



z/OS Workload Manager V1R11 Update

Speaker Name: Glenn Anderson
IBM Training for Systems



© 2009 IBM Corporation

Trademarks

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

APPN*	HiperSockets	OS/390*	VMESA*
CICS*	HyperSwap	Parallel Sysplex*	VSE/ESA
DB2*	IBM*	PR/SM	VTAM*
DB2 Connect	IBM eServer	Processor Resource/Systems Manager	WebSphere*
DirMaint	IBM e(logo)server*	RACF*	z/Architecture
e-business logo*	IBM logo*	Resource Link	z/OS*
ECKD	IMS	RMF	z/VM*
Enterprise Storage Server*	Language Environment*	S/390*	z/VSE
ESCON*	MQSeries*	Sysplex Timer*	zSeries*
FICON*	Multiprise*	System z	z10
GDPS*	NetView*	System z9	z9
Geographically Dispersed Parallel Sysplex	On demand business logo	System z10	
		TotalStorage*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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

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

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

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

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

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

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.



All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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





**ENABLING BUSINESS.
A THROUGH Z.**

2

WLM Update and Latest News


© 2009 IBM Corporation



Content

- ❑ **Blocked Workload Support**
- ❑ **Group Capacity Support and Enhancements for z/OS R11**
- ❑ **Manage selected components in service class SYSTEM**
- ❑ **Work-dependent enclaves**
- ❑ **Hiperdispatch**
- ❑ **Routing Support**
- ❑ **Enhanced Storage Monitoring**
- ❑ **WLM Tools**
- ❑ **Overview: Latest WLM Functions/APARs**
- ❑ **ITSO Redbook Page and Documentation**

3 | WLM Update and Latest News | © 2009 IBM Corporation



Blocked Workload Support: Objectives

- ❑ **Problem**
 - Work competes for resources, serialized by locks and latches
 - Low import work may hold a resource and high important work may have to wait for it
- ❑ **Solution**
 - Resource Manager monitor the resource and detect lock and latch contention
 - Use WLM interfaces (SYSEVENT ENQHOLD/ENQRLSE) to promote the lock holder in order to resolve the contention quick
- ❑ **But**
 - Not all resources can be monitored
- ❑ **System Assistance (WLM): Blocked Workload Support**
 - Recognizes blocked work
 - Work which doesn't show any progress for an elongated period of time
 - Allows this work to use a small amount of CPU periodically
 - With the hope to resolve existing (potential) resource contentions

4 | WLM Update and Latest News | © 2009 IBM Corporation

Blocked Workload Support: User Interface: IEAOPT

BLWLTRPCT	<p>Percentage of the CPU capacity of the LPAR to be used for promotion</p> <ul style="list-style-type: none"> <input type="checkbox"/> Specified in units of 0.1% <input type="checkbox"/> Default is 5 (=0.5%) <input type="checkbox"/> Maximum is 200 (=20%) <input type="checkbox"/> Would only be spent when enough units of work exist which need promotion
BLWLINTHD	<p>Specifies threshold time interval for which a blocked address space or enclave must wait before being considered for promotion.</p> <ul style="list-style-type: none"> <input type="checkbox"/> Minimum is 5 seconds. Maximum is 65535 seconds. <input type="checkbox"/> Default is 20 seconds.

5

WLM Update and Latest News

© 2009 IBM Corporation

Blocked Workload Support: User Interface: RMF

```

CPU ACTIVITY
...
BLOCKED WORKLOAD ANALYSIS
OPT PARAMETERS: BLWLTRPCT (%) 0.5 PROMOTE RATE: DEFINED 50000 WAITERS FOR PROMOTE: AVG 0.001
                  BLWLINTHD 60          USED (%) 95          PEAK 15

```

- Extensions of RMF Postprocessor CPU Activity and WLMGL reports with information about blocked workloads and the temporary promotion of their dispatching priority
- SMF record 70-1 (CPU activity) and SMF 72-3 (Workload activity)

