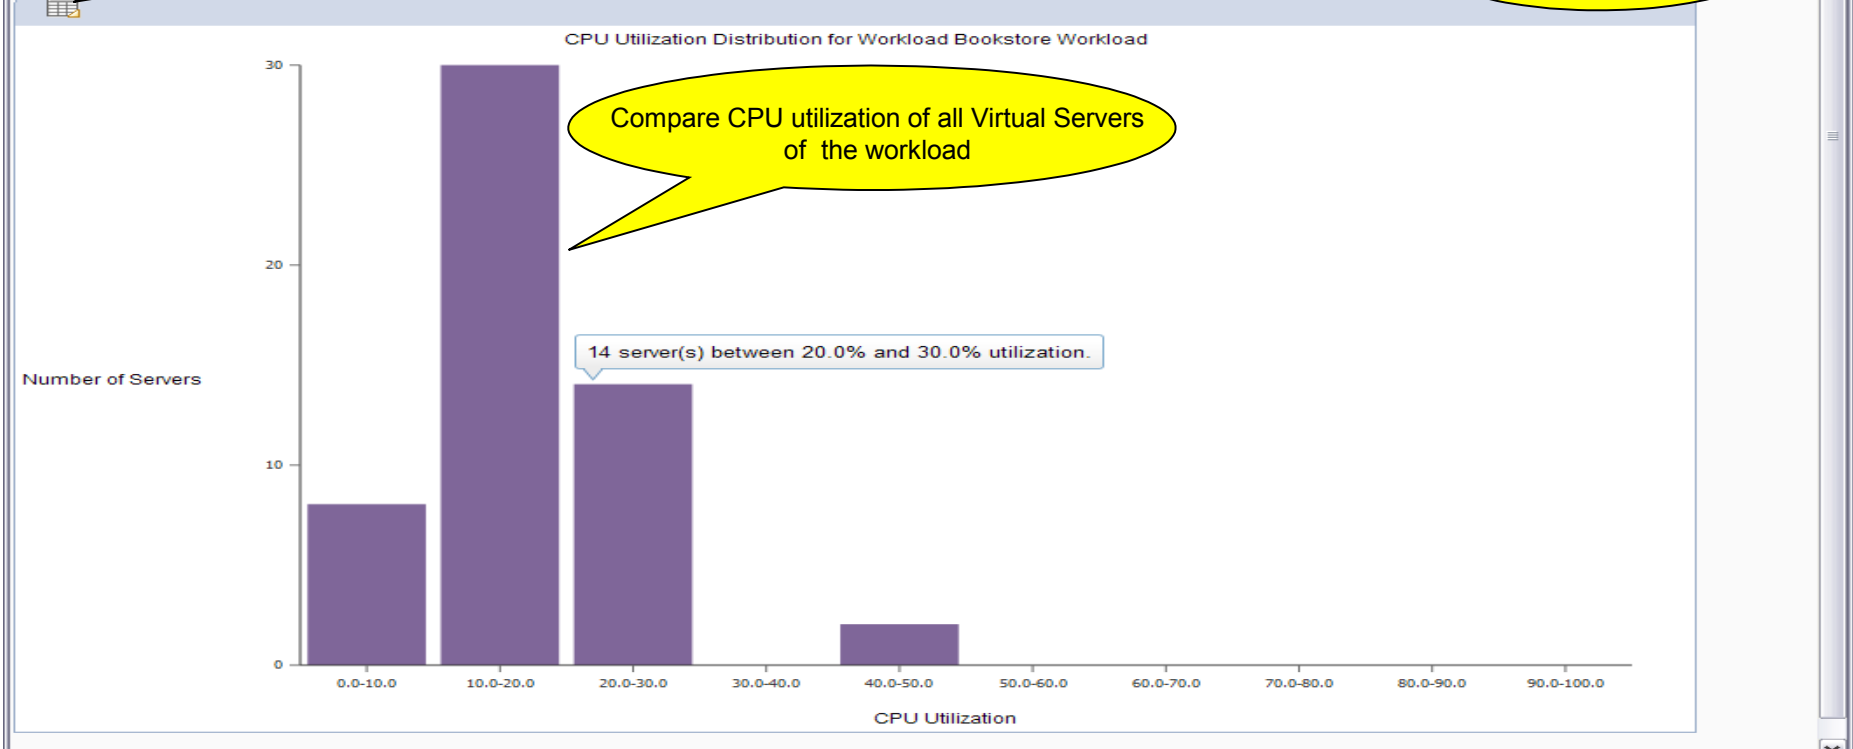
Select	Workload	Service Class With Largest PI (PI)	Performance Policy
<input type="radio"/>	AIXBL5AND7	SCAIXB5and7 (3.52)	PoAIXB5and7
<input type="radio"/>	AIXSASPWKLD	AIXSASP_SC (0.70)	AIXSASP_Policy
<input type="radio"/>	Blade 6 workload	v1_v2 SC (1.25)	Blade 6 policy
<input type="radio"/>	Blade 9 10 11 workload	High priority SC (1.00)	Cross Blade Center Policy
<input type="radio"/>	BladeC.1.4 SASP	SASP SC Gold (1.00)	SASP Policy
<input type="radio"/>	BladeCenter2 Workload	Highest SC (1.55)	BC2 Policy
<input checked="" type="radio"/>	Bookstore Workload	A-Team (2.04)	Bookstore Policy
<input type="radio"/>	Default	Default (0.72)	Default

Spread sheet of monitoring data with "Export" support

Workload health overview

Active Performance Policy

Workload Charts  
 Charts: CPU Utilization



Compare CPU utilization of all Virtual Servers of the workload

# 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

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	Service Class	Performance Policy	Current Performance	PI	Goal	Business importance
<input checked="" type="radio"/>	A-Team	Bookstore Policy	Moderate	1.90	Velocity - Fast	High
<input type="radio"/>	B-Team	Bookstore Policy	Slow	1.46	Velocity - Moderate	Medium
<input type="radio"/>	BuyNode	Bookstore Policy			Velocity - Fastest	Highest
<input type="radio"/>	C-Team	Bookstore Policy	Slow	1.00	Velocity - Slow	Low
<input type="radio"/>	Dead Meat	Bookstore Policy	Slow	0.37	Velocity - Slowest	Lowest
<input type="radio"/>	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



