
z/OS Workload Manager (WLM) Update for z/OS V1.13 and V1.12

November 2012

Horst Sinram, *z/OS Workload Management*

IBM Germany Research & Development

[Email: sinram@de.ibm.com](mailto:sinram@de.ibm.com)





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

AIX*	DS8000*	Language Environment*	SystemPac*	z10
BladeCenter*	FICON*	Parallel Sysplex*	System Storage	z10 BC
DataPower*	HiperSockets	POWER7*	System z	z10 EC
DB2*	Hyperwap	PrintWay	System z9	z/OS*
DFSMS	IBM*	ProductPac*	System z10	zEnterprise
DFSMSdss	IBM eServer	RACF*	System z10 Business Class	zSeries*
DFSMSshsm	IBM logo*	REXX	WebSphere*	
DFSMSrmm	ibm.com	RMF	z9*	
DFSORT	Infiniband*	ServerPac*		
DS6000*	InfoPrint			

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

InfiniBand is a registered trademark of the InfiniBand Trade Association (IBTA).

Intel is a trademark of the Intel Corporation in the United States and other countries.

Linux is a trademark of Linux Torvalds in the United States, other countries, or both.

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

All other products may be trademarks or registered trademarks of their respective companies.

The Open Group is a registered trademark of The Open Group in the US and other countries.

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.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.



- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals

- HiperDispatch and other relevant APARs

- WLM Support for I/O Priority Manager in DS8K Series

- WLM Support for IBM zEnterprise 196

- Temporary Capacity Reporting via SYSEVENT REQLPDAT

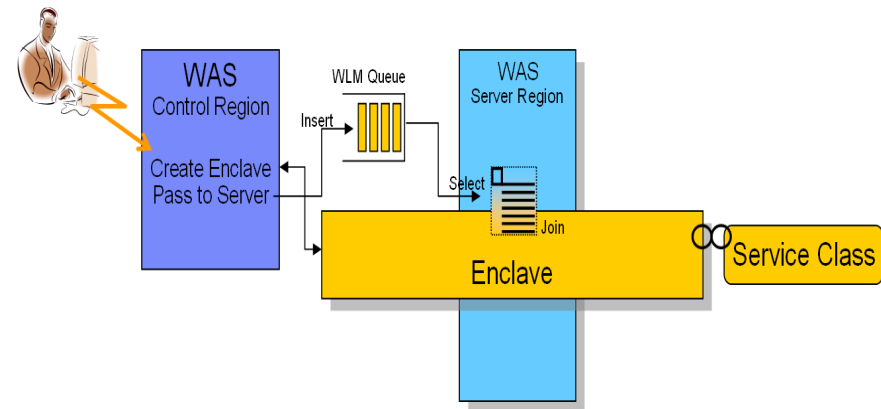
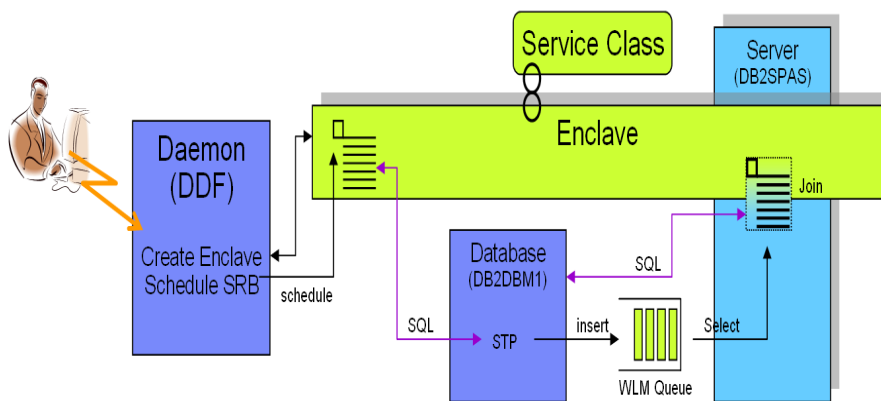
- z/OSMF Workload Management

- WLM support for Unified Resource Manager

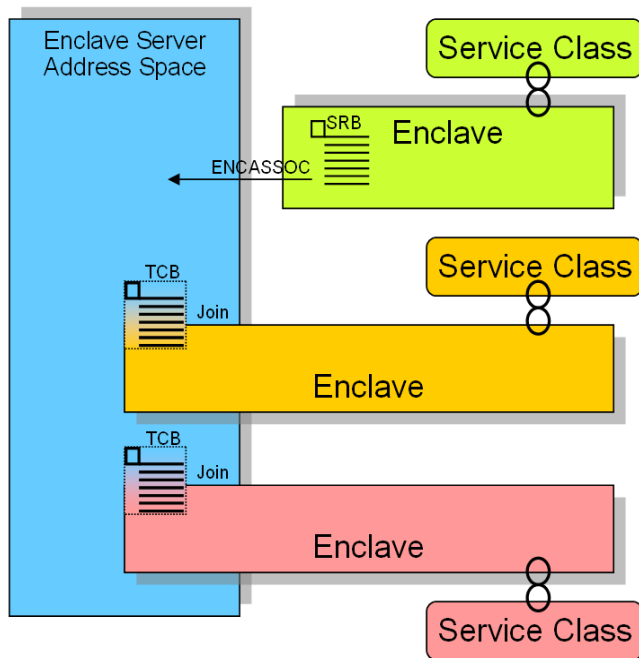
- Capacity Provisioning Update Summary

- WLM Tools Overview

- An **enclave** is a transaction that can span multiple dispatchable units (SRBs and tasks) in one or several address spaces and is reported on and managed as one unit
- The enclave is managed separately from the address spaces it runs in
 - CPU and I/O resources associated with processing the transaction represented by the enclave are managed by the transaction's performance goal
 - Storage (MPL level, paging) of the address space is managed to meet the goals of the enclaves it serves (if enclave server address space) or to the performance goal of the address space (if no server address space)
- Typical exploitation through DB2 and WebSphere



WLM Enclave Server Management: A short review




- An address space becomes an enclave server when
 - An enclave SRB issues SYSEVENT ENCASSOC
 - A TCB of the address space joins an enclave, and does not specify ENCLAVESERVER=NO (which is typically not the case)

- **Original** implementation assumed:
 - All work being executed within the address space is related to enclaves
 - Therefore no significant amount of work (TCBs) executing in such address spaces unrelated to enclaves

 - Enclave Server Management (old)
 - CPU and I/O DP is derived from service class of most important enclave
 - Hence no CPU and I/O management exists for these server address spaces
 - Storage management is done to meet the served enclave's goals

WLM Enclave Server Management Rationale and Changes with z/OS 1.12

- However the original assumptions did not always hold true - there **are** examples of significant work unrelated to an enclave:
 - Garbage collection for a JVM (WAS)
 - Common routines which provide service for the enclave TCBs
- Problems exist when no enclaves are running in server address spaces and the address space is swapped out
- Solution by new capability for managing non-enclave-related work.
- Controlled by new IEAOPT Parameter
 - **ManageNonEnclaveWork = {No|Yes}**
 - Yes: Work in the address space not associated to an enclave is managed towards the goals of the external service class to which the address space has been classified to
 - No: Non enclave work is managed based on the most important enclave
- With ManageNonEnclaveWork = Yes the importance and goal of the service class for the address space is more important than it used to be. Housekeeping and address space(s) restart performed under the address space classification
- Recommendation:
 - Verify goal/importance settings for server address spaces and when acceptable specify ManageNonEnclaveWork = Yes



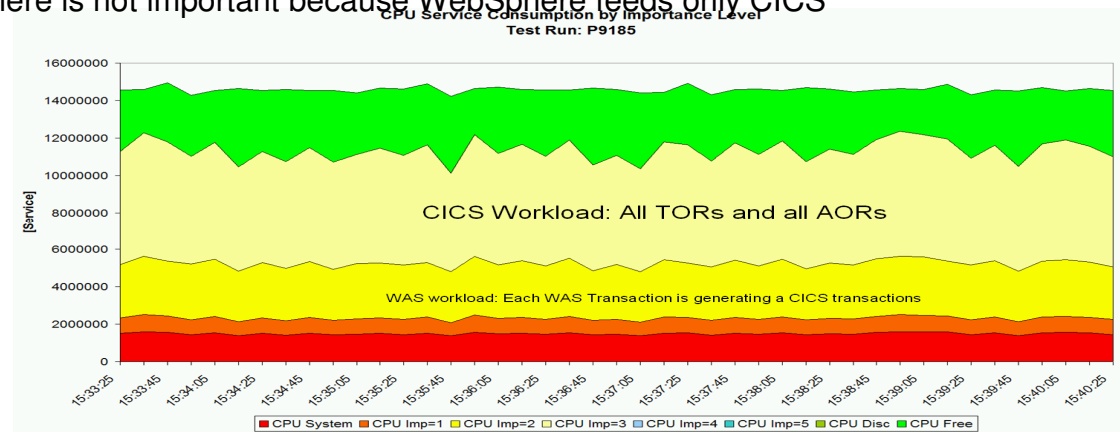
OA35428: New management option for CICS environments

Problem scenario and analysis

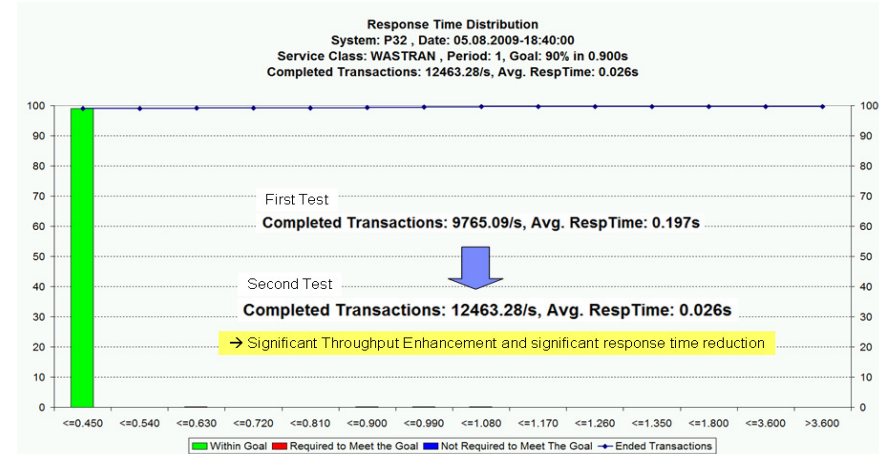
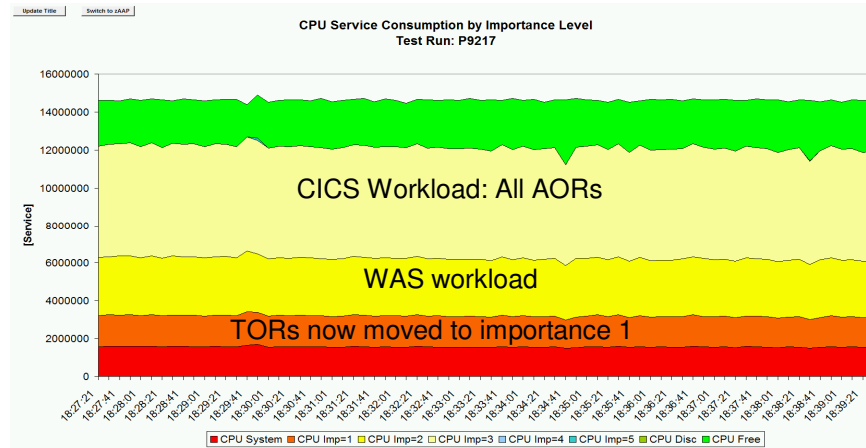
- Problem can show up when a workload predominantly consists of a CICS workload with little displaceable other work
- Sample environment:
 - System: 2097-764, 1 LPAR only → 8 nodes with 4 High processors each
 - Workload: WebSphere → CICS → DB2
 - WebSphere receives work, sends it to CICS TORs which send it to AORs which execute DB2 calls
 - Classification: WebSphere Imp=2 and all CICS Imp=3, managed towards response time goals
- Problem:
 - Low system throughput; relatively high response times. System utilization could not exceed 80%**
- Note: In this scenario, the presence of WebSphere is not important because WebSphere feeds only CICS

Problem Analysis

- TORs and AORs run at the same dispatch priority
- AORs heavily consumes CPU. TORs need to wait too long to receive work and return results to the caller
- HiperDispatch can amplify the situation because it runs the work at higher utilization on nodes with typically 4 processors



OA35428 Two possible circumventions



- Circumvention:

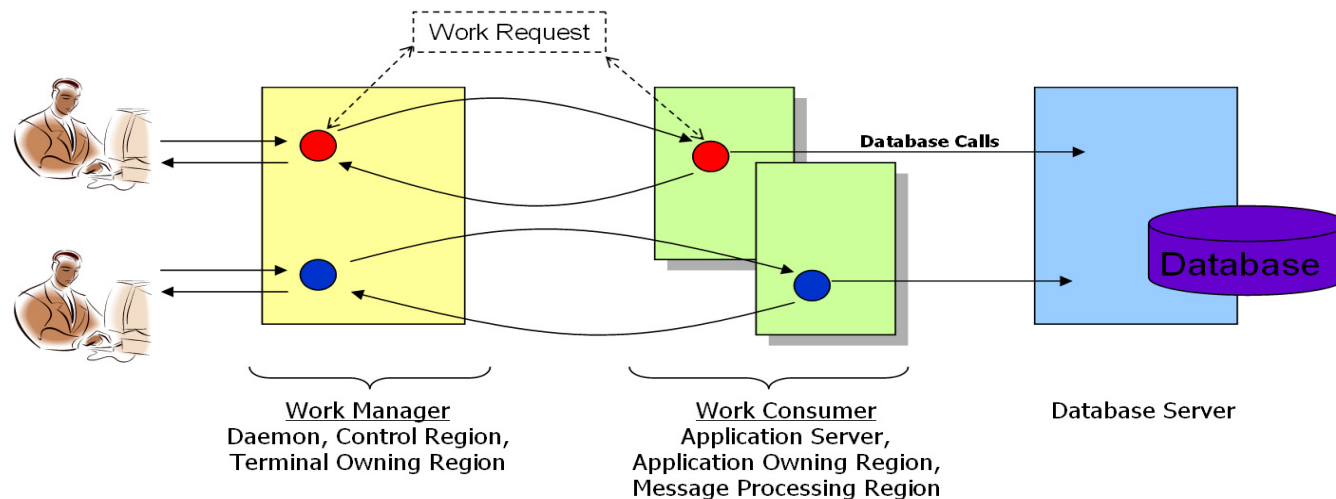
- Move TORs to a service class with higher importance than AORs

- Option 1: Exempt **all** regions from being managed by response time goals and classify TORs to a service class with higher importance than AORs.

Disadvantage: No response time data present

- Option 2: Exempt only AORs and move them to a service class with lower importance than the CICS service classes with response time goals.

Disadvantage: Response time data cover only a small portion of the execution path because AORs consume much more than TORs.