6

WLM Update and Latest News

© 2009 IBM Corporation

Blocked Workload Support: Availability

Function	z/OS V1.11 <small>as previewed 2/2009</small>	z/OS V1.10	z/OS V1.9	z/OS V1.8
Promotion Enhancements for Chronic Contention	+	+		
Blocked Workload Support	+	OA27046	OA27046	OA27046
Promotion of Cancelled Jobs	+	+	+	

Related Support

- Promotion for Chronic Contentions
 - Introduces a new WLM service (IWMCTCN) which allows resource holders to tell WLM that long lasting contention situation exists
 - WLM will then manage the holder according to the goals of the most important waiter
 - Exploitation planned by DB2 in the future
- Promotion of Cancelled Jobs
 - The majority of the cancel process runs within the address space being cancelled
 - This support swaps in the address space and gives it a higher DP in order to assure that the work completes fast and frees system resources

7

WLM Update and Latest News

© 2009 IBM Corporation

Group Capacity: Objectives

Base: Defined Capacity for single LPAR

- Allows to control the software licenses based on a 4 hour rolling average consumption of the CPU service

Desired Extension

- Combine multiple partitions and control them together with a capacity limit
- When WLM decides that capping is required
 - Allow partitions to use the **unused** CPU service from those partitions which show only little demand

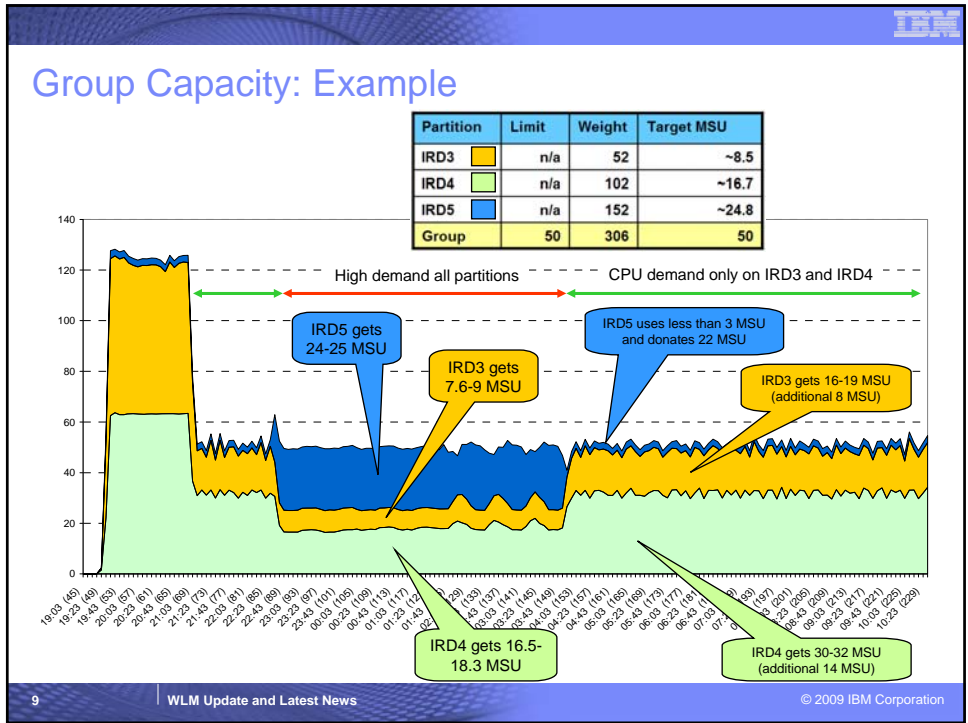
Solution

- Group Capacity Support
 - Makes better use of CPU capacity by guaranteeing that software licenses can be maintained

8

WLM Update and Latest News

© 2009 IBM Corporation



Group Capacity: Availability

Function	z/OS V1.11 as previewed 2/2009	z/OS V1.10	z/OS V1.9	z/OS V1.8
Group Capacity plus OA24096 Enhancements	+	+	OA24096 OA23230	OA24096 OA23230
z/OS Capacity Provisioning	+	+	OA20824	

- ❑ **OA24096**
 - Changes the behavior when then group limit is changed according to the behavior for an individual defined capacity limit
- ❑ **OA23230**
 - Corrects a storage overlay which will occurs when SMF 99 data is collected and a partition is dynamically activated via HCD
- ❑ **Short Comings of the existing Group Capacity Report**
 - Reporting was not sufficient to understand capping of partitions within a group
- ❑ **Related z/OS Functions**
 - z/OS Capacity Provisioning
 - Allows to activate additional CPU capacity via OoCoD in a controlled manner

10 | WLM Update and Latest News | © 2009 IBM Corporation

RMF z/OS 1.11 Enhancements for Group Capacity...