A-Team - Time: 2/26/11 20:14:30, PI: 2.13

# 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
<input type="radio"/>	GSSF26-27-28	<input checked="" type="checkbox"/>
<input checked="" type="radio"/>	GSSH18Cluster	<input checked="" type="checkbox"/>
<input type="radio"/>	GSSP17	<input checked="" type="checkbox"/>
<input type="radio"/>	GSSP17-J80	<input checked="" type="checkbox"/>
<input type="radio"/>	LinuxDaytraderZGF	<input checked="" type="checkbox"/>
<input type="radio"/>	LinuxGA1toGA2	<input checked="" type="checkbox"/>

Service Class(es):

Select	Service Class	Available
<input type="checkbox"/>	Default:Default	<input checked="" type="checkbox"/>
<input type="checkbox"/>	GSSH18POL:Default	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	GSSH18POL:GSSH18SC	<input checked="" type="checkbox"/>
<input type="checkbox"/>	GSSH18POL:GSSH18SLOW	<input checked="" type="checkbox"/>

Alert settings based on Service level impact

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



SHARE  
Technology • Connections • Results

# HMC1: Service Classes Report - Mozilla Firefox

http://9.60.31.154:8080/hmc/wcl/T1615

## Service Classes Report - Weinheimer Agriculture Parts

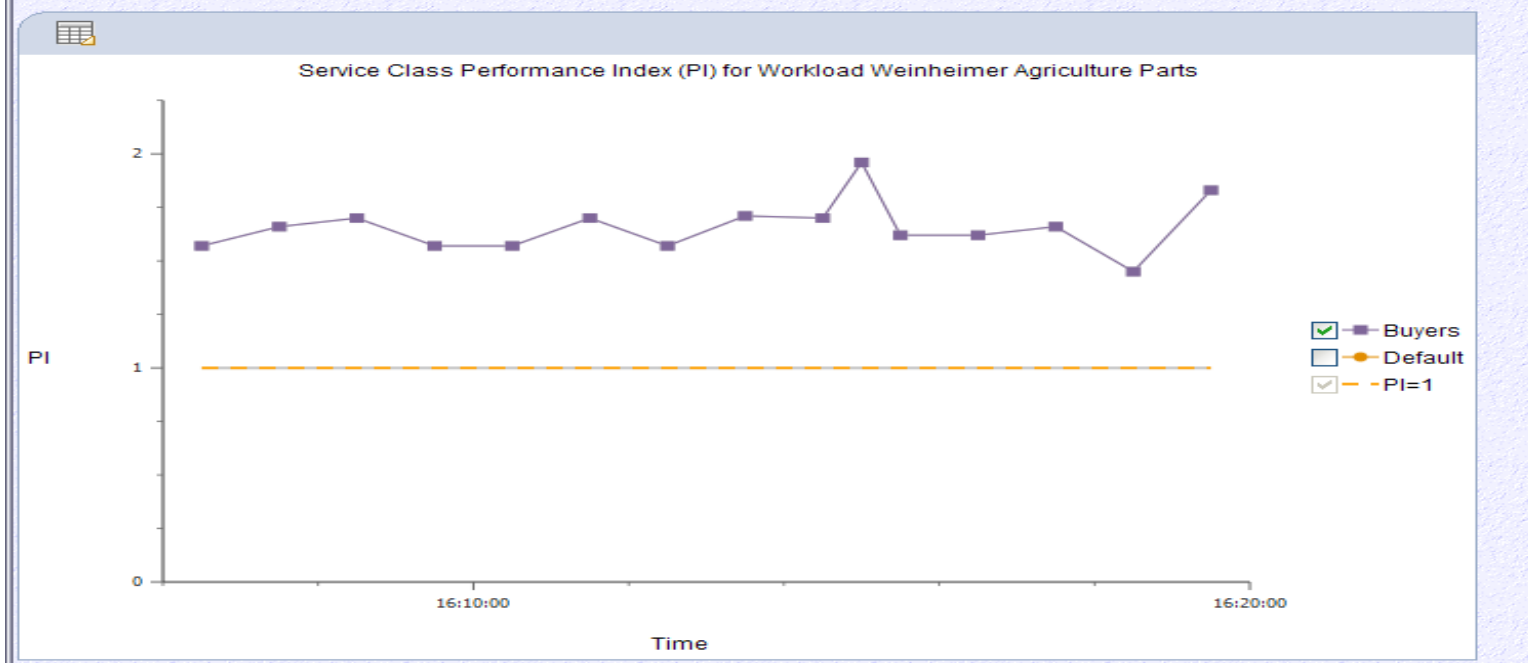
Report Interval: Last 15 minutes [Modify](#)

Select	Service Class	Performance Policy	Current Performance	PI	Alert	Business Importance
<input checked="" type="radio"/>	Buyers	Peak Period	Fast	1.65	Velocity - Fastest	Highest
<input type="radio"/>	Default	Peak Period	Fast	0.77	Velocity - Moderate	Medium