- Adjust WLM CICS management to follow “Work Manager/Consumer” model
 - A TOR is a region which consumes typically little CPU and just functions as a work receiver and result sender.
 - This is the work manager - needs quick access to CPU
 - An AOR is a server region which typically is much more resource intensive.
 - This is the work consumer. An AOR doesn’t require the same instantaneous access to CPU than TORs
- WLM already has experience with that model:
 - WebSphere Application Server and DB2/DDF work
 - The control regions are managed towards execution velocity goals
 - The work is managed towards response time goals (via enclaves and the servers processing the enclaves are tight to their goals)

OA35428: Enhancement of WLM Management

Solution to solve the Work Manager/Consumer Model

- New “Manage Regions by Goals Of” option in WLM service definition: “BOTH”
 - Use option “BOTH” for TORs
 - Define a STC service class for TORs which has a higher importance than the CICS service class with response time goals for the CICS work and AORs
 - Stay with “Manage Regions by Goals Of Transaction” for AORs.

- Result:
 - WLM will manage the TORs towards the goals of the STC service class
 - And WLM will ensure bookkeeping of transaction completions to the correct CICS response time service class
 - The CICS transactions are managed towards CICS response time goals and the AORs are also managed towards these goals like today

```

Subsystem-Type  Xref  Notes  Options  Help
-----
Command ==>  Modify Rules for the Subsystem Type  Row 1 to 3 of 3
                Scroll ==> PAGE

Subsystem Type . : JES          Fold qualifier names?  Y (Y or N)
Description . . : Batch Work

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

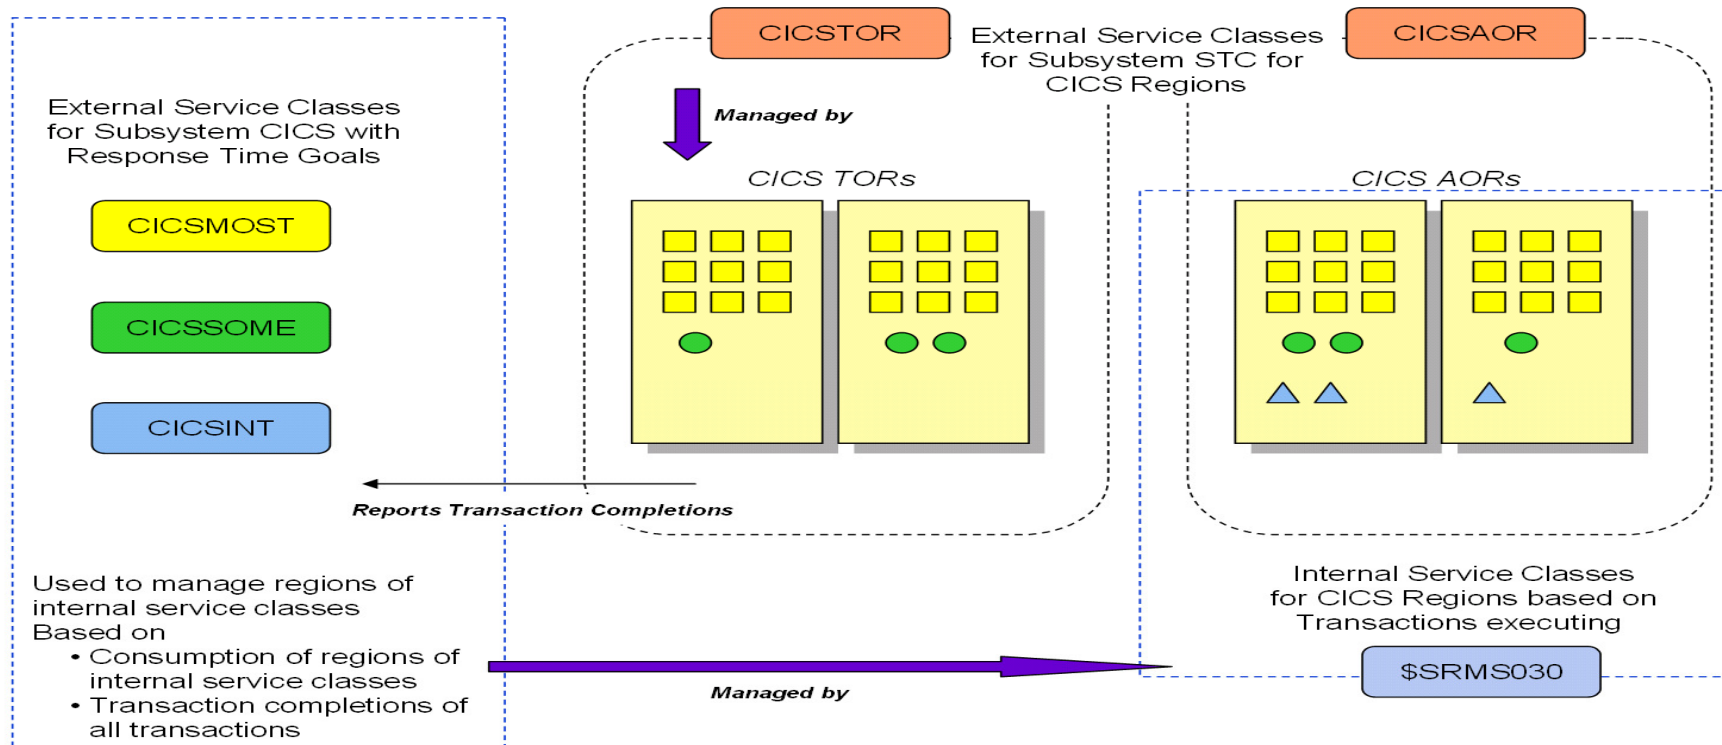
Action      -----Qualifier-----      Storage      Manage Region
Type      Name      Start      Critical      Using Goals Of

_____  1  TN      CICSTOR*  _____  NO      BOTH
_____  1  TN      CICSAOR*  _____  NO      TRANSACTION
_____  1  TN      CICS*     _____  NO      TRANSACTION
***** BOTTOM OF DATA *****

```

OA35428: New WLM Management Option

Structure of Service Classes



- TORs are now managed towards the goal of the service class **CICSTOR**
 - They still report their transaction completions for management
- AORs are still managed towards the goals of the CICS service classes and the consumption of the internal service class for the region
- Recommendation: **CICSTOR** should be defined at a higher importance than the CICS service classes

Background: WLM CICS Management Summary

Options for managing CICS work

1. CICS managed by Response Time Goals

- All Regions defined as managed towards TRANSACTION goals
 - Existing Method
 - Works well for most environments
 - Older CICS environments which do not have Multi Region Option
 - All environments which are not exclusively CICS workload or don't have any problem

2. CICS managed by Region Goals

- If response time goals have not been defined all CICS regions are managed towards REGION goals (exempted from transaction management)
 - Existing Method
 - Works also well for most environments
 - But: Execution velocity goals are more sensitive to hardware and software changes
 - Usually no transaction reporting available
 - o This can be enabled for report classes but requires additional definitions

3. CICS managed by Region and Response Time Goals

- CICS TORs defined as managed towards BOTH goals
- CICS AORs defined as managed towards TRANSACTION goals
 - [New Method introduced with OA35428](#)
 - Works well for most environments, too. Especially suited for systems where CICS is the predominant workload.

Response Time Distribution for Velocity Goals (z/OS V1.13)



- Currently WLM reporting does not provide a response time distribution (ended transactions) for workloads with velocity goals
- Sometimes it is desirable to have a response time distribution for all transactional workloads, even if they have a velocity goal
 - More data to analyze workload behavior and to detect problems
 - Better support for migration of goal definitions to response time goals
- With z/OS V1.13
 - the IWMRCOLL answer area IWMWRCAA provides also a response time distribution for service class periods with an execution velocity goal,
 - **the RMF Postprocessor Workload Activity report (WLMGL) displays the new response time distributions**
 - Response Time distributions also added to SMF 99 subtype 2 data

Response Time Distribution for Velocity Goals: The Mid-Point Change Algorithm




- Rationale

- Velocity goals do not have a „reference“ response time
- The mid-point (MP) should be set to values which matches the current workload conditions
 - Mid point may change drastically from time to time, but WLM expects them to be consistent for a time long enough, so that it can compute sensible mid-point (MP) values

- Algorithm

- The model behind the algorithm is a gaussian RT distribution, with $d = M/3$ (M is the mean, d is the standard deviation). Ideal would be $MP = M$
- Mid point will be (re)computed after 10-70 min
 - Depending on significance of deviation



This is the ideal case:

The mid-point set by WLM is equals the average response time (M) of the transactions, so we get a recognizable gaussian distribution



Response Time Distribution for Velocity Goals: RMF WLMGL Enhancement

REPORT BY: POLICY=POLICY01 WORKLOAD=STC SERVICE CLASS=STCDEF RESOURCE GROUP=*NONE PERIOD=1 IMPORTANCE=5
CRITICAL =NONE

-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD I/O--	---SERVICE---	SERVICE TIME	---APPL %---	--PROMOTED--	----STORAGE----
AVG	28.04	ACTUAL	16.629	SSCHRT 89.0	IOC 524944	CPU 1.453	CP 0.22	BLK 0.000
MPL	28.04	EXECUTION	15.724	RESP 0.2	CPU 649332	SRB 0.277	AAPCP 0.00	ENQ 0.000
ENDED	2	QUEUED	904	CONN 0.1	MSO 14840	RCT 0.010	IIPCP 0.00	CRM 0.000
END/S	0.00	R/S AFFIN	0	DISC 0.0	SRB 123890	IIT 0.197	LCK 0.000	
#SWAPS	100	INELIGIBLE	0	Q+PEND 0.1	TOT 1313K	HST 0.000	AAP 0.00	
EXCTD	0	CONVERSION	0	IOSQ 0.0	/SEC 1459	AAP 0.000	IIP 0.00	
AVG ENC	0.00	STD DEV	0			IIP 0.000		
REM ENC	0.00				ABSRPTN 52			
MS ENC	0.00				TRX SERV 52			
								-PAGE-IN RATES-
								SINGLE 0.0
								BLOCK 0.0
								SHARED 0.0
								HSP 0.0

GOAL: EXECUTION VELOCITY 20.0% VELOCITY MIGRATION: I/O MGMT 88.2% INIT MGMT 88.2%

...

-----RESPONSE TIME DISTRIBUTIONS-----									
SYSTEM: SYSD -----INTERVAL: 14.59.998 -----MRT CHANGES: 0 ---					SYSTEM: SYSE -----INTERVAL: 01.22.123 -----MRT CHANGES: 1 ---				
----TIME----	-NUMBER OF	TRANSACTIONS-	-----PERCENT-----		----TIME----	-NUMBER OF	TRANSACTIONS-	-----PERCENT-----	
HH.MM.SS.TTT	CUM TOTAL	IN BUCKET	CUM TOTAL	IN BUCKET	HH.MM.SS.TTT	CUM TOTAL	IN BUCKET	CUM TOTAL	IN BUCKET
< 00.00.00.200	581	581	94.2	94.2	< 00.00.00.300	581	581	94.2	94.2
<= 00.00.00.240	584	3	94.7	0.5	<= 00.00.00.360	584	3	94.7	0.5
<= 00.00.00.280	586	2	95.0	0.3	<= 00.00.00.420	586	2	95.0	0.3
<= 00.00.00.320	586	0	95.0	0.0	<= 00.00.00.480	586	0	95.0	0.0
<= 00.00.00.360	588	2	95.3	0.3	<= 00.00.00.640	588	2	95.3	0.3
<= 00.00.00.400	591	3	95.8	0.5	<= 00.00.00.600	591	3	95.8	0.5
<= 00.00.00.440	592	1	95.9	0.2	<= 00.00.00.660	592	1	95.9	0.2
<= 00.00.00.480	592	0	95.9	0.0	<= 00.00.00.720	592	0	95.9	0.0
<= 00.00.00.520	593	1	96.1	0.2	<= 00.00.00.780	593	1	96.1	0.2
<= 00.00.00.560	596	3	96.6	0.5	<= 00.00.00.840	596	3	96.6	0.5
<= 00.00.00.600	596	0	96.6	0.0	<= 00.00.00.900	596	0	96.6	0.0
<= 00.00.00.800	599	3	97.1	0.5	<= 00.00.01.200	599	3	97.1	0.5
<= 00.00.01.600	604	5	97.9	0.8	<= 00.00.02.400	604	5	97.9	0.8
> 00.00.01.600	617	13	100	2.1	> 00.00.02.400	617	13	100	2.1



Response Time Distribution for Velocity Goals: RMF WLMGL Enhancement

REPORT BY: POLICY=BASEPOL		WORKLOAD=STC_WLD		SERVICE CLASS=STCLOW CRITICAL =NONE		RESOURCE GROUP=*NONE		PERIOD=1 IMPORTANCE=3							
-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD I/O--	---SERVICE---	SERVICE TIME	---APPL %---	--PROMOTED--	----STORAGE----							
AVG	5.24	ACTUAL	7.753	SSCHRT	1.5	IOC	6570	CPU	7.909	CP	22.19	BLK	0.000	AVG	612.39
MPL	5.24	EXECUTION	7.198	RESP	0.3	CPU	4503K	SRB	125.182	AAPCP	0.00	ENQ	0.000	TOTAL	3208.80
ENDED	436	QUEUED	552	CONN	0.1	MSO	0	RCT	0.009	IIPCP	0.00	CRM	0.000	SHARED	0.00
END/S	0.73	R/S AFFIN	2	DISC	0.0	SRB	71278K	IIT	0.026			LCK	0.000		
#SWAPS	0	INELIGIBLE	0	Q+PEND	0.2	TOT	75788K	HST	0.000	AAP	0.00	SUP	0.000	-PAGE-IN RATES-	
EXCTD	0	CONVERSION	14	IOSQ	0.0	/SEC	126314	AAP	0.000	IIP	0.00			SINGLE	0.0
AVG ENC	0.00	STD DEV	669					IIP	0.000					BLOCK	0.0
REM ENC	0.00					ABSRPTN	24K							SHARED	0.0
MS ENC	0.00					TRX SERV	24K							HSP	0.0

GOAL: EXECUTION VELOCITY 90.0% VELOCITY MIGRATION: I/O MGMT 95.5% INIT MGMT 42.8%

SYSTEM	RESPONSE TIME EX	PERF	AVG	--EXEC USING%--				EXEC DELAYS %				-USING%-		--- DELAY % ---			%		
	VEL%	INDX	ADRSP	CPU	AAP	IIP	I/O	TOT	CPU				CRY	CNT	UNK	IDL	CRY	CNT	QUI
*ALL	--N/A--	95.5	0.9	5.3	4.4	0.0	0.0	0.0	0.2	0.2			0.0	0.0	0.4	95	0.0	0.0	0.0
TRX1		97.2	0.9	2.6	4.3	0.0	0.0	0.0	0.1	0.1			0.0	0.0	0.4	95	0.0	0.0	0.0
TRX2		94.0	1.0	2.7	4.4	0.0	0.0	0.0	0.3	0.3			0.0	0.0	0.4	95	0.0	0.0	0.0