GROUP CAPACITY REPORT

z/OS V1R11 SYSTEM ID TRX1 DATE 02/26/2009 INTERVAL 05.00.000
 RPT VERSION V1R11 RMF TIME 11.00.00 CYCLE 1.000 SECONDS

GROUP-CAPACITY NAME	LIMIT	PARTITION	SYSTEM	-- MSU -- DEF ACT	WGT	---	CAPPING DEF WLM% ACT%	---	ENTITLEMENT - MINIMUM MAXIMUM		
RMFGRP	60	TRX1	TRX1	100	4	400	NO	25	23	40	60
	60	TRX2	TRX2	100	13	200	NO	100	46	20	60
TOTAL				17		600					

Field Heading	Meaning
CAPPING WLM%	Percentage of time when WLM considers to cap the partition
CAPPING ACT%	Percentage of time when capping actually limited the usage of processor resources for the partition

RMF z/OS 1.11 Enhancements for Group Capacity...

PARTITION DATA REPORT

z/OS V1R11 SYSTEM ID TRX1 DATE 02/26/2009 INTERVAL 05.00.000
 RPT VERSION V1R11 RMF TIME 11.00.00 CYCLE 1.000 SECONDS

MVS PARTITION NAME	TRX1	NUMBER OF PHYSICAL PROCESSORS	24	GROUP NAME	RMFGRP
IMAGE CAPACITY	60	CP	3	LIMIT	60
NUMBER OF CONFIGURED PARTITIONS	58	APP	1	AVAILABLE	43
WAIT COMPLETION	NO	IFL	18		
DISPATCH INTERVAL	DYNAMIC	ICF	1		
		IIP	1		

----- PARTITION DATA -----		-- LOGICAL PARTITION PROCESSOR DATA --				-- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --									
NAME	S	WGT	DEF	ACT	DEF	WLM%	NUM	TYPE	EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
TRX1	A	400	100	4	NO	0.0	3.0	CP	00.00.11.049	00.00.11.371	1.23	1.26	0.04	1.23	1.26
H05LP45	A	10	0	2	NO	0.0	2	CP	00.00.04.720	00.00.05.690	0.79	0.95	0.11	0.52	0.63
TRX2CFA	A	100	0	1	YES	0.0	1	CP	00.00.02.950	00.00.03.070	0.99	1.03	0.01	0.33	0.34
H05LP59	A	100	0	1	NO	0.0	3	CP	00.00.02.700	00.00.03.501	0.30	0.39	0.09	0.30	0.39
H05LP6A	A	10	0	9	NO	0.0	2	CP	00.00.23.742	00.00.26.331	3.96	4.39	0.29	2.64	2.93

Field Heading	Meaning
AVAILABLE	Long-term average of CPU service units which would be allowed by the limit of the capacity group but are not used by its members. If the value is negative, this capacity group is subject to capping.

RMF z/OS 1.11 Enhancements for Group Capacity...



Monitor III CPC Capacity report enhancements for capacity groups:

- The threshold applied when calculating the *remaining time until capping* is changed for partitions being member of a capacity group. Then, the MSU limit is derived from defined capacity and group capacity.
- A new metric *Remaining time until group capping* is provided which is the projected time until the usage of processor resources for one or more members of the group might be limited.
- The *4h unused group capacity average* is provided. This is the average available capacity for the group during the last 4 hours. If the value is negative, that is the members consumed more service than allowed, the group is subject to capping.
- Both new metrics are not displayed in the ISPF version of the report but within the Monitor III data portal. The *Remaining time until group capping* is also available as DDS performance metric.