Total: 2 Filtered: 2 Selected: 1

Service Class Alert was triggered

### Service Class Charts



Close Help

Done

# 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

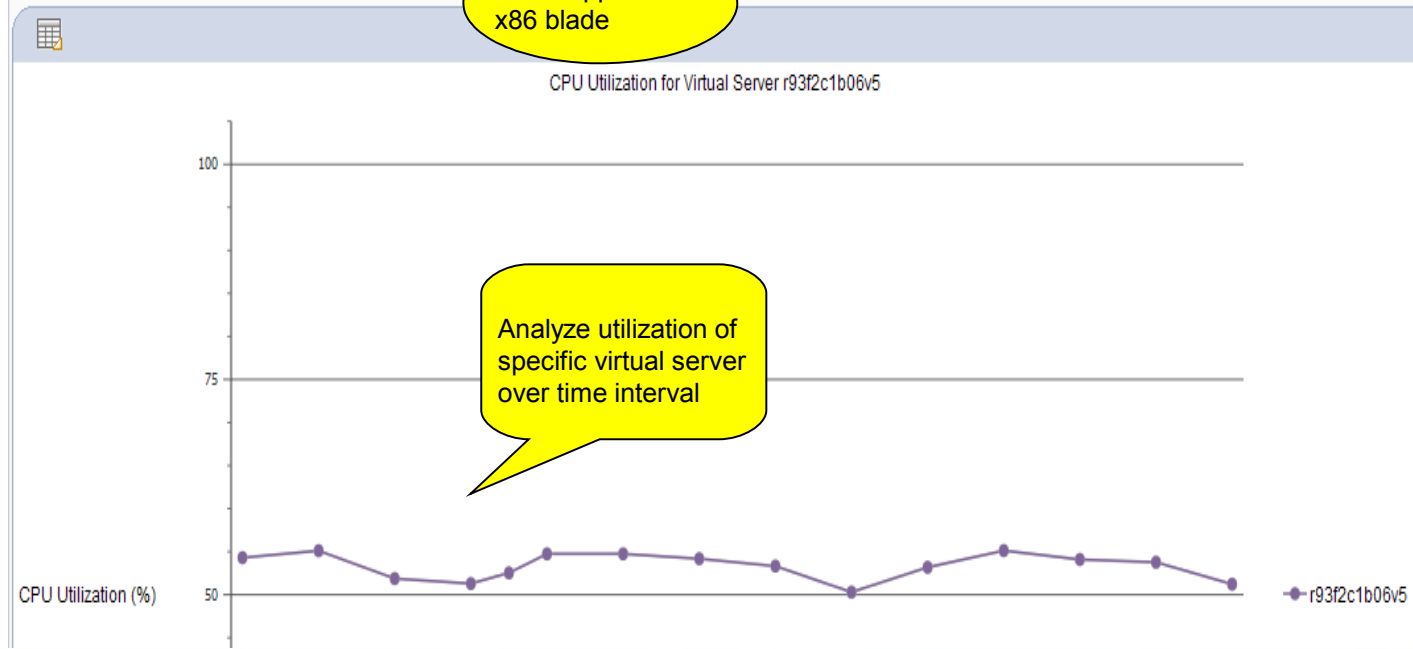
Resource utilized by VS

Select	Virtual Server	Hypervisor	Hypervisor Type	OS Name	OS Type	OS Level	Hostname	Virtual Proce...	Allocated Memory (MB)	Physical CPU Utilization...	Hypervisor CPU Delay (%)	Idle Time (%)	Other Time (%)	Service Class (PI)	OS Proces... Total CPU Using Samples (%)
<input type="radio"/>	GSSF23	VML1	zVM	GSSF23	Red Hat Enterprise Linux Server	5.5	GSSF23	1	821	0.7	0.0	99.3	0.0	SvcClsForModer	0.0
<input type="radio"/>	GSSF24	VML1	zVM	GSSF24	SUSE Linux Enterprise Server	10.4	GSSF24	1	667	20.6	0.5	77.4	0.0	SvcClsForFastes	4.0
<input type="radio"/>	r93f2c1b03v5	C.1.03	x Hyp	r93f2c1b03v5	SUSE Linux Enterprise Server	11.1	r93f2c1b03v5	2	4,096	4.0	0.3			SvcClsForLinux	0.0
<input type="radio"/>	r93f2c1b03v6	C.1.03	x Hyp	r93f2c1b03v6	SUSE Linux Enterprise Server	11.1	r93f2c1b03v6	2	4,096	4.6	0.3			SvcClsForLinux	0.0
<input checked="" type="radio"/>	r93f2c1b06v5	C.1.06	PowerVM	r93f2c1b06v5	AIX	7.1.0.2	r93f2c1b06v5.pokprv.stg	2	4,096	53.3	4.6			Default (0.42)	43.7
<input type="radio"/>	r93f2c1b06v6	C.1.06	PowerVM	r93f2c1b06v6	AIX	7.1.0.2	r93f2c1b06v6.pokprv.stg	2	4,096	1.2	2.6			Default (0.42)	8.8
<input type="radio"/>	r93f2c1b07v1	C.1.07	PowerVM	r93f2c1b07v1	AIX	6.1.6.3	r93f2c1b07v1.pokprv.stg	2	3,072	20.4	6.0			SvcClsForFastes	27.5
<input type="radio"/>	r93f2c1b07v2	C.1.07	PowerVM	r93f2c1b07v2	AIX	6.1.6.3	r93f2c1b07v2.pokprv.stg	2	4,096	45.0	6.0			SvcClsForFastHi	41.3

New support for x86 blade

Monitoring data provided by GPMP running on VS

Virtual Server Charts



Analyze utilization of specific virtual server over time interval

# 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

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:

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

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
- GPMP shipped as firmware component, installed from HMC for zLinux, AIX, xLinux and Windows Operating systems. z/OS ships GPMP as part of OS
- 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
- Provides instrumentation to support SASP Load Balancing (AIX, Linux, Windows)

# Workload Monitoring with GPMP

## 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

Detail hop report with transaction avg response time



--- 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 Webseving 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

Do we have any failure?

How much time it took for transaction execution?

Tasks Zoom Layout



Drill down to specific server to get details about application and resource utilization

Virtual Server Topology derived from ARM data provided by GPMP

Total: 5 Selected: 0

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

# Co-operative management with z/OS WLM

- z/OS provides differentiated service to PPM classified work
- Requirements:
  - GPMP must be running on z/OS
  - 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



Session B - [24 x 80]

File Edit View Communication Actions Window Help

Subsystem-Type Xref Notes Options Help

-----

Modify Rules for the Subsystem Type Row 3 to 10 of 16  
 Command ==> \_\_\_\_\_ Scroll ==> CSR

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

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

Action	Type	Qualifier	Name	Start	Class	Service	Report
_____	1	ESC	SrvClsFo	1	DEFAULTS:	EWLMDEFA	_____
_____	2	ESC	rFastest	9	_____	_____	_____
_____	3	ESC	Highest	17	_____	GPFSTHST	_____
_____	1	ESC	SrvClsFo	1	_____	_____	_____
_____	2	ESC	rFastHig	9	_____	_____	_____
_____	3	ESC	h	17	_____	GPFASHIG	_____
_____	1	ESC	SrvClsFo	1	_____	_____	_____
_____	2	ESC	rModerat	9	_____	_____	_____

MA b 21/040

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

# Monitoring with RMF

Session B - [24 x 80]

File Edit View Communication Actions Window Help

RMF V1R12 Sysplex Summary - ZGPLEX Line 1 of 18  
Command ==> \_ Scroll ==> CSR

WLM Samples: 400    Systems: 5    Date: 09/28/10    Time: 12.38.20    Range: 100    Sec

>>>>>>> [Progress Bar] <<<<<<<<

Service Definition: WLMPOL                                    Installed at: 09/28/10, 12.18.06  
Active Policy: WLMPOL    Activated at: 09/28/10, 12.18.10