-----RESPONSE TIME DISTRIBUTIONS-----									
SYSTEM: TRX1 -----INTERVAL: 14.59.995 -----MRT CHANGES: 0 ---					SYSTEM: TRX2 -----INTERVAL: 14.59.995 -----MRT CHANGES: 0 ---				
----	----	----	----	----	----	----	----	----	----
TIME	NUMBER OF	TRANSACTIONS-	PERCENT		TIME	NUMBER OF	TRANSACTIONS-	PERCENT	
HH.MM.SS.TTT	CUM TOTAL	IN BUCKET	CUM TOTAL	IN BUCKET	HH.MM.SS.TTT	CUM TOTAL	IN BUCKET	CUM TOTAL	IN BUCKET
< 00.00.05.317	6	6	2.7	2.7	< 00.00.05.317	7	7	3.2	3.2
<= 00.00.06.381	45	39	20.5	17.8	<= 00.00.06.381	38	31	17.5	14.3
<= 00.00.07.444	99	54	45.2	24.7	<= 00.00.07.444	88	50	40.6	23.0
<= 00.00.08.508	152	53	69.4	24.2	<= 00.00.08.508	139	51	64.1	23.5
<= 00.00.09.571	194	42	88.6	19.2	<= 00.00.09.571	193	54	88.9	24.9
<= 00.00.10.635	219	25	100	11.4	<= 00.00.10.635	217	24	100	11.1
<= 00.00.11.698	219	0	100	0.0	<= 00.00.11.698	217	0	100	0.0
<= 00.00.12.762	219	0	100	0.0	<= 00.00.12.762	217	0	100	0.0
<= 00.00.13.825	219	0	100	0.0	<= 00.00.13.825	217	0	100	0.0
<= 00.00.14.889	219	0	100	0.0	<= 00.00.14.889	217	0	100	0.0
<= 00.00.15.952	219	0	100	0.0	<= 00.00.15.952	217	0	100	0.0
<= 00.00.21.270	219	0	100	0.0	<= 00.00.21.270	217	0	100	0.0
<= 00.00.42.540	219	0	100	0.0	<= 00.00.42.540	217	0	100	0.0
> 00.00.42.540	219	0	100	0.0	> 00.00.42.540	217	0	100	0.0

Response Time Distribution for Velocity Goals

IWMRCOLL enhancements for Service and Report Class Periods

Section	Field	Response time goals	Execution velocity goals
RCAEIHDR (RCAE period header)	RCAEIMID (mid-point in milliseconds)	Same as goal value (milliseconds)	0 after policy activation/refresh/IPL New value computed when WLM detects that current workload distribution deviates too much from RCAEIMID for a too long time
	RCAEIRCT (running count)	N/A (value always 0)	Total number of RCAEIMID changes since last policy activation
	RCAEITST (timestamp of last change)	Policy activation time	Time of last RCAEIMID change or time of last policy activation
RCAEDIST	RCAEDENT	No change Distribution centered around goal value	Centered around RCAEIMID Reset after each RCAEIMID change

- **RCAEIRCT**
 - Is reset to 0 after each policy activation/refresh/IPL
 - Is incremented each time the report class period becomes heterogeneous (when RCAEPLSC and RCAEPMCI are updated)
 - Is incremented each time a transaction is reported with a new mid-point/timestamp
- The 14 buckets of the report class period's response time distribution are reset to 0 when RCAEIRCT is updated
- RCAEIMID is copied from the current service class period's RCAEIMID each time the report class period's RCAEIRCT is incremented
- RCAEITST is copied from the current service class period's RCAEITST each time the report class period's RCAEIRCT is incremented

Transaction Management Enhancements: Availability

Function	z/OS V1.13	z/OS V1.12	z/OS V1.11	z/OS V1.10
OA38320: Reduce Sampling Overhead (WLM address space utilization)	OA38320	OA38320	OA38320	
Enclave Server Management (Non Shell Server Management)	+	+		
CICS Region/RT Management	+	OA35428	OA35428	OA35428
RT Distribution for Execution Velocity Goals	+			

- OA23320: DB2 and CICS subsystems can create many PB/PBDE control blocks for monitoring environments. Code reduces WLM CPU consumption on z196 and z114 when many such control block exists.
- Enclave Server Management
 - Is enabled via OPT parameter ENCLAVESERVER=YES
- CICS Region/RT Management
 - OA35248 supersedes OA34801
 - OA34801 was a temporary solution introduced for a customer to maintain response time reporting while temporarily moving from response time to region management
 - OA34801 introduced a new OPT parameter REPORTCOMPLETIONS={YES|NO}
 - With OA35428 and with z/OS 1.13 WLM will still accept the new OPT parameter but the reporting functionality introduced by OA34801 is no longer supported.

Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals

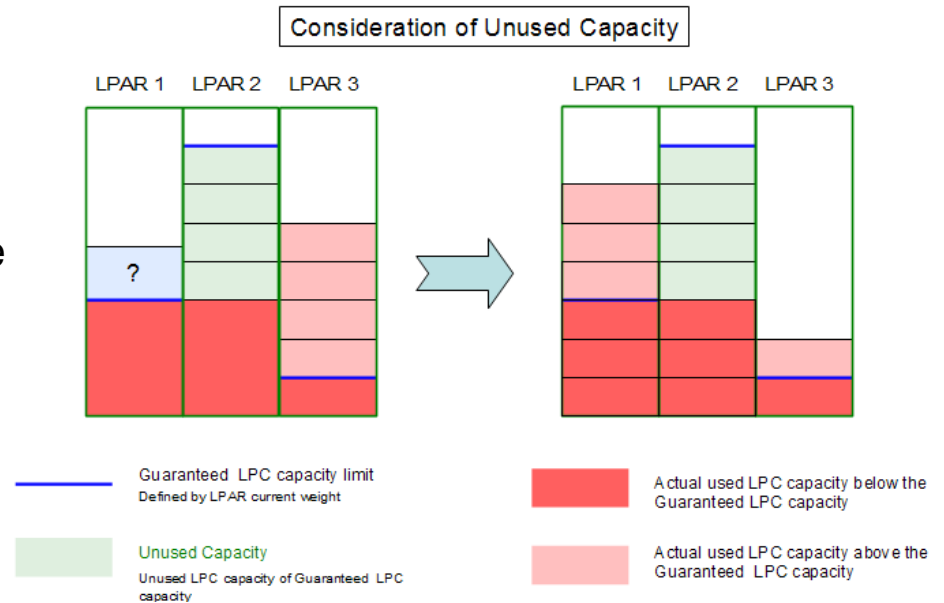


- HiperDispatch and other relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
- WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview

OA37736: HiperDispatch Enhancements for “Unused Capacity”

- Problem addressed:
 - If a large LPAR consumes below its weight-entitlement it is possible that a low-weight LPAR unparks many Vertical Low (VL) processors
 - Therefore small LPAR could “dominate” larger LPARs because those could not unpark their VLs

- Solution:
 - HiperDispatch considers now also the “unused capacity share” for a partition to unpark VLs.
 - This share is calculated by dividing the unused capacity (guaranteed but not used) of all partitions in the CEC by the share of the partitions which can use more capacity.




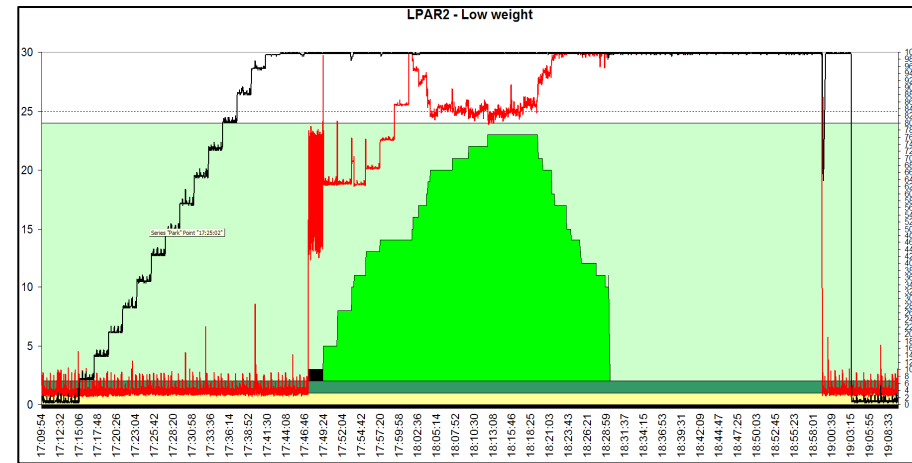



- Availability
 - OA37737 z/OS V1.12, V1.13 closed 5/2012

OA37736 Test Scenario



- z196 with 24 CPs; 4 LPARs with 24 logical CPs
- LPAR1/2: low weight (1 VH+1 VM); LPAR2/3: high weight (9VH + 1 VM)
- Start work on LPAR1 first, then LPAR2, LPAR3, LPAR4
- Result shows that unparked VLs are reduced as high weight LPARs have demand



HiperDispatch related WLM APARs

APAR	Description	Close Date	Affects
OA37736	Unused capacity enhancements	5/2012	Small plus large partitions with unused capacity.
OA36459	Modify PARK/UNPARK algorithm to become more sensitive for smaller partitions	10/2011	Smaller partitions at low CEC utilizations
OA35428	Introduces new option to manage CICS environments in a work receiver/consumer model	09/2011	Installation running CICS- only workloads Can be amplified by HiperDispatch
OA35860	Correct calculation of CEC free capacity. The CPU consumption of the *PHYSICAL* partition was not included. This can lead to too many “unpark” operations under rare cases.	06/2011	Systems with high *PHYSICAL* time
OA35989	Correct overflow condition of CEC free capacity	05/2011	Small systems running on big CECs with very high unused capacity

- **HiperDispatch=YES is the default when running z/OS 1.13 on z196 or above**
- Older releases will still have NO as default - even on z196

OA39941: Eliminate large frame shortage critical messages (currently open – subject to change)

- Prior to OA39941 SRM issued critical messages
 - IRA120E LARGE FRAME SHORTAGE
 - IRA121E CRITICAL LARGE FRAME SHORTAGE
 - IRA122I LARGE FRAME SHORTAGE RELIEVEDwhen the LFAREA allocation exceeded 80 or 95 percent, respectively.

- OA39941
 - eliminates these messages and introduces
IRA137I 100% OF THE LARGE FRAME AREA IS ALLOCATED

- Recommendation
 - Check for existing automation of IRA120E/IRA121E
 - Monitor LFAREA usage
 - RMF Mon I Paging Activity report
 - RMF Mon III STORx reports
 - SMF71LRX, SMF71LRA

OA38280: Reduced Sampling Overhead

- Prior to OA38280 SRM used an IVSK instruction for sampling the “Performance Block” (PB)
- IVSK instruction more costly on IBM z196 or z114 than previously
- Resulted in increased WLM address space CPU utilization when *many* PBs had to be sampled
 - Usually only the case when many subsystem instances started, typically in test systems
- Recommendation
 - Apply PTF
 - Beyond that:
 - In some environments it may be worth while to review how many subsystems are started and require sampling of many PBs:
 - CICS (MAXTASK)
 - DB2 MSTR (1000+)
 - DBM1 (500+)

OA27940: New IEAOPTxx Parameter RTPIFACTOR

- WLM Sysplex routing services are not only based on capacity but can also reflect server specific performance and health
 - Namely IWM4SRSC / IWMSRSRS SPECIFIC requests honor server performance index
- For very challenging goals the performance index (PI) may be pretty bad (>>1)
 - May not be true for all sysplex members, or routing of long running work leads to oscillation effects across systems
- New parameter RTPIFACTOR allows to control the influence of the performance index on routing recommendations
- Recommendation: Use RTPIFACTOR *only if* a negative effect of the PI on routing has been proven.
Defining achievable goals should be first step.
- The RTPIFACTOR value can be accessed in the IRARMCTZ control block as
RMCTZ_ROUTING_PI_FACTOR_PERCENTAGE

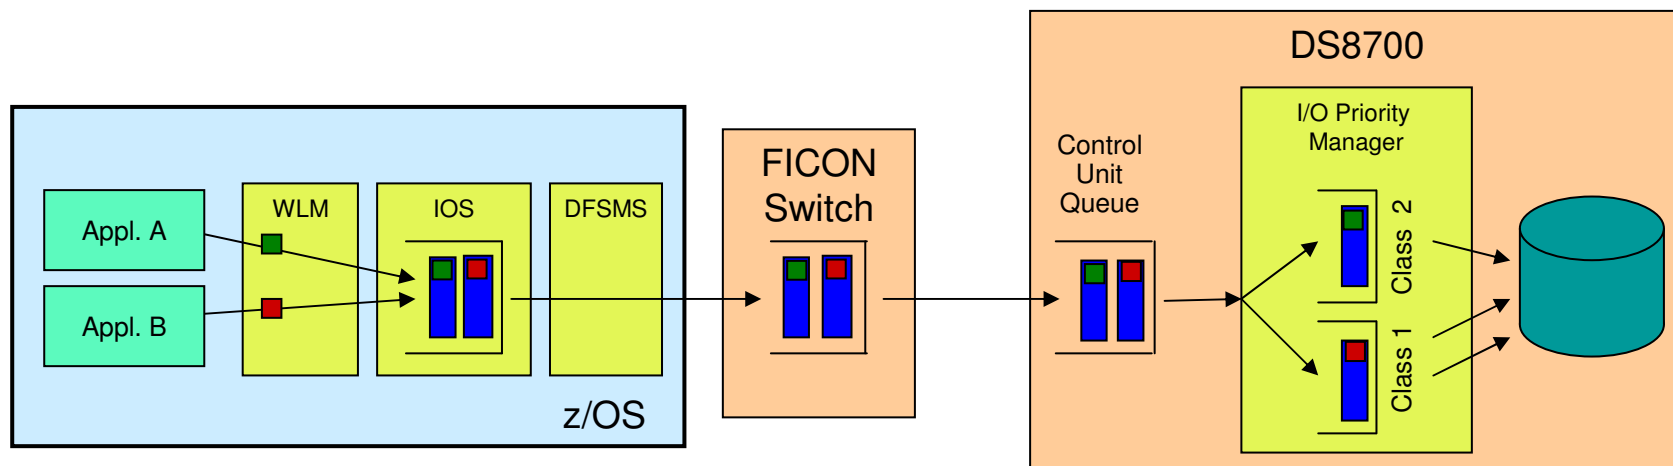
RTPIFACTOR=	
<u>100</u>	Performance index is fully considered in routing recommendations
0	Performance index is not considered in routing recommendations
1-99	Influence of PI on routing recommendations is reduced

Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals
- HiperDispatch and Other Relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
- WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview

WLM Support for I/O Priority Manager in DS8K Series

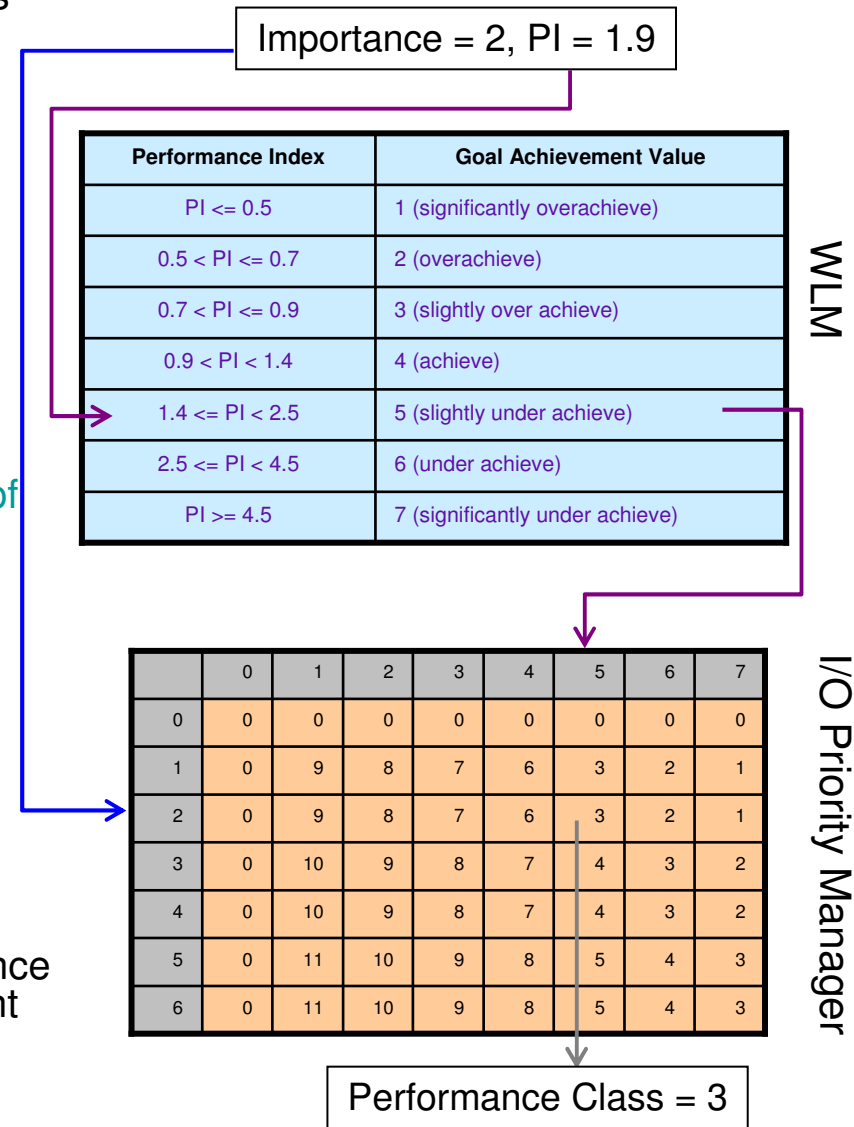
- WLM collaborates with the I/O Priority Manager in DS8700 & DS8800 storage servers.
 - This feature is supported on IBM System Storage® DS8700 and DS8800 series, and requires a DS8000 licensed machine code
- WLM sends I/O Priority Manager information about the goal fulfillment and importance of z/OS workloads (service classes).
- Passing these performance parameters to the storage server enables the I/O Priority Manager to determine which I/O requests are more important than others and which I/O requests need to be processed faster to fulfill the performance goals defined for the corresponding workload in z/OS.
- Using the passed information from WLM, the I/O Priority Manager throttles I/O requests of workloads which exceed their goals to help I/O requests of workloads which do not fulfill their goals.
- New IEAOPT parameter **STORAGESERVERMGT={YES|NO}**



WLM Support for I/O Priority Manager in DS8K Series

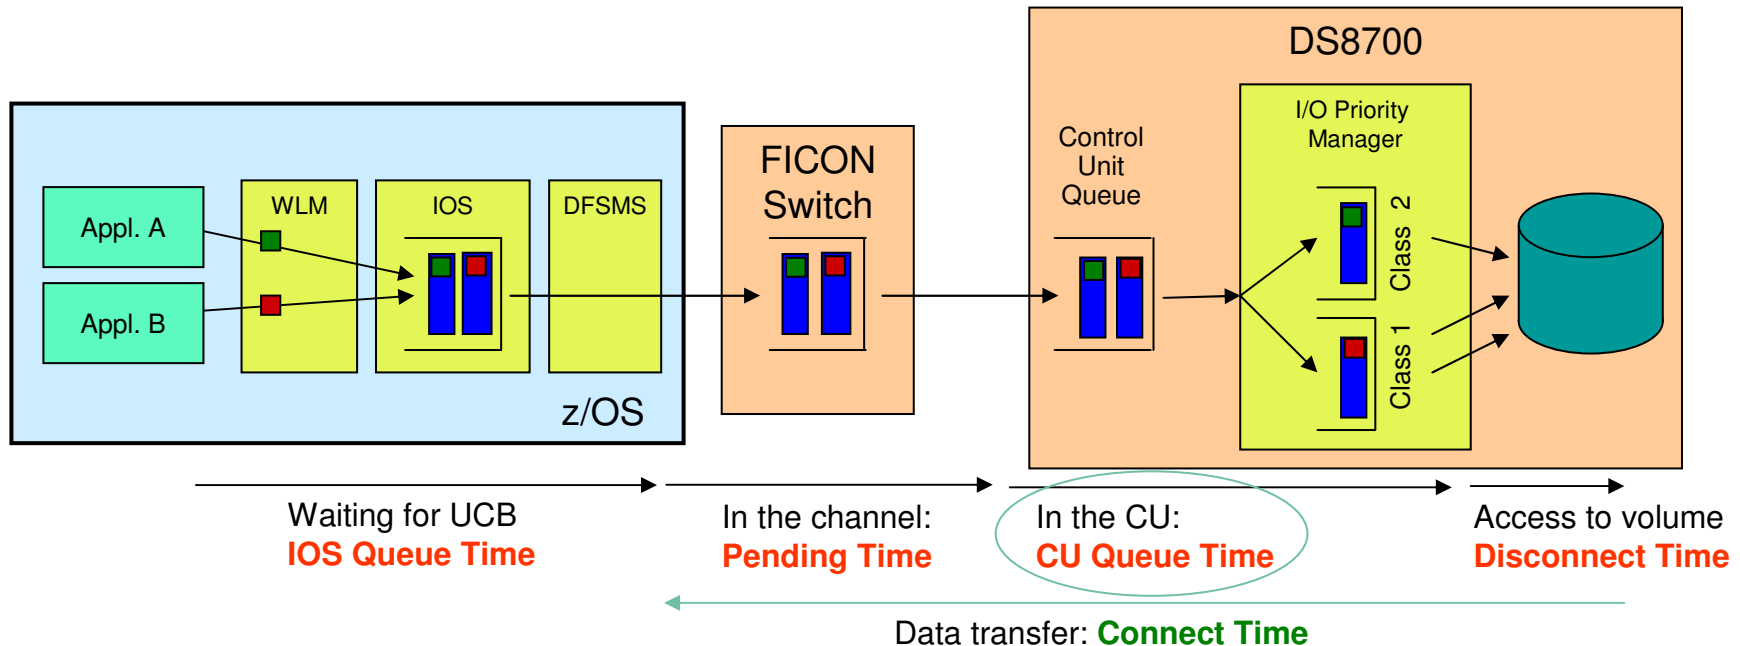
Goal Achievement Data

- WLM derives goal achievement data for service class periods as follows
 - Response Time goal:
 - Importance of period
 - Goal Achievement Value is derived from Performance Index (PI)
 - **Dynamic management considering goal achievement of service class**
 - Velocity goal:
 - Importance of period
 - Velocity level
 - **Static management considering specified goal of service class**
 - System service classes
 - Importance: 0
 - Goal Achievement Value: 0 (no monitor)
 - **No management**
 - Discretionary goal:
 - Importance: 6
 - Goal Achievement Value: 1
 - **Static management considering goal type**
- I/O Priority Manager assigns I/O request a Performance Class corresponding to the passed Goal Achievement Data
- Each Performance Class is associated with a certain maximum throttling level



WLM Support for I/O Priority Manager in DS8K Series

Throttling



- I/O Priority Manager
 - impacts how long an I/O request has to wait for access to the **volume**
 - does **not** impact an I/O request **if it is served from the storage server cache**
 - induced throttling delays are reported as **CU Queue Time**
- WLM excludes CU Queue Time when calculating I/O delays for service class periods with velocity goal to avoid oscillations of the performance index
 - **If WLM support for I/O Priority Manager is turned on, you may have to adjust the velocity goals if you have significant CU Queue Times in your environment**

WLM Support for I/O Priority Manager in DS8K Series

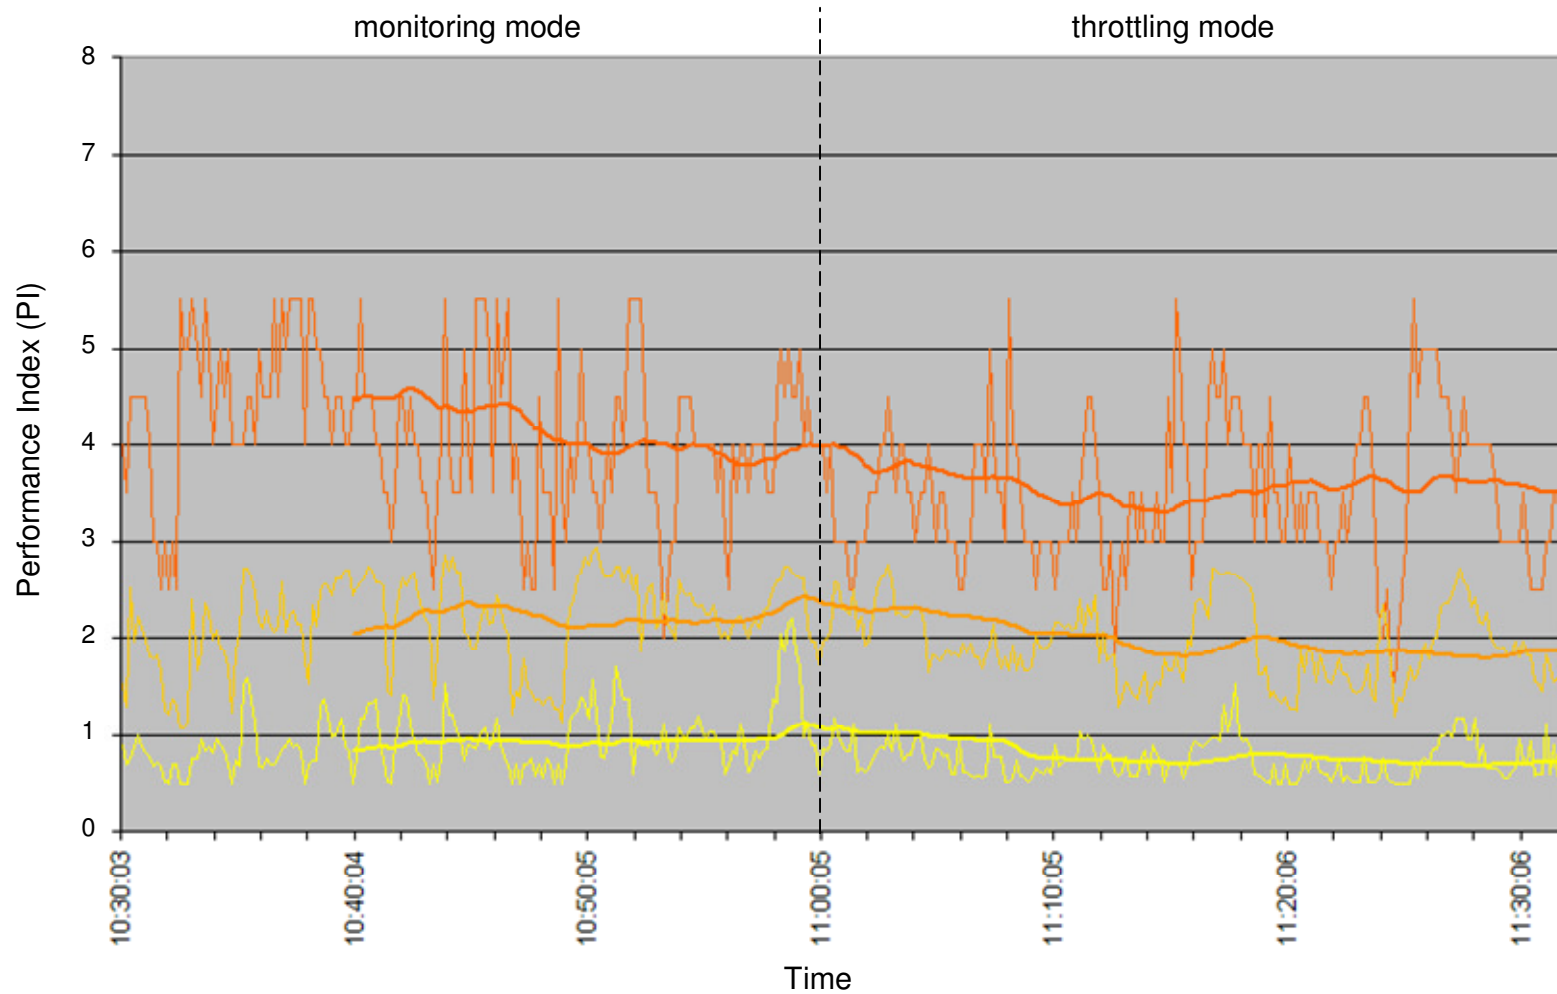
Test Results



- Test Environment
 - Five CPCs running z/OS V1.13 in nine LPARs
 - DS8000
 - Workloads running simultaneously
 - ATM transaction workload (CICS, IMS, DB2)
 - Middleware OLTP workloads (CICS, IMS, VSAM, DB2)
 - IMS data sharing workloads
 - CICS VSAM/RLS data sharing workload
 - CICS VSAM/NRLS workload
 - DB2 data sharing workloads
 - DB2 large storage utilization workload

- Test results were described in a white paper:
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102074>