13

WLM Update and Latest News

© 2009 IBM Corporation

RMF z/OS 1.11 Enhancements for Group Capacity...

- The remaining time until group capping is available as DDS performance metric.
- It can be displayed by the Monitor III Data Portal or with the RMF PM workstation client.

14

WLM Update and Latest News

© 2009 IBM Corporation

RMF z/OS 1.11 Enhancements for Group Capacity...

The screenshot displays the RMF Monitor III Data Portal for z/OS. The main report is titled "RMF Report [TRX2,MVS_IMAGE] : CPC (Central Processor Complex)". It shows various system metrics and a detailed table of LPARs. Three callouts highlight specific enhancements:

- Monitor III CPC report in Monitor III Data Portal displays the projected remaining time until image/group capping in the report header:** This points to the "Time until Group Capping: 14400" and "Time until Image Capping: 14400" fields in the report header.
- Average available capacity for the group during last 4 hours:** This points to the "# Average Group Capacity Average: 142" field.

LPAR Name	Defined MSU/h	Actual MSU/h	Capping Option	# Logical Processors Online	Logical Effective %	Logical Total %	LPAR Physical Mgmt %	Physical Effective %	Physical Total %	Line Type	# Online Processors Shared	# Online Processors Dedicated	Current LPAR Weight	Logical Processor Share %	Hyper Dispatch # High	Dispatch # Medium	Dispatch # Low	Operating System Name	LPAR Cluster Name	Initial Weight	Mr Weight
*CP				14.0			2.5	4.8	7.3	CS	14	0	820								
HDLP#6	0	1	NO	2.0	0.4	0.5	0.1	0.2	0.3	CP	2	0	10	2.4	N/A	N/A	N/A	BOEH0545			
HDLP#9	0	1	NO	3.0	0.3	0.4	0.1	0.2	0.3	CP	3	0	100	10.2	N/A	N/A	N/A	BOEH0559			
HDLP#0	0	8	NO	2.0	3.8	4.1	0.2	1.9	2.1	CP	2	0	10	2.4	N/A	N/A	N/A	BOEH0560			
TRX1	50	4	NO	3.0	1.2	1.3	0.0	0.9	1.0	CP	3	0	415	51.2	1	2	0	TRX1	TRX1PLEX	400	9
TRX2	60	4	NO	3.0	1.3	1.4	0.0	1.0	1.0	CP	3	0	185	90.2	0	1	2	TRX2	TRX1PLEX	200	9
TRX2CFA	0	2	YES	1.0	2.1	2.1	0.0	0.5	0.5	CP	1	0	100	48.7	N/A	N/A	N/A				
PHYSICAL								2.2		CY											
*CFCPOOL				1.0			1.4	56.8	58.2	IS	100	1	0								

Manage selected components in service class SYSTEM

- Problem**
 - Incorrect WLM classification of system work can lead to serious system problems and even outages
- Solution**
 - Manage certain system address spaces ALWAYS in service class SYSTEM independent of their classification in the WLM service definition
- Address spaces now forced into SYSTEM are:**
 - XCFAS, GRS, CONSOLE, IEFSCHAS, IXGLOGR, SMF, CATALOG, SMSPDSE, SMSPDSE1
- Annotation**
 - While the service class definition in the policy is ignored the report class definition is still honored

Work-dependent Enclaves

Background

- zIIPs allow middleware components to run a certain amount of their work “offloaded” from regular processors
 - The offload percentage is defined by the middleware component via a WLM interface (not generally published)
 - For example DB2/DDF exploits this feature t of some of their DDF workload

Behavior

- The offload percentage is an attribute of the enclave under which the unit of work runs
 - It is not possible to specify different offload percentages for different units of work running under the same enclave

Desired Behavior

- DB2/DDF wants to specify different offload percentages for the different units of work of a parallel query
- AND still wants to maintain the transactional context to run the units of work under the same “SRM Transaction” (enclave)