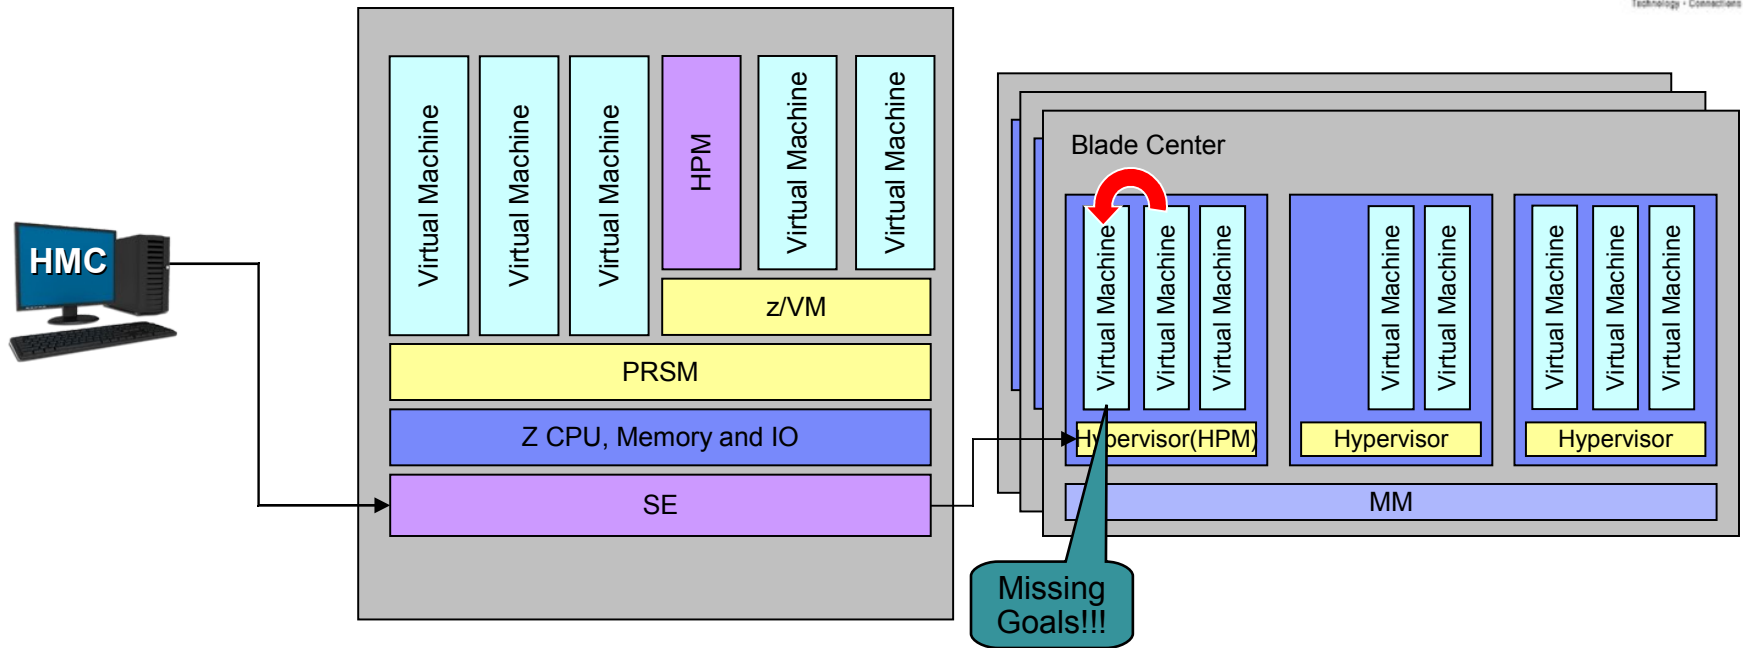
Name	T	I	Goal	Act	Response Goal	Response Actual	Perf	Indx	Trans Ended Rate	--Avg. WAIT Time	Resp. EXECUT Time	Time- ACTUAL Time
DATABASE	W			0.4					2765	0.000	0.009	0.009
DDF	S			0.0					0.010	0.000	0.016	0.016
		1 3		0.0	0.500	80%	100%	0.50	0.010	0.000	0.016	0.016
GPFASHIG	S	2		0.0	0.500	AVG	0.008	AVG	922.3	0.000	0.008	0.008
GPFSTHST	S	1		0.0	0.500	AVG	0.009	AVG	899.3	0.000	0.009	0.009
GPMODMED	S	3		0.9	0.500	AVG	0.009	AVG	943.7	0.000	0.009	0.009
VERYHIGH	S	2	40	0.0				N/A	0.000	0.000	0.000	0.000
STCTASKS	W			78					0.060	0.490	0.336	0.826
STC	S	3	30	78			0.38		0.060	0.490	0.336	0.826
SYSTEM	W			94					0.180	0.061	0.009	0.070