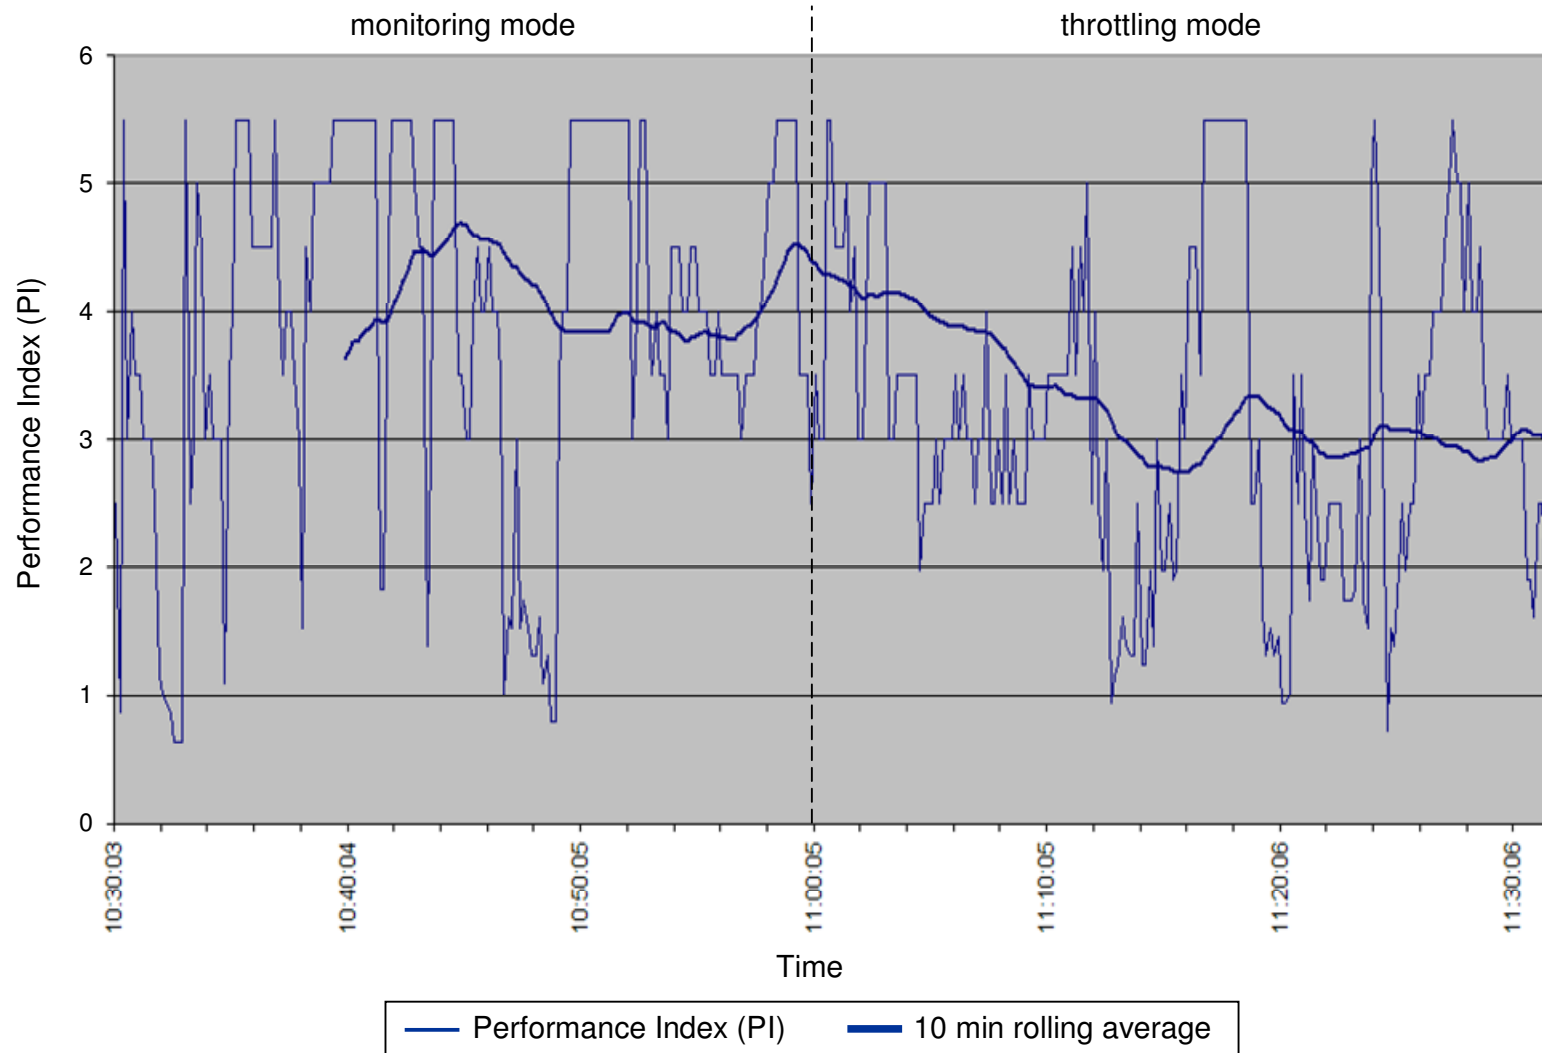
WLM Support for I/O Priority Manager in DS8K Series Test Results



- PI of Imp 1 RT service classes improved by 10%
- PI of Imp 2 RT service classes improved by 22%
- PI of Imp 3 RT service classes improved by 23%

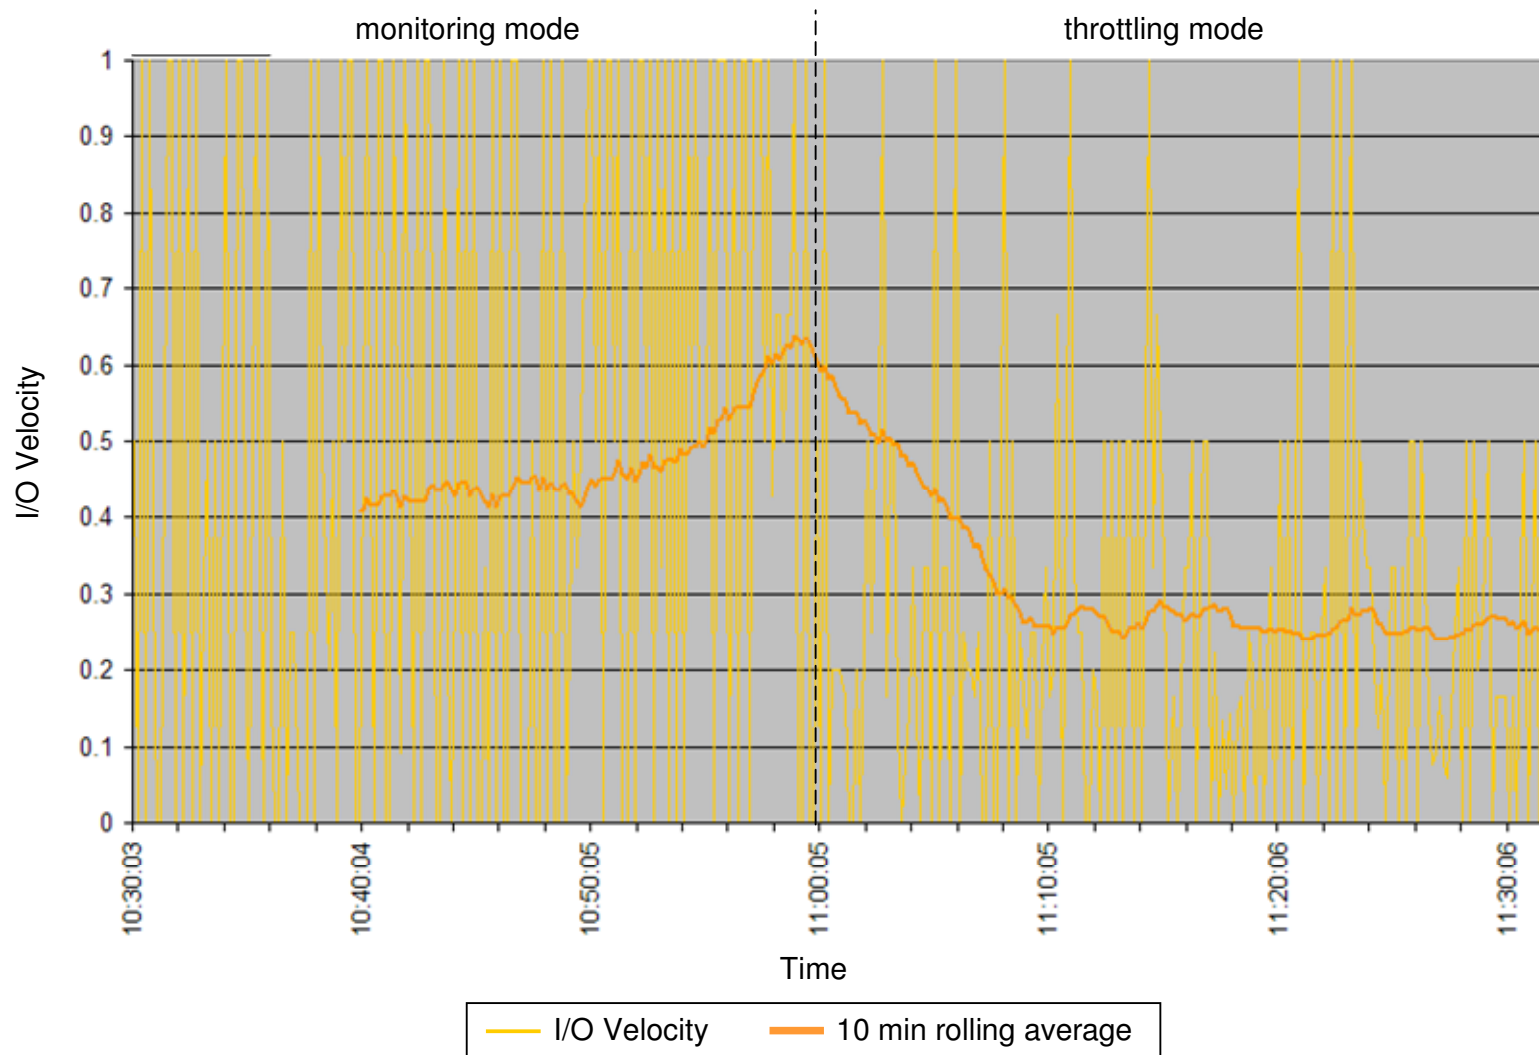
— Importance 1 — Importance 2 — Importance 3 — 10 min rolling average

WLM Support for I/O Priority Manager in DS8K Series Test Results



- PI of CICS service class improved from 4.3 to 3.0
- Achieved response time reduced from 800ms to 570ms

WLM Support for I/O Priority Manager in DS8K Series Test Results



- PI of Velocity service class STCI2V40 stayed in same range
- But the I/O velocity decreased from 0.63 to 0.27


WLM Support for I/O Priority Manager in DS8K Series Availability

Function	z/OS V1.13	z/OS V1.12	z/OS V1.11	Older Releases
WLM Support for I/O Priority Manager in DS8000 series	OA32298	OA32298	OA32298	

- WLM Support for I/O Priority Manager
 - Has to be enabled via OPT parameter STORAGESEVERMGT=YES
 - Default is STORAGESEVERMGT=NO
 - Only active when WLM policy specifies I/O Priority Management = YES

- The I/O Priority Manager feature is associated with DS8K R6.1.5

Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals
- HiperDispatch and Other Relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
-  ▪ WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview

WLM Support for IBM zEnterprise 196

- IBM zEnterprise 196 (z196)
 - STSI instruction no longer returns the alternate CPU capability
 - CPU adjustment factors are now calculated based on the Model Capacity Ratings by the machine
 - Supplies additional information about speed change
 - Speed changes may occur due to model changes (*capacity level*), or to physical processor tact (*cycle steering*)

- WLM
 - uses the new MSU values to calculate pricing adjustment factors
 - introduces message IWM064I to explain the reason for a processor speed change
 - makes new HW information available via public data areas IRARCT, IRARMCT, IRARMCTZ and via SYSEVENT QVS

- Existing Message:

IWM063I WLM POLICY WAS REFRESHED DUE TO A PROCESSOR SPEED CHANGE

- Depending on the reason for the speed change one of the following messages will be issued on *when running on z196 or later hardware*:

- IWM064I THE SYSTEM IS RUNNING AT NOMINAL CAPACITY.

- IWM064I THE SYSTEM IS RUNNING AT NOMINAL CAPACITY; MODEL CONVERSION OCCURRED.

- IWM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF A MANUAL CONTROL SETTING.

- IMM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF A MACHINE EXCEPTION CONDITION.

- IWM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF A NON-EXCEPTION MACHINE CONDITION.

- IWM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF AN EXCEPTION CONDITION EXTERNAL TO THE MACHINE.

- IRARCT
 - RCTPCPUA; RCTPCPUA_actual; RCTPCPUA_nominal; RCTPCPUA_scaling
 - See previous chart
- IRARMCTZ
 - RMCTZ_Capacity_Change_Time
 - Time when the capacity was last changed
 - RMCTZ_Capacity_Adjustment_Indication
 - When zero, the indication is not reported. When in the range 1-99, some amount of reduction is indicated. When 100, the machine is operating at its normal capacity. Primary CPUs and all secondary-type CPUs are similarly affected
 - RMCTZ_Capacity_Change_Reason
 - Indicates the reason which is associated with the present value contained in RMCTZ_Capacity_Adjustment_Indication
 - RMCTZ_CAI_IPL
 - Capacity adjustment indication at IPL
 - RMCTZ_CCR_IPL
 - Capacity change reason at IPL
 - RMCTZ_nominal_CPMP
 - Nominal CPU adjustment factor (similar to RMCTCPMP but for nominal speed)

- Sysevent QVS: (IRAQVS and IWMQVS.H)
 QVSCECCapacityStatus possible values:

Constant	Mnemonic	Description
0	QvsCecCapStatUndef	QvsCecCapacityStatus is undefined (not supported by hardware)
1	QvsCecCapStatNominal	Machine is running at nominal capacity
2	QvsCecCapStatRedIntentional	Machine is running with reduced capacity due to a manual control setting. (e.g. power saving mode, customer initiated)
3	QvsCecCapStatRedMachEx	Machine is running with reduced capacity due to a machine exception condition (e.g. cooling problem)
4	QvsCecCapStatRed MachNonEx	Machine is running with reduced capacity due to a machine non-exception condition (e.g. firmware update)
5	QvsCecCapStatRed EnvCond	Machine is running with reduced capacity due to an exception condition external to the machine (e.g. ambient temperature exceeded specified maximum)

- IWMRCOLL

- IWMWRCAA

- RCAAADJCCPU CPU adjustment factor
 - RCAAADJCCPUNOM Nominal CPU adjustment factor
 - RCAAADJCCEC CEC adjustment factor

WLM Support for IBM zEnterprise 196

SMF record type 70 subtype 1 (CPU Activity) – CPU control section				
Offset	Name	Length	Format	Description
204 xCC	SMF70NCR	4	Binary	Nominal model-capacity rating in MSU/hour. When non-zero, this value is associated with the nominal model capacity as identified in field SMF70MDL. When field SMF70CAI contains a value of 100, this value equals the value in field SMF70MCR.
208 xD0	SMF70NPR	4	Binary	Nominal permanent model-capacity rating. When non-zero, this value is associated with the nominal permanent model capacity as identified in field SMF70MPC. When field SMF70CAI contains a value of 100, this value equals the value in field SMF70MPR.
212 xD4	SMF70NTR	4	Binary	Nominal temporary model-capacity rating. When non-zero, this value is associated with the nominal temporary model capacity as identified in field SMF70MTC. When field SMF70CAI contains a value of 100, this value equals the value in field SMF70MTR.
216 xD8	SMF70CAI	1	Binary	Capacity-adjustment indication. When zero, the indication is not reported. When in the range from 1 to 99, some amount of reduction is indicated. When 100, the machine is operating at its normal capacity. Temporary capacity changes that affect machine performance (for example, CBU or OOCoD) are not included.
217 xD9	SMF70CCR	1	Binary	Capacity-change reason. Valid if SMF70CAI is non-zero. When 0, no capacity change took place. When 1, the capacity change is due to the setting of a manual control. When greater than 1, the capacity change is due to an internal machine condition or due to an external machine exception.