Solution


- Define a work-dependent enclave which is a spawn from an independent enclave but which still is closely related to it and maintains the scope of a “SRM Transaction” with the independent enclave
- BUT it allows the middleware to specify different offload values for each instance of the work-dependent enclaves

17

WLM Update and Latest News

© 2009 IBM Corporation

Work-Dependent Enclaves: Attributes



Independent and their associated work-dependent enclaves form one entity

- The CPU consumption is the total of all entities
 - A period switch also takes place for all entities at the same time
- The elapse time is the measured for the “independent” enclave
- A service class reset for the independent enclave also resets the associated work-dependent enclaves

A work-dependent enclave can only exist related to an independent enclave

- If it is created from a dependent enclave it is already created as a dependent enclave because dependent enclaves are already extensions to address space work

18

WLM Update and Latest News

© 2009 IBM Corporation

Work-dependent Enclaves: User Interface: SDSF

With SDSF support installed

Without SDSF support installed

Work-dependent Enclaves: Availability

Function	z/OS V1.11	z/OS V1.10	z/OS V1.9	z/OS V1.8
Work-dependent Enclaves	+	OA26104	OA26104	OA26104

- New function becomes available with APAR OA26104
- DB2 exploitation with APAR PK76676
- SDSF support with APAR PK74125
- RMF support with z/OS 1.11
- Expected availability: April 2009

Hiperdispatch: Motivation

Cache and memory latency on a hypothetical server

1 machine cycle
4 machine cycles
100+ machine cycles
200+ machine cycles
600+ machine cycles

CP L1 L1.5 L2 Main Memory

Book

Book

Compared to real world

1 cycle	Stuttgart inner city	1 km
4 cycles	Stuttgart train station to television tower	4 km
100+ cycles	Stuttgart – Heidelberg	120 km
200+ cycles	Stuttgart – Kaiserslautern	210 km
600+ cycles	Stuttgart – Paris	620 km

21 | WLM Update and Latest News | © 2009 IBM Corporation

Dispatching in an LPAR Environment

without Hiperdispatch or Hiperdispatch=NO

LPAR1 LPAR2

Work Units (TCBs and SRBs)

z/OS Dispatcher

Logical Processors

Logical Processors

PR/SM Dispatcher

Physical Processors

PR/SM

22 | WLM Update and Latest News | © 2009 IBM Corporation

Hiperdispatch: Introduction

□ Design Objective

- Keep work as much as possible local to a physical processor to optimize the usage of the processor caches
- As a result systems with high number of physical processors provide a much better scalability

□ Function: Hiperdispatch

- Interaction between z/OS and the PR/SM Hypervisor to optimize work unit and logical processor placement to physical processors
- Consists of 2 parts
 - In z/OS (sometimes referred as Dispatcher Affinity)
 - In PR/SM (sometimes referred as Vertical CPU Management)

23

WLM Update and Latest News

© 2009 IBM Corporation

HiperDispatch mode

□ PR/SM

- Supplies topology information/updates to z/OS
- Ties *high priority* logicals to physicals (gives 100% share)
- Distributes remaining share to *medium priority* logicals
- Distributes any additional service to unparked *low priority* logicals

□ z/OS

- Ties tasks to small subsets of logical processors
- Dispatches work to *high priority* subset of logicals
- Parks *low priority* processors that are not need or will not get service

- Hardware cache optimization occurs when a given unit of work is consistently dispatched on the same physical CPU

24

WLM Update and Latest News

© 2009 IBM Corporation


Hiperdispatch: z/OS Part

z/OS obtains the logical to physical processor mapping in Hiperdispatch mode

- Whether a logical processor has high, medium or low share
- On which book the logical processor is located

z/OS creates dispatch nodes

- The idea is to have 4 high share CPs in one node
- Each node has TCBs and SRBs assigned to the node
- Optimizes the execution of work units on z/OS



Hiperdispatch: Availability

Function	z/OS V1.11 <small>as previewed 2/2009</small>	z/OS V1.10	z/OS V1.9	z/OS V1.8
HiperDispatch Enhancements for zIIPs	+			
HiperDispatch	+	+	OA20418	OA20418
New Park/Unpark algorithm for smaller LPARs	+	OA26789	OA26789	OA26789