MA b 02/015

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

# Management Functions

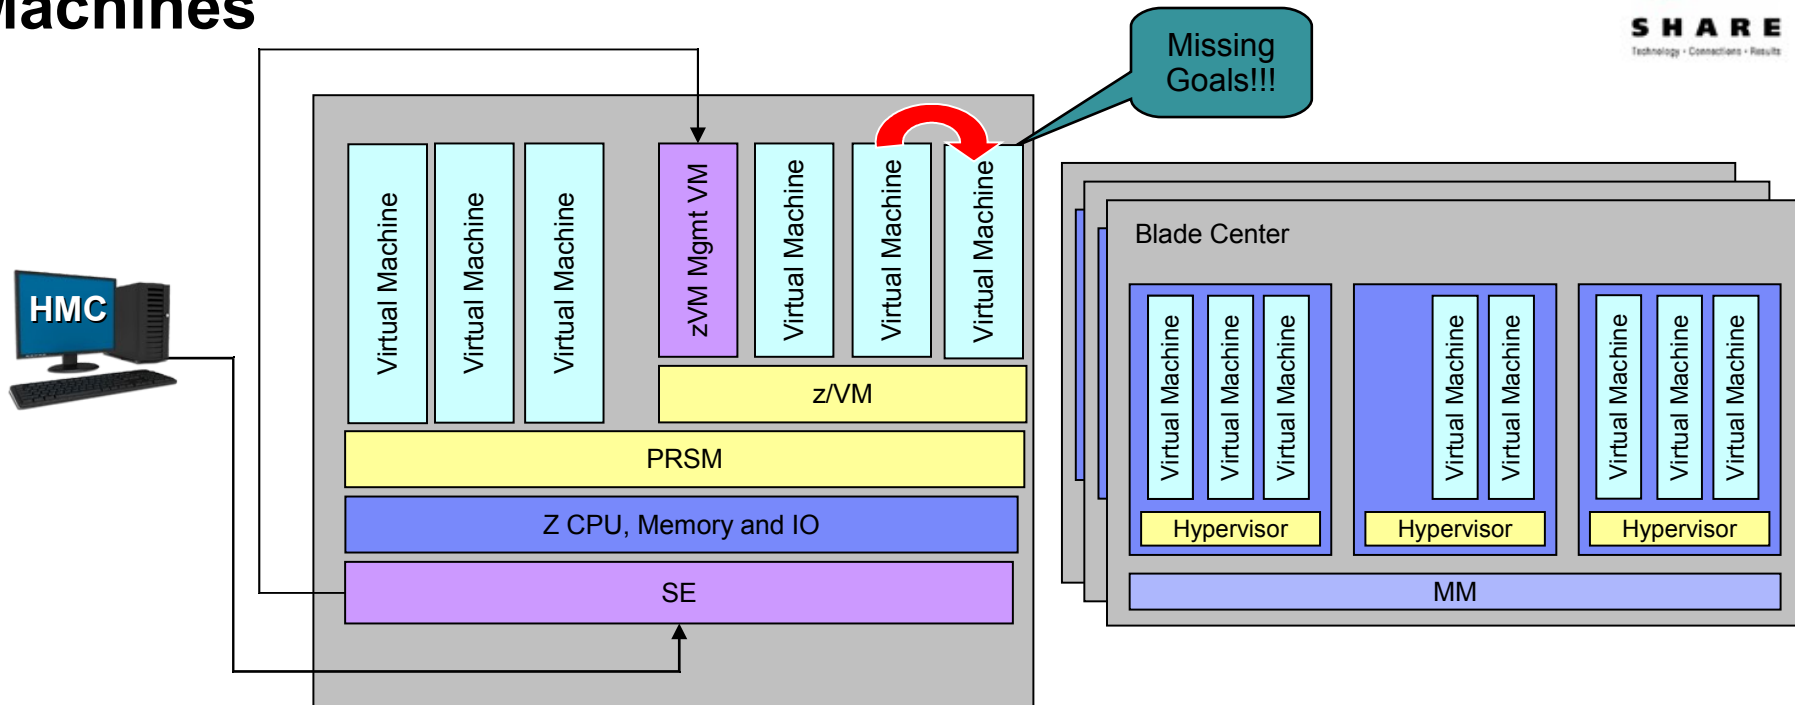
# Managing Resources across Virtual Servers on Power Blade



- 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



HMC1: Workloads Report - Mozilla Firefox  
 http://9.60.31.154:8080/hmc/wcl/T14d0

Virtual Server Resource Adjustments Report - Buyer 1

Report Interval: Last 15 minutes [Modify](#)

Successful Adjustments

Receiver Virtual Servers ^	Receiver Workload ^	Receiver Service Class ^	Receiver Processing Units After (Before) ^	Donor Virtual Servers ^	Donor Workload ^	Donor Processing Units After (Before) ^	Time ^
Buyer 1	Weinheimer Agriculture Parts	Buyers	0.52 (0.50)	Payroll App	Payroll	0.49 (0.50)	Jul 11, 2010 4:13:18 PM
Buyer 1	Weinheimer Agriculture Parts	Buyers	0.52 (0.50)	Vendor 1	Weinheimer Agriculture Parts	0.49 (0.50)	Jul 11, 2010 4:13:18 PM
Total: 2 Filtered: 2							

Failed Adjustments

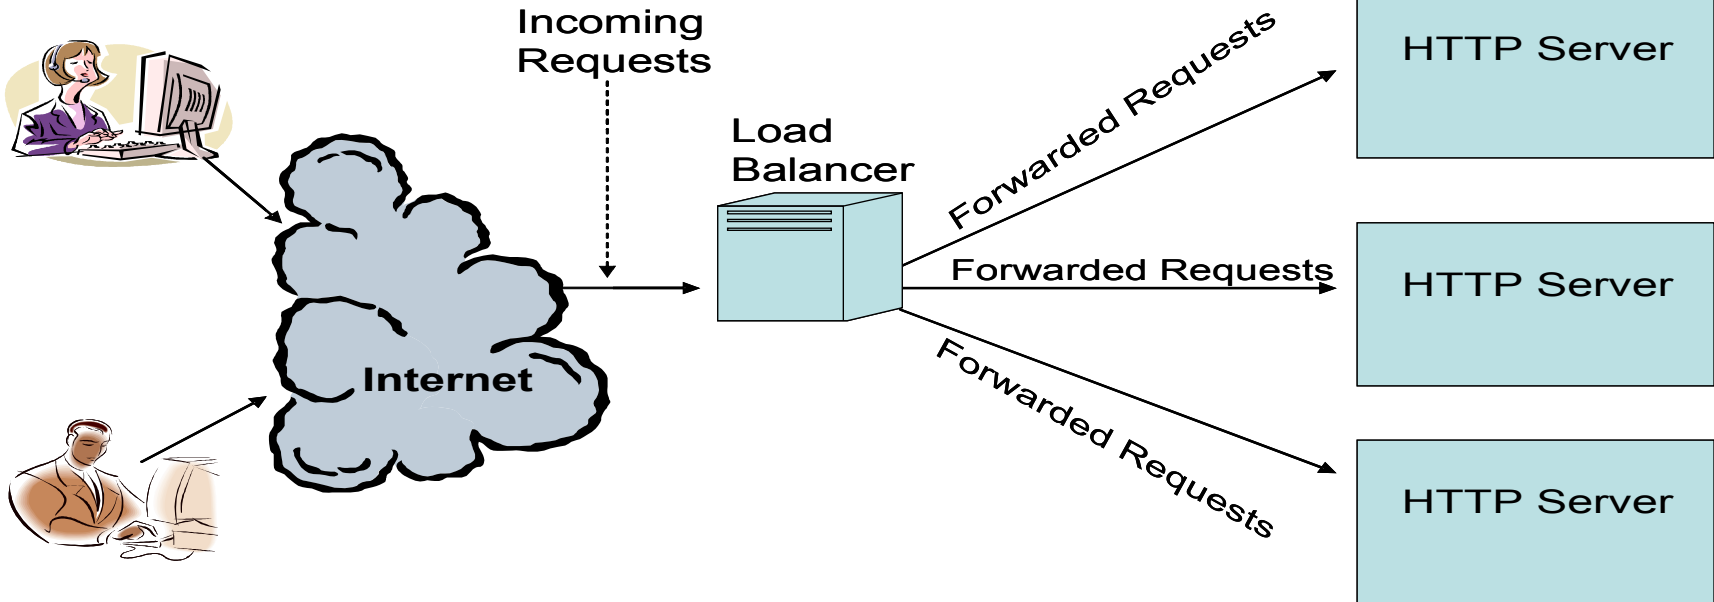
Receiver Virtual Servers ^	Receiver Workload ^	Receiver Service Class ^	Failure Reason ^	Time ^
Total: 0 Filtered: 0				