WLM Support for IBM zEnterprise 196

SMF72 Enhancements

SMF record type 72 subtype 3 (Workload Activity) – Workload manager control section				
Offset	Name	Length	Format	Description
172 xAC	R723MADJ	4	Binary	Adjustment factor for CPU rate
248 xF8	R723NADJ	4	Binary	Nominal adjustment factor for CPU rate

WLM Support for IBM zEnterprise 196: Availability

Function	z/OS V1.13	z/OS V1.12	z/OS V1.11	z/OS V1.10	z/OS V1.9
New message IWM064I API enhancements	+	OA30968	OA30968	OA30968	
New MSU computation	+	OA30968	OA30968	OA30968	OA30968
New Programming Interface (IRARMCTZ)	+	+	OA31201	OA31201	

Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals
- HiperDispatch and Other Relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
- WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview



Temporary Capacity Reporting via SYSEVENT REQLPDAT

- **SYSEVENT REQLPDAT** was changed to return capacity information about IBM z10 (and later) capacity settings:
 - permanent capacity information
 - The base capacity of the machine
 - temporary capacity data
 - Replacement Capacity: Capacity Backup (CBU), or Planned Event (CPE)
 - Additional Capacity: On/Off Capacity on Demand (OOCoD)
 - *The differentiation is relevant for potential license cost or entitlement impact*

Temporary Capacity Reporting via SYSEVENT REQLPDAT

▪ IRALPDAT new data fields

- LPDATMODELCAPIIDENT
 - The 16-character EBCDIC model-capacity identifier of the configuration.
- LPDATMODEL
 - The 16-character EBCDIC model identifier of the configuration. If not valid, field LPDatModelCapIdent represents both the model-capacity identifier and the model.
- LPDATMODELPERMCAPIIDENT
 - The 16-character EBCDIC model-permanent capacity identifier of the configuration.
- LPDATMODELTEMPCAPIIDENT
 - The 16-character EBCDIC model-temporary capacity identifier of the configuration.

Permanent+OOCoD + (CUBU+PE)

H/W model

Permanent configuration

Permanent + OOCoD

Temporary Capacity Reporting via SYSEVENT REQLPDAT

▪ IRALPDAT new data fields

➤ LPDATMODELCAPRATING

- When non-zero, an unsigned integer (“MSU rating”) as identified by the model-capacity identifier. There is no formal description of the algorithm used to generate this integer.

Permanent+OOCoD + (CUBU+PE)

➤ LPDATMODELPERMCAPRATING

- When non-zero, an unsigned integer (“MSU rating”) as identified by the model-permanent-capacity identifier

Permanent configuration

➤ LPDATMODELTEMPCAPRATING

- When non-zero, an unsigned integer (“MSU rating”) as identified by the model-temporary-capacity identifier.

Permanent + OOCoD

Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals
- HiperDispatch and Other Relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
- WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview

z/OSMF Workload Management

The new WLM Control Center in z/OSMF V1.12 and V1.13

- Policy editor
 - Simplified creation and editing of WLM policies supported by **best practice checks**
 - Support for review and investigation of WLM policies
- Policy repository
 - WLM policies are stored in a repository integrated in the z/OSMF file system
 - Policies can be exported to the local workstation or a host data set as well as imported from a file or a host data set
 - Policies or best-practice recommendations can be printed for further study
 - Integrated operation history makes manual tracking superfluous
- Installation and activation of WLM policies
- Monitoring of the WLM status in the sysplex
- Administration and operation tasks can be performed simultaneously
 - Simplified migration: Policy elements can be copied from one service definition to another
 - Simplified operation: You can start to edit a policy, interrupt the editing to activate a policy, and then continue with the editing without losing the context
- z/OSMF Workload Management synchronizes automatically with z/OS WLM
- Different authorization levels: View, Install, Modify (V1.13)

z/OSMF Workload Management Service Definition Repository



- Integrated repository for service definitions
- Service definitions can be
 - Imported
 - Exported
 - Printed
 - Viewed or edited
 - Created or Copied
 - Installed on the sysplex
- Indications
 - If service definition is installed and active
 - If service definitions are being viewed or edited
 - If messages exist for a service definition

The screenshot shows the IBM z/OS Management Facility Workload Management console. A table of service definitions is displayed with columns for Name, Description, Activity, Sysplex, Messages, Last Modified (GMT), and Modified By. The 'TESTFIX1' row is selected, and a context menu is open over it, with 'Install and Activate...' highlighted. A callout box points to this menu item with the text 'Click to view, edit, print, install a service definition'. Another callout box at the top of the table area says 'Store all service definitions in one repository'.

Name	Description	Activity	Sysplex	Messages	Last Modified (GMT)	Modified By
R12RGRP2	D10 WLM ZOSMF.POLICY.R12RGRF			Error	Feb 23 2010 2:30:24 PM	bmor
RTDST3	Copy of RTDST3				Mar 21 2001 8:23:19 PM	bmair
RTDST3	SDS1 copy 5			Warning	Jan 31 2010 10:49:38 PM	wirag
SampdeF	Sample WLM Service Definition 62				Sep 24 2007 8:48:22 AM	tblatt
SampdeF (Installed & Active)	Sample WLM Service Definition 57		WLMMPLEX	Information	Feb 1 2010 8:52:56 PM	wirag
SPMinTst				Information	Jan 26 2010 3:50:46 PM	wirag
T13DEC07	add/remove SAP DB2s				Dec 13 2007 9:01:59 PM	ks56551
TEST15				Information	Jan 12 2010 12:43:29 PM	wirag
TESTFIX1					Oct 3 2006 11:40:35 AM	sig011
TESTSD1				Error	Dec 30 2009 6:42:37 PM	wirag
WLM_BOF1	Large			Error	Feb 19 2010 5:12:08 PM	debug22
WLM_BOF2						
WLM_DESC	WL De					
WLM001	Service					
WLM000						
WLM700						
wImpol01	policy					
WLMPOL03				Warning	Jan 13 2010 9:19:00 AM	wirag
WLMPOL04				Warning	Feb 2 2010 12:09:54 AM	wirag
WLMSTT	AVT R10+R11RAS				Jul 8 2008 10:38:57 AM	bmor
WSCWLMDE	WSC Sample WLMServiceDefinition			Error	Jan 27 2010 4:05:01 AM	p3asru

z/OSMF Workload Management Editing Service Definitions

- Simplified creation, modification and review of service definitions
 - Policy elements are presented in tables
 - Tables can be filtered and sorted
 - Direct editing of policy elements within tables
 - Best-practice hints are displayed automatically while specifying policy elements
 - Several service definitions can be opened simultaneously
 - Cut, Copy, Paste of policy elements between service definitions

The screenshot shows the IBM z/OSMF Workload Management interface in Mozilla Firefox. The main content area displays a table of Service Classes. A context menu is open over the selected row (STTCL7), showing options like 'New Period', 'Expand', 'Collapse', 'Cut to Clipboard', 'Copy to Clipboard', 'Paste Periods', 'Move Periods', 'Delete...', and 'View Cross References'. Three callout boxes provide instructions: one points to a warning icon in the Velocity column, another points to the 'Copy to Clipboard' option, and a third points to the 'Delete...' option.

Name	Period	Importance	Duration	Goal Type	Response Time	Percentile	Velocity	CPU Critical	Resource Group
*STTCL3	1	*3		*Velocity			*91	*No	
*STTCL4	1	*2		*Average Response Time	*00:00:03.000			*No	
*STTCL5	1	*2		*Average Response Time	*00:00:30.000			*No	
*STTCL6	1	*2		*Average Response Time	*00:02:00.000			*No	
*STTCL7	1	*3	*3000	*Vel					
*STTCL8	1	*3	*3000	*Vel					
*STTCL8	2	*3		*Vel					

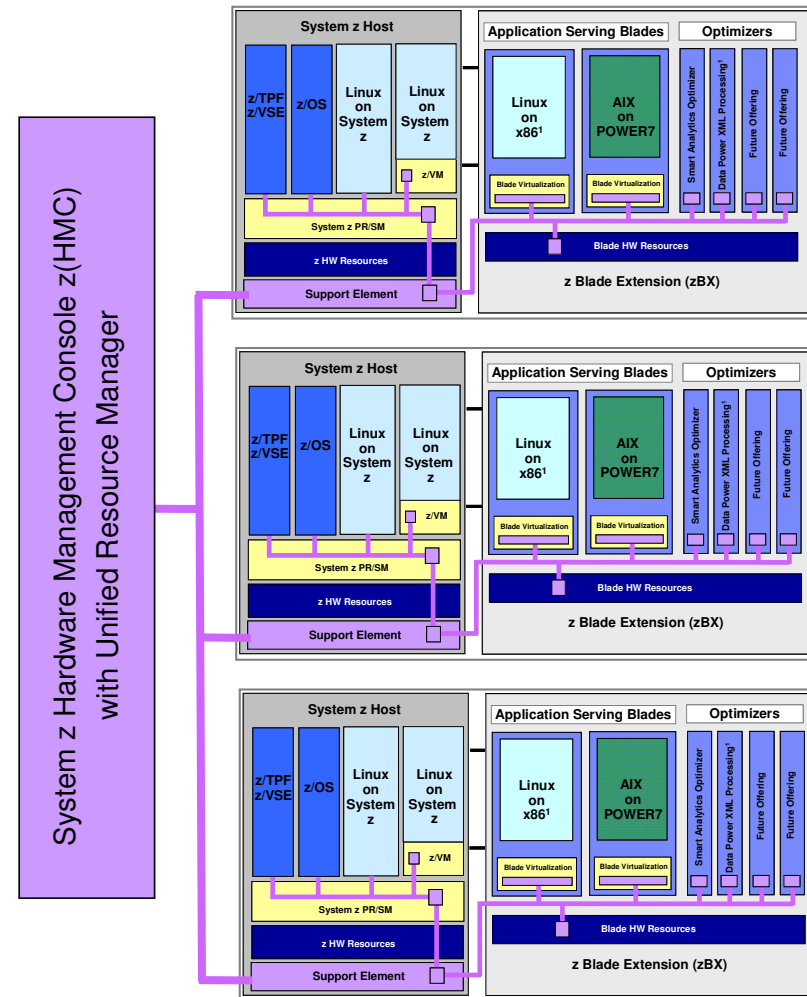
Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals
- HiperDispatch and Other Relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
- WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview



- Ensemble
 - A zEnterprise Ensemble is a collection of zEnterprise Nodes managed as a single virtualized pool of server resources
 - Native LPAR and z/VM Virtual Images
 - Power VM Virtual images
 - IBM Smart Analytics Optimizer for DB2
 - A zEnterprise Node can be a member of at most one Ensemble

- zEnterprise Unified Resource Manager
 - allows for the management and optimization of a zEnterprise Ensemble as a single resource pool
 - System z Hardware Management Console (HMC) is management console
 - Ensemble-wide scope of responsibility

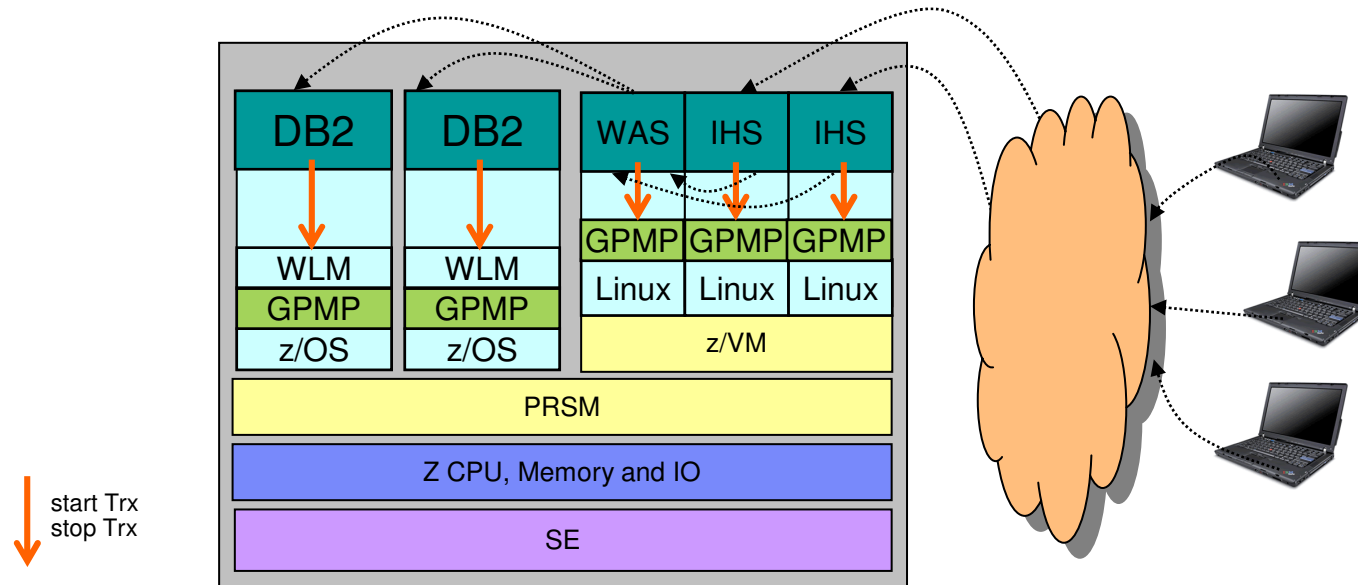


zEnterprise Platform Performance Manager

- Platform management component responsible for goal-oriented resource monitoring, management, and reporting across the zEnterprise Ensemble
 - Core component responsible for definition and implementation of goal-oriented management policy
 - Workload monitoring and reporting based on management policy
 - Extend goal oriented approach of z/OS WLM to platform managed resources
 - Orchestration of autonomic management of resources across virtual servers
 - Provide Intelligent Resource Director like function across the zEnterprise
 - Management functions will evolve over time
 - Pushes management directives to the Support Element, Hypervisors, and OS agents as required across the zEnterprise
- Integration of HMC console support
 - Integrated UI for monitoring, display of workload topology relationships, status alerts, etc.
 - Definition of Performance Management Goals and Policy Administration
- Functionality integrated into the zEnterprise Unified Resource Manager
 - Code structured and packaged as System z firmware
 - Inter-Component communication over trusted internal platform management network

zEnterprise Platform Performance Manager

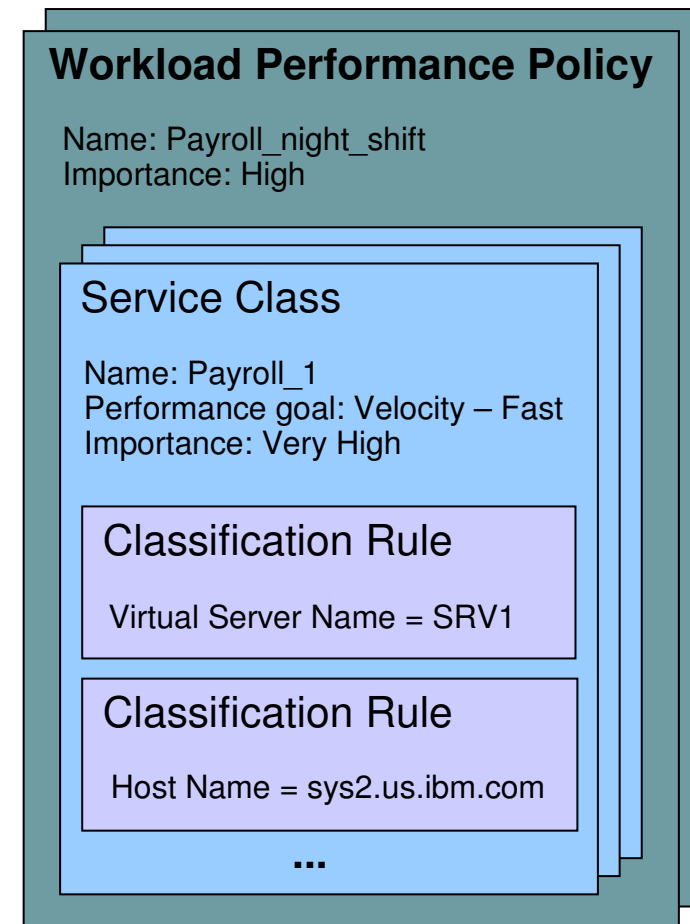
Resource management based on understanding of overall workload flow