- Base Function: OA20418 or included with z/OS 1.10**
- Important: OA26789: New Park/Unpark Algorithm for smaller LPARS**
- For full information:**
 - Refer to presentation: xxxxxx on mm/dd/yy hh:mm

Routing Support Enhancements

□ Background

- WLM provides programming interfaces which allow routing applications to obtain information where to route work in a parallel sysplex environment

□ Problems

- Originally the routing was purely capacity based which didn't reflect whether the work was healthy on the systems
- The capacity based algorithm attempted to return the systems with the highest free capacity or most low capacity first
- The introduction of offload processors was not reflected in the routing recommendation

□ Solution

- Since z/OS 1.7 a series of changes have been introduced for WLM routing services

27

WLM Update and Latest News

© 2009 IBM Corporation

Routing Services: Summary of Changes (z/OS 1.7 and 1.8)

Service	Function	Remark	Description
IWMSRSRS	Existing function SELECT		Returns capacity for the system relative to other systems in the Sysplex (accounts for the number of registered server instances)
IWMSRSRS	New parameter SPECIFIC	(1)	Returns capacity for the registered server relative to all other registered servers of the same type in the Sysplex. With R9, contains information on general purpose, ZAAP and zIIP capacities. Considers: <ul style="list-style-type: none"> □ Goal Achievement (PI) □ Queue Time for Enclaves □ Health Indicator
IWMSRSG	New Parameter HEALTH		Allows the server which registers to provide a health indicator from ok=100 to not ok=0. The factor is considered as part of the weight. IWMSRSG HEALTH can be updated by the server at any time
IWM4HLTH	New Service		For address spaces which are not registered and which want to set a health status. This status is factored into IWM4SRSC return data
IWM4SRSC	New service New parameter ABNORM_COUNT	(1)	Returns capacity for another address space to which the request is provided by the registered server. Considers <ul style="list-style-type: none"> □ Goal Achievement (PI) □ abnormal termination rate expressed as the number of abnormal terminations (as passed to WLM by the IWMRPT interface) per 1000 total terminations.


(1) Goal achievement is derived from the service class the working is running in. Can be an enclave service class, a transaction manager service class or the service class to which the address space is classified too

28

WLM Update and Latest News

© 2009 IBM Corporation

Routing Services: Summary of Changes (z/OS 1.9)



Programming Interface: IWMWSYSQ

- Mostly used by customer applications to retrieve capacity information by importance level from the system
- Changes with z/OS 1.9
 - With EXTENDED_DATA=YES additional information is returned in the output area
 - The system level (0) which contains the total system capacity was added
 - Data now returned for all processor types
 - In addition:
 - Uniprocessor speed of a single processor
 - ZAAP and ZIIP normalization factors (deviation from regular processor speed if applicable)
 - EXTENDED_DATA=NO returns the output area as before (pre V1R9 layout)

29

WLM Update and Latest News

© 2009 IBM Corporation

Routing Services: Summary of Changes (z/OS 1.11)

Enhancements

- Introduce a new algorithm to compare systems by their equivalent usage of regular (or "pure") CPU SU
- Introduce cost factors for crossover of work from ZAAPs and ZIIPs to regular CPs
- Introduce different options for weighting of lower important work running on the same system

New parameters are added on services IWMSRSRS and IWM4SRSC

- METHOD=PROPORTIONAL|EQUIVALENT
- COST_ZAAP_ON_CP=factor and COST_ZIIP_ON_CP=factor
- IL_WEIGHTING=weighting_method
 - Default= is 0 (no weighting),
 - possible values are 1=square_root, 2=linear, 3=quadratic

Notice: All functions will be exploited by Sysplex Communication Server

30

WLM Update and Latest News

© 2009 IBM Corporation

Routing Services: Availability

Function	z/OS V1.11	z/OS V1.10	z/OS V1.9	z/OS V1.8
Capacity-based Routing for WAS	+	+	+	OA16486
Routing (IWMSRSRS) Support for >96 Servers	+	+	+	OA18531
Routing Enhancements: Cost factor and weighting	+			
Use same free capacity calculation for <90% as it already exists for >90%	+	OA27032	OA27032	OA27032