Done

Examine resource adjustment action performed by PPM to help work meet its performance goal

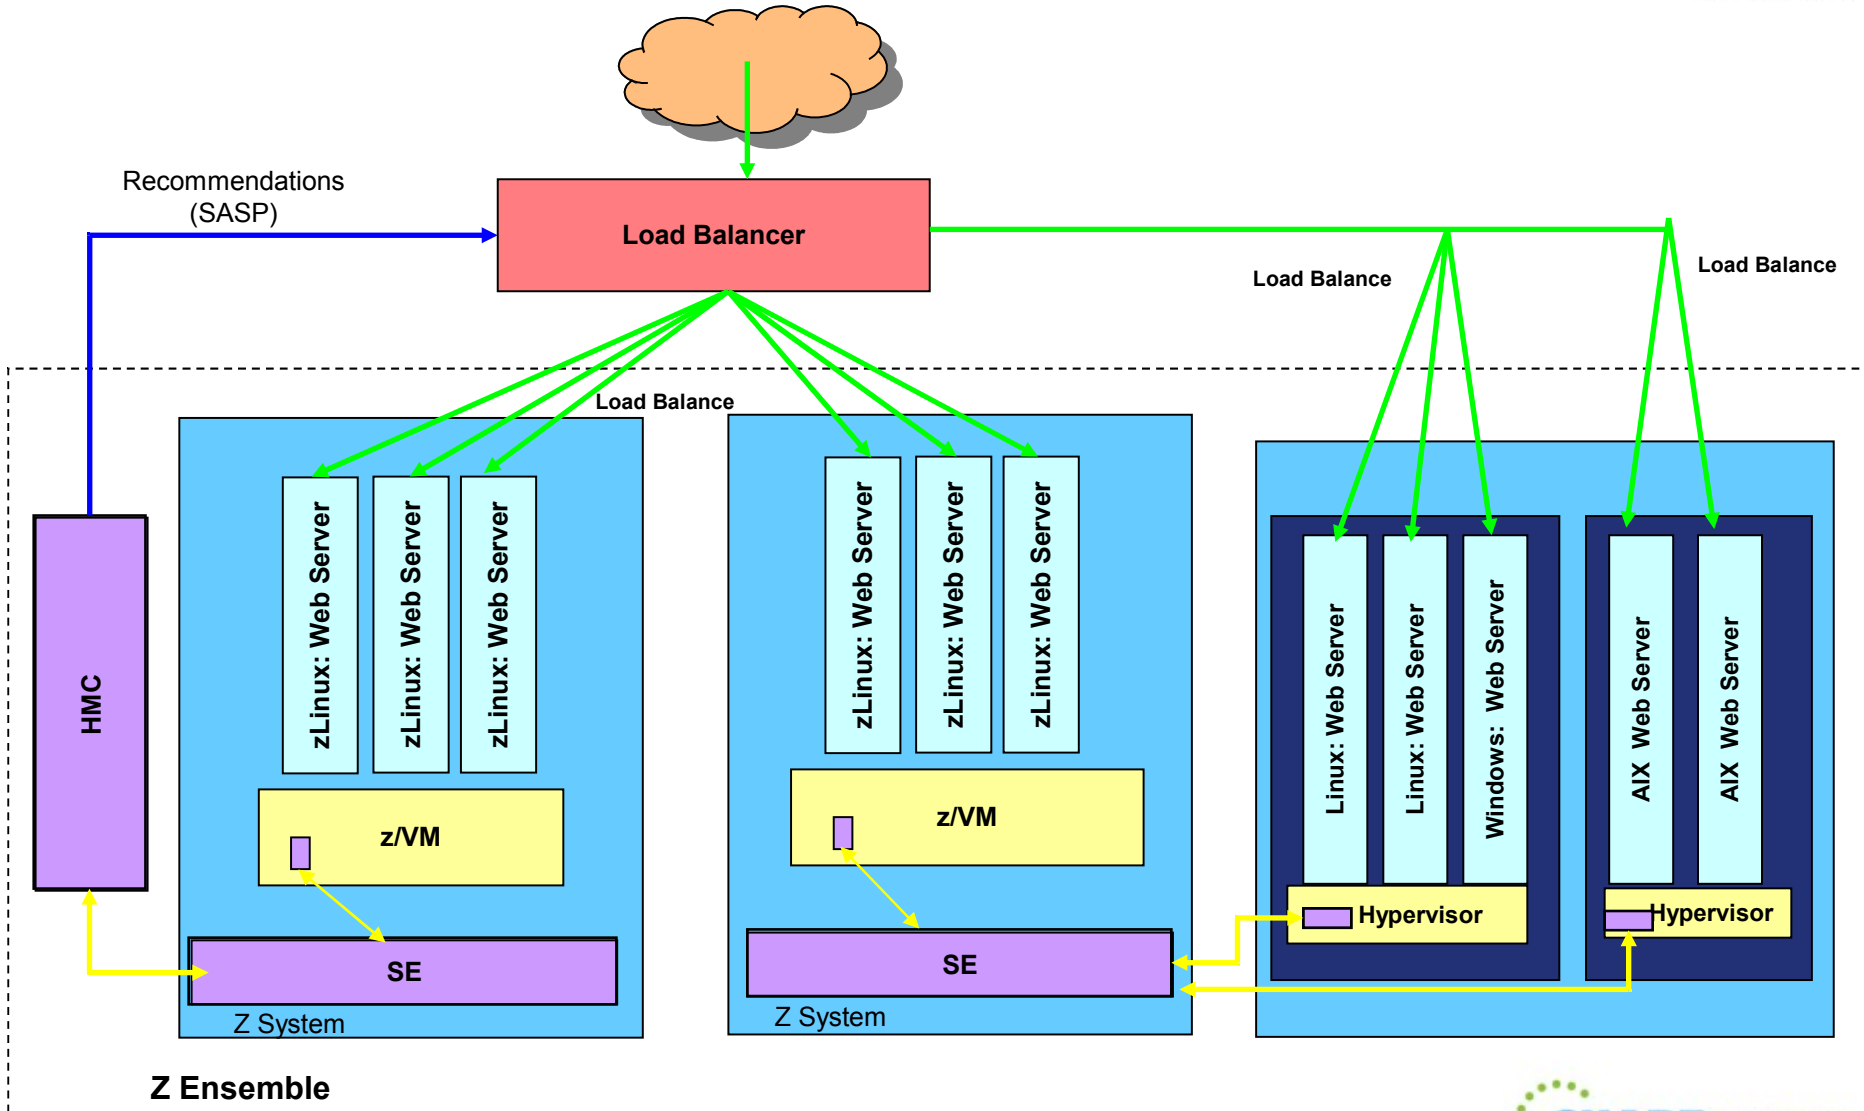
Explains why resource adjustment action was not performed