- Applications / middleware has to be instrumented with ARM – Application Response Measurement (Open Group Standard) to collect transaction statistics
 - Enables to monitor the flow of transactions
 - Enables to monitor transaction response times and processing statistics

- OS Agent – guest platform management provider (GPMP)
 - is required to identify individual units of work
 - collects data about processes / address spaces and transactions
 - passes data to Platform Performance Manager
 - On z/OS the data is collected by WLM

- Defines performance goals for virtual servers in a workload
 - Conceptually like a 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



- The *guest platform management provider* (GPMP) is the interface between the Unified Resource Manager and the z/OS Workload Manager

- GPMP
 - passes to WLM information about the platform wide performance goals of workloads in which the z/OS is participating
 - sends data provided by WLM to the HMC for platform performance monitoring
 - Server configuration and high level performance statistics collected on z/OS
 - Aggregated transaction response time and resource data for the ARM-instrumented applications

- WLM
 - supports GPMP configuration and management by new WLM service definition options, commands, and messages
 - manages the GPMP address space (start, stop, and restart)
 - displays GPMP status information
 - collects and aggregates performance measurements for GPMP

WLM support for Unified Resource Manager Service Definition Enhancements for GPMP

- z/OS V1R12 introduces WLM functionality level LEVEL025 to support Unified Resource Manager and GPMP
- Unified Resource Manager Service Classes can be classified to WLM service and report classes by specifying classification rules for subsystem EWLM
 - ~~Work qualifier ETC (EWLM transaction class name) is no longer supported~~
 - Work qualifier type ESC (EWLM service class name) is used to correlate Unified Resource Manager service classes with WLM service or report classes

Action	-----Qualifier-----			-----Class-----	
	Type	Name	Start	Service	Report
	1	ESC	Booking		
	2	ESC	System	9	
	3	ESC	GoldServ	15	
	4	ESC	ice	23	SERVCLS3

DEFAULTS:

- Although z/OS V1R12 simply disregards ETC classification rules, you have to delete them the next time you modify the EWLM subsystem type classification rules
 - Message **IWMAM726 ETC is not a recognized qualifier type** is displayed when pressing F3=Exit
 - Rows with ETC rules have to be deleted before F3 becomes successful

WLM support for Unified Resource Manager Service Definition Enhancements for GPMP

To configure GPMP

- 1) Select option 11 on the Definition Menu
- 2) Specify Guest Platform Management Provider settings

```

File Utilities Notes Optio
-----
Functionality LEVEL025 Def... WLM Appl LEVEL025
Command ==>

Definition data set . . . : none

Definition name . . . . . MYDEF01 (Required)
Description . . . . . Production 01

Select one of the
following options. . . . . 11
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
    
```

After GPMP settings defined, functionality level changes to 25

```

GPMP-Settings Notes Options Help
-----
Guest Platform Management Provider (GPMP) Settings
Command ==>


Guest Platform Management Provider activation: 2 1. NO
                                                2. YES

Names of systems to be excluded:

SYS04   SYS09   _____
_____
_____
_____
_____
    
```

Specifies the systems in the sysplex on which the GPMP should not be started automatically

Specifies whether you want to start the GPMP address space automatically when a WLM policy is activated



WLM support for Unified Resource Manager

GPMP Configuration and Management



- On **policy activation**
 - WLM checks whether the service definition has valid GPMP settings
 - If activate=yes and system name not specified on excluded-list, GPMP is started automatically

- Also, you can use the **MODIFY WLM** command
 - To start the GPMP on a system
 - To stop the GPMP on a system
 - Intended for recovery actions. Recommended is to manage GPMP through WLM

- Once you stopped the GPMP manually, the GPMP switches into “manual mode”. It is not automatically restarted even if a WLM policy with a valid GPMP configuration gets activated
 - Status maintained until next IPL

WLM support for Unified Resource Manager

GPMP related Commands



- Use the **MODIFY WLM,GPMP** command to start, stop, and modify the guest platform management provider:

- **F WLM,GPMP,START**

- Indicates that you want to start the GPMP

```
16.55.59 WLMG          f wlm, gpmp, start
16.55.59 WLMG STC00752 $HASP373 HVEMCA  STARTED
16.55.59 WLMG STC00752 IEF403I HVEMCA -  STARTED -  TIME=16.55.59
```

- **F WLM,GPMP,STOP**

- Indicates that WLM stops the currently active GPMP instance

```
17.03.39 WLMG          f wlm, gpmp, stop
17.03.39 WLMG STC00753 IEF404I HVEMCA -  ENDED -  TIME=17.03.39
17.03.39 WLMG STC00753 $HASP395 HVEMCA  ENDED
```

- **F WLM,GPMP,TRACE=NONE|LOW|MEDIUM|HIGH,DEST=FILE| MEMORY**

- Enables you to change the GPMP internal tracing level “on the fly” and to change the destination of the trace (file or memory)

WLM support for Unified Resource Manager

GPMP related Commands



- DISPLAY WLM command extensions:

```
IWM025I 11.42.45 WLM DISPLAY 231
ACTIVE WORKLOAD MANAGEMENT SERVICE POLICY NAME: BASEPOL
ACTIVATED: 2010/02/18 AT: 12:57:55 BY: BMAI FROM: TRX2
DESCRIPTION: Base policy for system test
RELATED SERVICE DEFINITION NAME: Ra1fpol
INSTALLED: 2010/02/18 AT: 12:57:48 BY: BMAI FROM: TRX2
WLM VERSION LEVEL: LEVEL025
WLM FUNCTIONALITY LEVEL: LEVEL025
WLM CDS FORMAT LEVEL: FORMAT 3
STRUCTURE SYSZWLM_WORKUNIT STATUS: CONNECTED
STRUCTURE SYSZWLM_EBAE2097 STATUS: CONNECTED
```

To display system and GPMP status information, enter:

D WLM,SYSTEMS,GPMP

```
STATE OF GUEST PLATFORM MANAGEMENT PROVIDER (GPMP): ACTIVE
```

SYSNAME	*MODE*	*POLICY*	*WORKLOAD MANAGEMENT STATUS*
TRX1	GOAL	BASEPOL	ACTIVE
TRX2	GOAL	BASEPOL	ACTIVE

SYSNAME	*GPMP STATUS*
TRX1	INACTIVE
TRX2	ACTIVE

To display whether ARM is enabled or disabled, enter:

D WLM,AM

```
IWM075I 11.45.43 WLM DISPLAY 233
ARM SERVICES ARE ENABLED
GUEST PLATFORM MANAGEMENT PROVIDER JOBNAME=HUEMCA ASID=0032
GPMP POLICY IS ACTIVE
NUMBER OF REGISTERED PROCESSES=3, APPLICATIONS=1
```

- The existing **MODIFY WLM,AM=DISABLE|ENABLE** command is not changed, but the logic for DISABLE/ENABLE changed in the following way:
 - Disabling ARM (Application Response Measurement) will terminate a running GPMP
 - Manually starting the GPMP (using the MODIFY WLM,GPMP,START command) when ARM is disabled will result in message IWM078E
 - Activating a WLM policy that contains valid GPMP settings will not result in the start of a GPMP instance, if ARM is disabled
 - The state of the GPMP will be displayed as “DISABLED”, if ARM is disabled
 - If ARM is enabled again, the state of the GPMP will change to “STOPPED”. To start the GPMP again, it has to be started manually

WLM support for Unified Resource Manager Prerequisites

- Hardware
 - z196 with zEnterprise Unified Resource Manager
 - Guest platform management provider on z/OS cannot be started on pre-z196 servers
 - If started on pre-z196 servers, message **IWM078E GUEST PLATFORM MANAGEMENT PROVIDER CANNOT BE STARTED, FUNCTION NOT AVAILABLE** is issued on the console
- Software
 - z/OS V1R12 and OA30928
 - For z/OS V1R10 and V1R11: OA30928

Agenda

- Transaction Management Enhancements
 - Non Shell Enclave Server Management
 - CICS Region / Response Time Management
 - Response Time Distribution for Execution Velocity Goals
- HiperDispatch and Other Relevant APARs
- WLM Support for I/O Priority Manager in DS8K Series
- WLM Support for IBM zEnterprise 196
- Temporary Capacity Reporting via SYSEVENT REQLPDAT
- z/OSMF Workload Management
- WLM support for Unified Resource Manager
- Capacity Provisioning Update Summary
- WLM Tools Overview

Capacity Provisioning Enhancements in z/OS V1.13



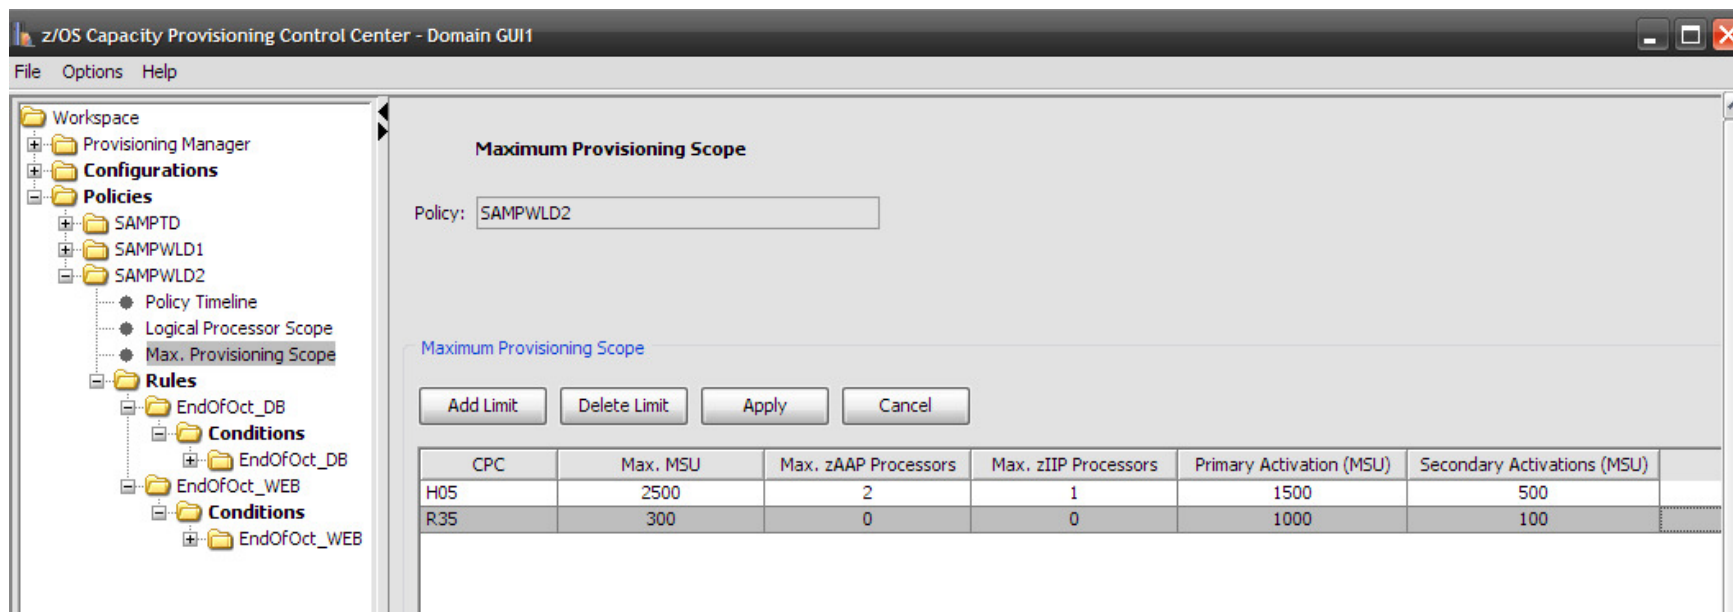
- Capacity Provisioning management enhancements
 - **Provisioning increments** allow for faster or more aggressive provisioning
 - **Recurring time condition** support allows to define recurring time windows
 - Allows to avoid ENABLE and DISABLE commands
 - Statement of Direction to withdraw support for the SNMP
 - z/OS BCPii is the recommend protocol

- Control Center Enhancements
 - Support the 32- and 64-bit versions of Microsoft Windows 7 Professional Edition