□ Capacity-based routing for WAS

- Before OA16486 only round robin routing was available
- New IEAOPT parameter: WASROUTINGLEVEL={0,1}

□ Free capacity calculation (OA27032)

- OA10006 changed the free capacity calculation for >90%
 - This was successful until Hiperdispatch required to adjust the calculation also for <90%

31

WLM Update and Latest News

© 2009 IBM Corporation

Enhanced Storage Monitoring

□ Problem

- Fixed, Pageable and Auxiliary storage shortages can lead to serious system problems (including system outages)

□ What needs to be done

- Identify a storage shortage when it occurs
- Identify the reason (causing application) of the storage shortage
- Give the installation a chance to react on a storage shortage when it occurs

□ Solution

- Introduce a new set of messages to warn the installation when auxiliary and pageable storage shortages occur
- Introduce a WTOR which allows the installation to cancel storage consumers
- Set storage consumers non-dispatchable to allow the installation to react on the situation
- Introduce a set of new programming interfaces (ENF signal) to allow applications to react on storage shortages

32

WLM Update and Latest News

© 2009 IBM Corporation

Enhanced Storage Monitoring: Availability

Function	z/OS V1.11	z/OS V1.10	z/OS V1.9	z/OS V1.8
Enhanced Storage Monitoring	+	+		
Support for >128GB Real Storage	+	+	+	+

- Enhanced Storage Monitoring is introduced for z/OS 1.10
- Support for >128GB Real Storage was introduced with z/OS 1.8
- Refer to “Enhanced Storage Monitoring Presentation” on Thursday for both parts

33

WLM Update and Latest News

© 2009 IBM Corporation

WLM Tools

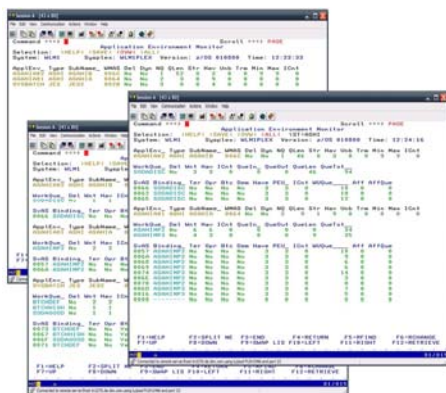
- WLM provides 3 tools**
 - WLMQUE: Application Environment Monitoring
 - WLMOPT: IEAOPT parameter viewer
 - WLM Policy Editor
- All tools can be downloaded “as is” from**
<http://www.ibm.com/servers/eserver/zseries/zos/wlm/tools/>

34

WLM Update and Latest News

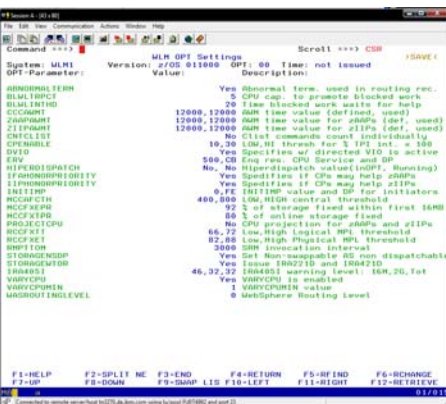
© 2009 IBM Corporation

WLM Tools: WLMQUE: Application Environment Monitoring



- ❑ **Middleware like WAS and DB2 exploits WLM Application Environments**
 - WLM Queue Manager Services
- ❑ **Problem**
 - No mechanism exists yet to monitor the work queues and server address spaces of WAS or DB2 using these environments
 - Over time the support became more and more complex which required the installation to get more knowledge about this mechanism
- ❑ **Solution**
 - Tool: WLMQUE
 - Gives an overview of all application environments in use for a system
 - Allows to monitor the activity and usage of work queues
 - Allows an installation to potentially change definitions to use WLM Queue Management Services more efficiently