# What is Load Balancing?



- Distributes incoming workload among a group of machines with similar functionality.
  - factoring in the availability of server resources,
  - avoiding over-utilized servers, where possible,
  - factoring in the *business importance* of the work (potential future consideration)
  - estimating the likelihood of meeting objectives, (potential future consideration)

# Ensemble Workload Balancing



# Motivation for PPM to provide recommendation

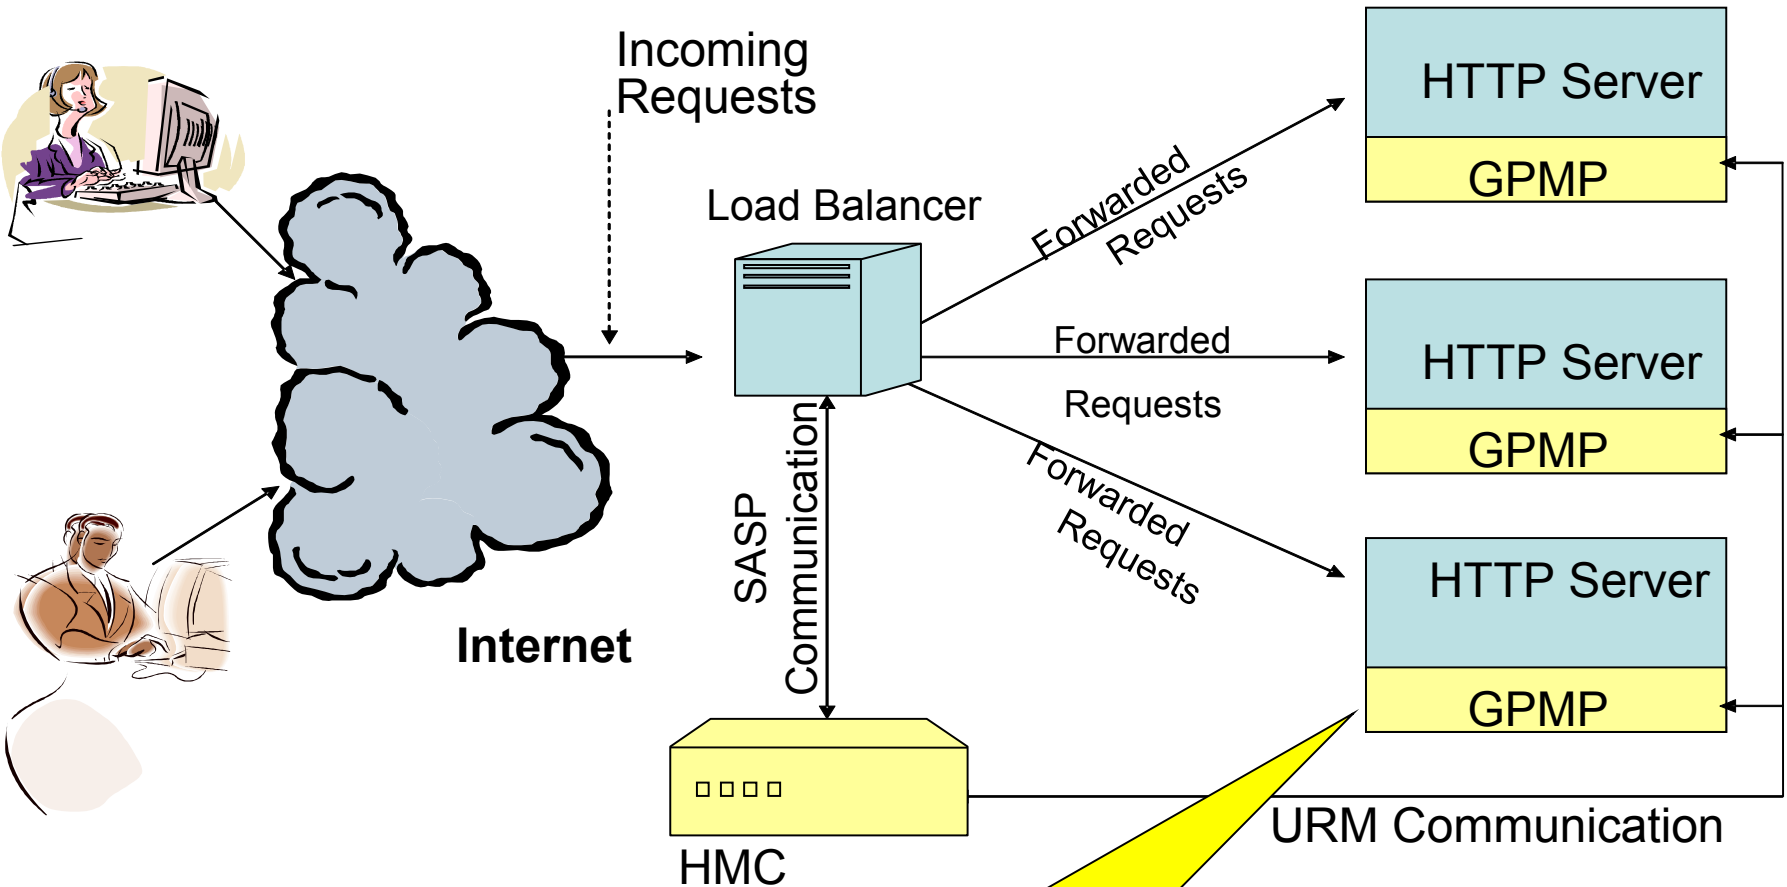


- Objective: Influence workload balancing decisions across a System z Ensemble
  - Use SASP protocol to make recommendations for workload balancers (e.g. IP switches / routers that load balance)
  - HMC hosts SASP function
  - Scope of recommendations is non-z/OS virtual servers within the ensemble
    - z/OS Load Balancing Advisor (LBA) provides SASP recommendations for z/OS
    - Same SASP client code can interact with both LBA and HMC SASP implementations to provide complete coverage of z environment
      - *Client code needs to be configured to different SASP server for z/OS (LBA) and non z/OS (HMC)*
  - HMC recommendations based on the platform performance manager's understanding of the current performance of the members of a load balancing group
    - Recommendation based on overall utilization and delays experienced by virtual servers
    - If IP address and port are used to register members of a load balancing group, port is used to determine application availability on each member of load balancing group. Weight of 0 given to members where port is not open