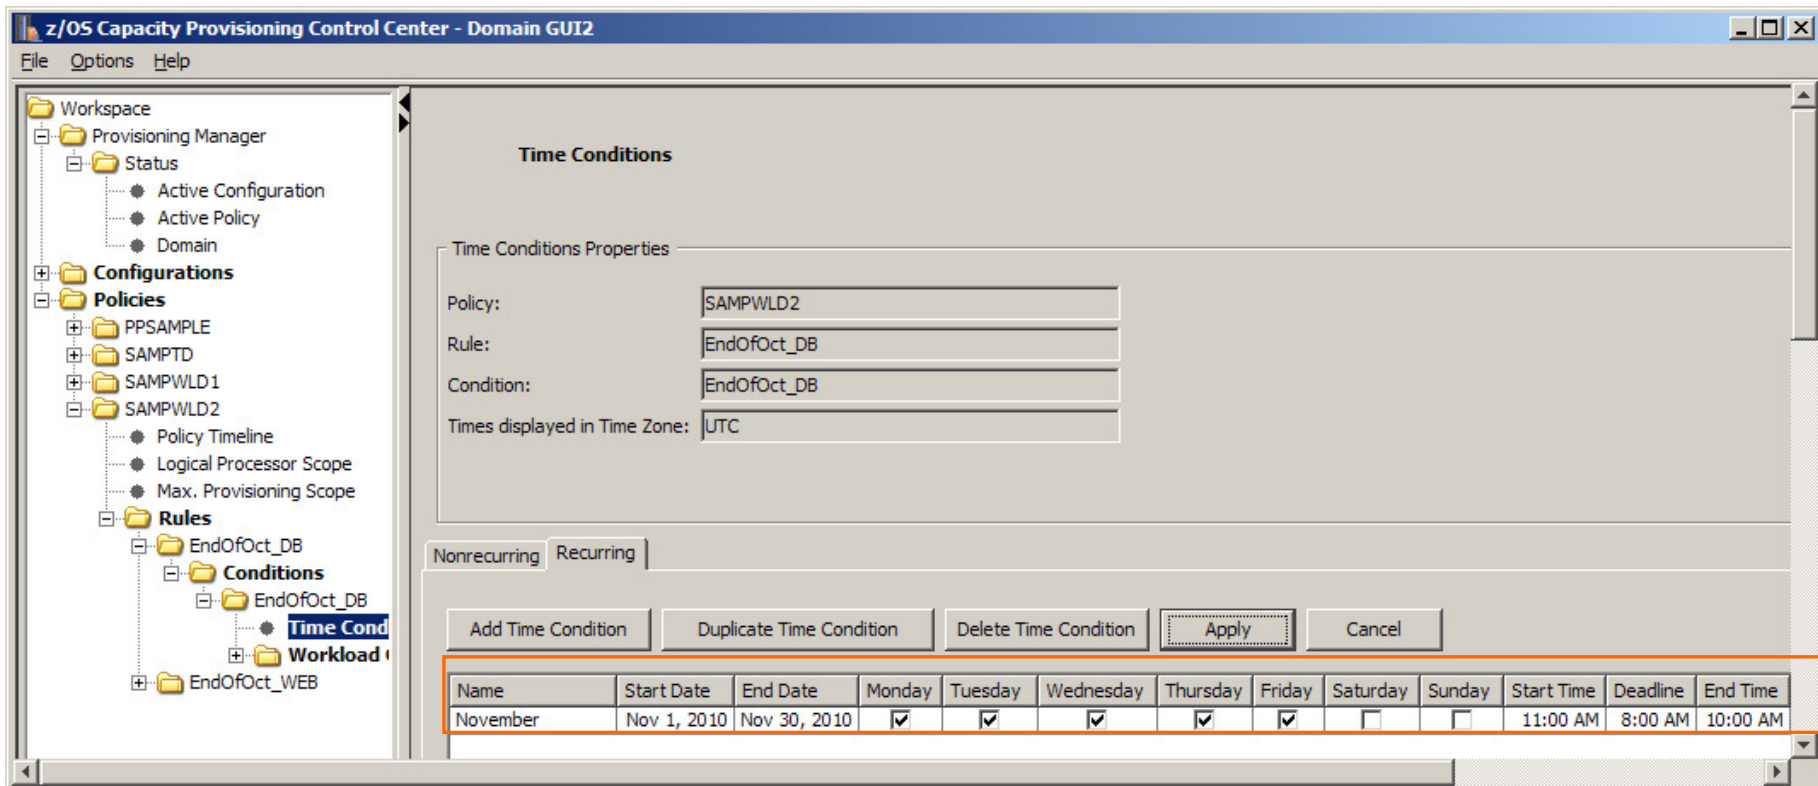
- **New with z/OSMF V1.13:**
 - Capacity Provisioning monitoring task



- Previously CPM increased capacity in small increments
 - On full speed models it would have usually added one processor at a time
- Now CPM supports primary and secondary activation quantum
 - Primary quantum added for first activation on a given CPC
 - Secondary added on subsequent activations
 - Definition is made on “Maximum Provisioning Scope” Panels.
 - Only general purpose capacity supports primary and secondary quantum at this time.
 - Retrofit to z/OS V1.11, V1.12



- Previously the CPM policy supported only fixed time intervals
 - Defined by start date/time and end date/time
- Starting with z/OS 1.13 CPM plans to support (weekly) recurring time conditions
 - Defined by start date, end date, start time, end time and day of week to which it applies
 - Plan to retrofit to z/OS V1.11, V1.12



Capacity Provisioning Support of zEnterprise 196 Static Power Save Mode



- Commands to disable or enable static power save mode:

Syntax

► DISABLE—POWERSAVE—CPC=*name*—
 |—D—PS—

► ENABLE—POWERSAVE—CPC=*name*—
 |—E—PS—

- Existing reports are extended to report on power-save capability, and whether power-save mode can currently be enabled

```
CPC R35 with record * is enabled (default enabled)
CPC is matched with serial 000020089F25 since 07/23/2010 13:32:13
Hardware is of type 2817 with model M49
Current model is 722 with 2119 MSU, 1 zAAPs, and 1 zIIPs
No usable OoCoD record available
Power save mode is enabled
```

- If power-save mode cannot be re-enabled in current period: “Power save mode is disabled and not allowed”
- For CPCs supporting static power save mode the Provisioning Manager will not consider adding capacity based on the active policy while in power save mode
 - Already activated temporary capacity may be deactivated
 - ACTIVATE RESOURCE and DEACTIVATE RESOURCE commands are not affected by power save mode.
- Requires Automate version of the zEnterprise Unified Resource Management suite
 - CPC Power Saving setting must be “Custom”
 - Also see “Controlling IBM zEnterprise 196 Static Power Save Mode via MVS Capacity Provisioning Manager (CPM)” at <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101869>

Capacity Provisioning New Function Overview



<i>z/OS (CPM) release</i> Function	V1.13	V1.12	V1.11	V1.10	V1.9
<i>Capacity increments Recurring time conditions</i>	+	OA35284	OA35284	OA35284 Toleration-only	
<i>CPCCC Windows 7 support</i>	+				
<i>z196 Static Power Save Mode</i>	+	OA30433	OA30433	OA30433	OA30433 Toleration-only
<i>Samples for security definitions (CIM, CPM, z/OSMF)</i>	+	+			
<i>Control Center reporting enhancements, Windows Vista™ support</i>	+	+			
<i>CICS/IMS transaction classes support</i>	+	+	OA29641	OA29641	
<i>RMF provider can locate DDS dynamically</i>	+	+	OA31118	OA31118	
<i>z/OS BCPii Support and Logical Processor Mgmt</i>	+	+	+	OA25426 OA24945	
<i>z/OSMF release</i> Function	V1.13	V1.12	V1.11	V1.10	V1.9
<i>CPM Status Monitoring</i>	+				

धन्यवाद

Hindi

多謝

Traditional Chinese

ขอบพระคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Thank You

English

Obrigado

Brazilian Portuguese

شكراً

Arabic

Grazie

Italian

多谢

Simplified Chinese

Danke

German

Bedankt

Dutch

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

감사합니다

Korean

WLM Tools: A Summary

Tool	Name	Description	Content	Support
SVDEF	Service Definition Formatter	Uses output from WLM Administrative Administration to display content of service definition in a workstation spreadsheet	Excel/workstation	Not updated anymore but still available on WLM Tools page
WSE	Service Definition Editor	Allows to create, modify, retrieve and install WLM service definitions	Java program on workstation	YES!! Available
WLMQUE	Application Environment Viewer	Allows to monitor WLM Application Environments	ISPF Tool	YES!! Available
WLMOPT	OPT Display	Display WLM/SRM OPT Parameters	IPF Tool	No!! Obsoleted by RMF in z/OS V1.11

These tools are located on the z/OS WLM homepage:
<http://www.ibm.com/servers/eserver/zseries/zos/wlm/tools/>

WLM Tools Service Definition Editor

Phased out. Superseded by z/OSMF WLM task.

The screenshot shows the WLM Service Definition Editor interface. The main window displays a table with columns for Name, ServiceClasses, Period, Goal, Im..., Duration, ResponseTime, Perce..., Level, ResourceGr..., CPU..., and Description. A context menu is open over the table, showing options: Insert, Insert Before, Insert After, Replace by, Copy, Cut, and Delete. The 'Insert Before' option is selected, and a sub-menu is visible with 'AverageResponseTime' and 'PercentileResponseTime' options. Below the table, there are two error messages:

No	Description	Element
1	Importance value can not be null	Workload "WKLTKK"/ServiceClass "B4V20STD"/Velocity (#1)
2	WLM may not distinguish between periods with equal importance and only slightly different velocity levels	Workload "WLDASC"/ServiceClass "A3V30STD"

WLM Tools

Display WLM/SRM OPT Parameter (WLM Tool, supported up to R10)

```

Command ==>                               Scroll ==> PAGE
.
.
.      WLM OPT Settings                      >SAVE<
System: AQFT      Version: z/OS 011100  OPT: FT  Time: not issued
OPT-Parameter:   Value:                 Description:
.
.
ABNORMALTERM      Yes  Abnormal term. used in routing rec.
BLWLTRPCT         5   CPU cap. to promote blocked work
BLWLINTHD         20  Time blocked work waits for help
CCCAWMT           3200,3200  AWM time value (defined, used)
ZAAPAWMT          3200,3200  AWM time value for zAAPs (def, used)
ZIIPAWMT          3200,3200  AWM time value for zIIPs (def, used)
CNTCLIST          No   Clist commands count individually
CPENABLE          10,30  LOW,HI thresh for % TPI int. x 100
DVIO              Yes  Specifies w/ directed VIO is active
ERV              1000,E6  Enq res. CPU Service and DP
HIPERDISPATCH    Yes,Yes  Hiperdispatch value(inOPT, Running)
IFAHONORPRIORITY  Yes   Specifies if CPs may help zAAPs
IIPHONORPRIORITY  Yes   Specifies if CPs may help zIIPs
INITIMP           0,FE  INITIMP value and DP for initiators
MCCAFCTH          400,800  LOW,HIGH central threshold
MCCFXEPR          92   % of storage fixed within first 16MB
MCCFXTPR          80   % of online storage fixed
PROJECTCPU        No   CPU projection for zAAPs and zIIPs
RCCFXTT           66,72  Low,High Logical MPL threshold
RCCFXET           82,88  Low,High Physical MPL threshold
RMPTTOM          1000  SRM invocation interval
STORAGENSDP       Yes  Set Non-swappable AS non dispatchable
STORAGEWTOR       Yes  Issue IRA221D and IRA421D
IRA405I           46,32,32  IRA405I warning level: 16M,2G,Tot
VARYCPU           No   VARYCPU is enabled
VARYCPUMIN        1   VARYCPUMIN value
WASROUTINGLEVEL    0   WebSphere Routing Level
.
.

```

WLM Tools

Display WLM/SRM OPT Parameter (RMF Monitor II OPT Report)

```

boewlm1 - wc3270
RMF - OPT Settings
Line 1 of 29
Command ==>
Scroll ==> PAGE
CPU= 4      UIC= 65K PR= 0      System= WLM1 Total

OPT: 00      Time: N/A
-- Parameter -- - Default - -- Value -- Unit ----- Description -----
ABNORMALTERM      Yes      Yes Y/N  Abnormal terminations in routing
BLWLINTHD         20      20 sec  Time blocked work waits for help
BLWLTRPCT         5        5 0/00  CPU cap. to promote blocked work
CCCAWMT          12000    12000 usec  Alternate wait management time
ZAAPAWMT         12000    12000 usec  AWM time value for zAAPs
ZIIPAWMT         12000    12000 usec  AWM time value for zIIPs
CNTCLIST          No       No Y/N   Clist commands count individually
CPENABLE         10,30|0,0  10,30 %  Threshold for TPI (low,high)
DVIO              Yes      Yes Y/N  Directed VIO is active
ERV              500      500/CB  SU      Enqueue residency CPU Service/DP
HIPERDISPATCH    No       No/No Y/N  Hiperdispatch is desired/active
IFAHONORPRIORITY  Yes      Yes Y/N  Allows CPs to help zAAPs
IIPHONORPRIORITY  Yes      Yes Y/N  Allows CPs to help zIIPs
INITIMP           0        0/FE #   INITIMP value/DP for initiators
IRA405I           70,50,50  70,50,50 %  Fixed storage of <16M,16M-2G,tot
MAXPROMOTETIME    6        6 *10s  Holder allowed to run promoted
VCCAFACTH        400,800   400,800 #   Threshold for storage (low,ok)
VCCFXEPR          92       92 %    Fixed storage threshold < 16 MB
VCCFXTPR          80       80 %    Fixed online storage threshold
PROJECTCPU        No       No Y/N   CPU projection for zAAPs, zIIPs
RCCFXET          82,88    82,88 %  Physical MPL threshold (low,high)
RCCFXTT          66,72    66,72 %  Logical MPL threshold (low,high)
RMPPTOM          1000|3000  3000 msec  SRM invocation interval
RTPIFACTOR        100      100 %   PI affects server routing weights
STORAGEENSWDP     Yes      Yes Y/N  Sets non-swap. ASID non-dispatch.
STORAGEWTOR       Yes      Yes Y/N  WTOR to cancel AS in shortage
VARYCPU           Yes      Yes Y/N  VARYCPU is enabled
VARYCPUMIN        1        1 #     VARYCPUMIN value
#ASROUTINGLEVEL    0        0 #     WebSphere routing level

F1=HELP      F2=SPLIT      F3=END      F4=RETURN      F5=RFIND      F6=SORT
F7=UP        F8=DOWN       F9=SWAP lis F10=LEFT      F11=RIGHT     F12=RETRIEVE
4B          X          T          IPY$1C09      002/015

```

WLM Tools

WLMOPT – WLM Application Environment Viewer

```

Command ==>                               Scroll ==> PAGE
Application Environment Monitor
Selection: >HELP< >SAVE< >OVW< >ALL< \AE=SYSBATCH
System: AQFT      Sysplex: MCLXCF01  Version: z/OS 011100  Time: 06:22:27

ApplEnv_ Type SubName_ WMAS Del Dyn NQ QLen Str Hav Unb Trm Min_ Max_ ICnt
SYSBATCH JES  JES2      0031 No No  3   0  0 12  0  0  0  0  0

WorkQue_ Del Wnt Hav ICnt QueIn_ QueOut QueLen QueTot_ Act_ Idl_
WLMLONG No  7  7  0  0  0  0  0  0  4  3
WLMSHORT No  3  3  0  0  0  0  0  0  2  0
COMBUILD No  2  2  0  0  0  0  0  0  1  1

SvAS Binding_ Ter Opr Btc Dem Have Jobname
0043 WLMLONG No No Yes No 1 BCNDEVD
0175 WLMLONG No No Yes No 1 ALLAEBS.2.SEAS.2.JBNI
0166 WLMLONG No No Yes No 1 SERV9956
0165 WLMLONG No No Yes No 1 SERV9955
015A COMBUILD No No Yes No 1 C90SPACE
0150 WLMLONG No No Yes No 1 INIT
0202 WLMLONG No No Yes No 1 INIT
0152 COMBUILD No No Yes No 1 INIT
0229 WLMSHORT No No Yes No 1 BMGX1$
0119 WLMLONG No No Yes No 1 INIT
0050 WLMSHORT No No Yes No 1 ALLAEBS.2.SEAS.11.JBNI
01A5 WLMSHORT No No Yes No 1 INIT

```

धन्यवाद

Hindi

多謝

Traditional Chinese

ขอบพระคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Thank You

English

Obrigado

Brazilian Portuguese

شكراً

Arabic

Grazie

Italian

多谢

Simplified Chinese

Danke

German

Bedankt

Dutch

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

감사합니다

Korean