WLM Tools: WLMOPT: View all IEAOPT parameters



- ❑ **Problem**
 - Currently no possibility exists to view the actual setting of all OPT parameters in the system
 - Also with Hiperdispatch some settings of the IEAOPT member might be changed during runtime
- ❑ **Solution**
 - Tool WLMOPT
 - Displays all actual settings and for some parameters also effected values of the system

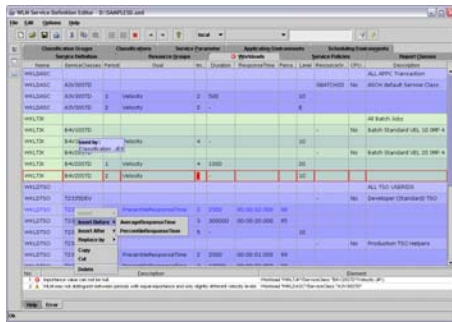
WLM Tools: WLM Policy Editor

Problem

- WLM Administrative Application is sometimes difficult to use

Solution

- WLM Policy Editor
 - Workstation based application which allows to manage WLM service definitions
 - Also supports the installation and activation of policies
 - Features
 - Information is represented as tables
 - Direct manipulation of tables
 - Display of relations between policy elements
 - Real-time error checking
 - Context-sensitive help



WLM Functions: CPU Management and Hardware

Function	z/OS V1.11 as previewed 2/2009	z/OS V1.10	z/OS V1.9	z/OS V1.8
HiperDispatch Enhancements for zLIPs	+			
HiperDispatch	+	+	OA20418	OA20418
Promotion Enhancements for Chronic Contention	+	+		
Blocked Workload Support	+	OA27046	OA27046	OA27046
Promotion of Cancelled Jobs	+	+	+	
Group Capacity plus OA24096 Enhancements	+	+	OA24096 OA23230	OA24096 OA23230
z/OS Capacity Provisioning	+	+	OA20824	

WLM Functions: Middleware Support

Function	z/OS V1.11	z/OS V1.10	z/OS V1.9	z/OS V1.8
Work-dependent Enclaves	+	OA26104	OA26104	OA26104
Start min. Number of Servers w/o Delay	+	+	+	
Capacity-based Routing for WAS	+	+	+	OA16486
Additional PB Delay States	+	+		

39 | WLM Update and Latest News | © 2009 IBM Corporation

WLM Functions: Storage and Routing Support

Function	z/OS V1.11	z/OS V1.10	z/OS V1.9	z/OS V1.8
Enhanced Storage Monitoring	+	+		
Capacity-based Routing for WAS	+	+	+	OA16486
Support for >128GB Real Storage	+	+	+	+
Routing (IWMSRSRS) Support for >96 Servers	+	+	+	OA18531
XCF Signaling Enhancements	+	+	OA20484	OA20484


40 | WLM Update and Latest News | © 2009 IBM Corporation

ITSO Redbook

System Programmer's Guide to: Workload Manager (SG24-6472)

- Content:
 1. Introduction
 2. How WLM works
 3. WLM functions
 4. Implementation and workload classifications
 5. Batch considerations
 6. TSO, STC, and APPC workloads
 7. DB2 workload considerations
 8. WebSphere Application Server workload considerations
 9. UNIX System Services considerations
 10. Transactional workload considerations
- Download via:

<http://www.redbooks.ibm.com/abstracts/sg246472.html?Open>



41 | WLM Update and Latest News | © 2009 IBM Corporation

Documentation

z/OS MVS Planning: Workload Management (SA22-7602)
<http://publibz.boulder.ibm.com/epubs/pdf/iea2w191.pdf>

z/OS MVS Programming: Workload Manager Services (SA22-7619)
<http://publibz.boulder.ibm.com/epubs/pdf/iea2w290.pdf>

Redbook – System Programmer's Guide to: Workload Manager (SG24-6472)
<http://www.redbooks.ibm.com/abstracts/sg246472.html?Open>

Internet Links

WLM <http://www.ibm.com/servers/eserver/zseries/zos/wlm/>

The Future Runs on System z



42 | WLM Update and Latest News | © 2009 IBM Corporation