# The New Server/Application State Protocol (SASP)

- Protocol Characteristics
  - TCP/IP Binary protocol
  - Open protocol documented in RFC4678:
    - <http://www.faqs.org/rfcs/rfc4678.html>
- Load Balancers can connect to URM for the following purposes:
  - Register Groups and Members
  - Deregister Groups and Members
  - Set quiesced/active state of Members
  - Set Load Balancer configuration parameters
    - Have weights pushed to LB
    - Allow members to register/deregister themselves
    - Don't send weights if there was no change
  - Get current weights

# How will URM perform load balancing?



- GPMP samples system statistics
- GPMP knows whether application (HTTP Server) is currently running or not
- URM utilizes these metrics from GPMP to provide recommendations to load balancer

# Enable Load Balancing support

**Ensemble Details - PPM1** i

Instance Information   Status   **Performance Management**   Energy Management   Network Information   Alternate

Processor performance management for z/VM:

Processor performance management for POWER hypervisors:

*Load Balancing*

Enable load balancing

Port:

IP Addresses:

IP Address:

*Note: A yellow callout bubble points to the IP address 9.12.43.140 with the text "IP address of switch with SASP support".*



# 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
- Load Balancing recommendation to SASP enabled routers.

# Back up

- Back up

# 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.

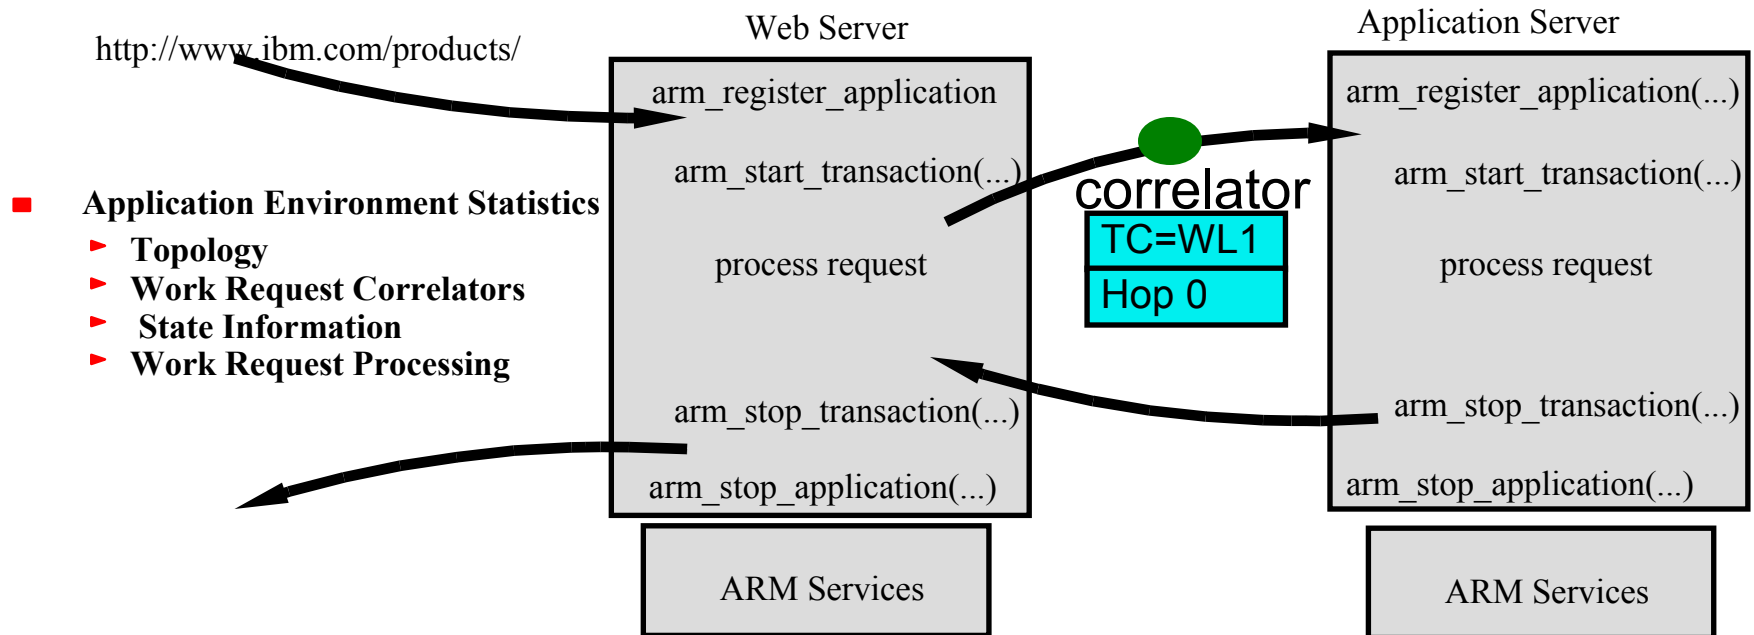


# Basic ARM calls

- 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