



Systems and Technology Group

z/OS Workload Manager

The Latest and Greatest

Horst Sinram

WLM Development

IBM Corporation, Böblingen, Germany

<mailto:sinram@de.ibm.com>

08/23/2005

ON DEMAND BUSINESS™ = *make it happen now*

Trademarks

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

APPN*	HiperSockets	OS/390*	VM/ESA*
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 z9	
GDPS*	NetView*	TotalStorage*	
Geographically Dispersed Parallel Sysplex	On demand business logo	Virtualization Engine	

* 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.**

Agenda

- **New function overview**
- Multiple subchannel sets
- Dynamic handling of processor speed changes
- zSeries Application Assist Processors , WLM zAAP reporting support
- Routing enhancements, server specific load balancing
- Contention Management Enhancements
- WLM Support for Enterprise Workload Management
- WLM-managed batch initiator enhancements
- Sub-capacity pricing support
- Service definition coefficients
- User Friendly Interface for WLM Administrative Application
- Summary

Overview of New Functions (1)

Enhancement	z/OS R7	z/OS R6	z/OS R5	z/OS R4	z/OS R3
EWLM Monitoring Support 1	☺	OA07196			
Sysplex Routing Enhancements	☺				
Enhancements in computation of free capacity in LPAR environments	☺	OA10006			
WLM Support of Multiple Subchannel Sets (z9-109)	☺				
Dynamically handle processor speed changes	☺ OA12155	OA07510 OA12155	OA07510		
zAAP Reporting Support	☺				

Overview of New Functions (2)

Enhancement	z/OS R7	z/OS R6	z/OS R5	z/OS R4	z/OS R3
DB2 Latch Contention	☺	OA08949			
Enhanced Enqueue Promotion	☺				
SCRT Extensions for Dedicated Processors	☺				
Sub-Capacity Pricing Enhancements	☺			OW55509	
Enhancements for WLM-established Stored-procedures	☺		OA04555		
Stateful session placement	☺		OA04699		
Dynamic Application Environments	☺			OW54622	

Overview of New Functions (3)

Enhancement	z/OS R7	z/OS R6	z/OS R5	z/OS R4	z/OS R3
WLM-managed Initiator Start Considers System Affinities	OA10814				
Controlling Initiator Dispatch Priority in Goal Mode	☺			OW55344	
Enhanced Initiator Balancing (z/OS R4)	☺				

Agenda


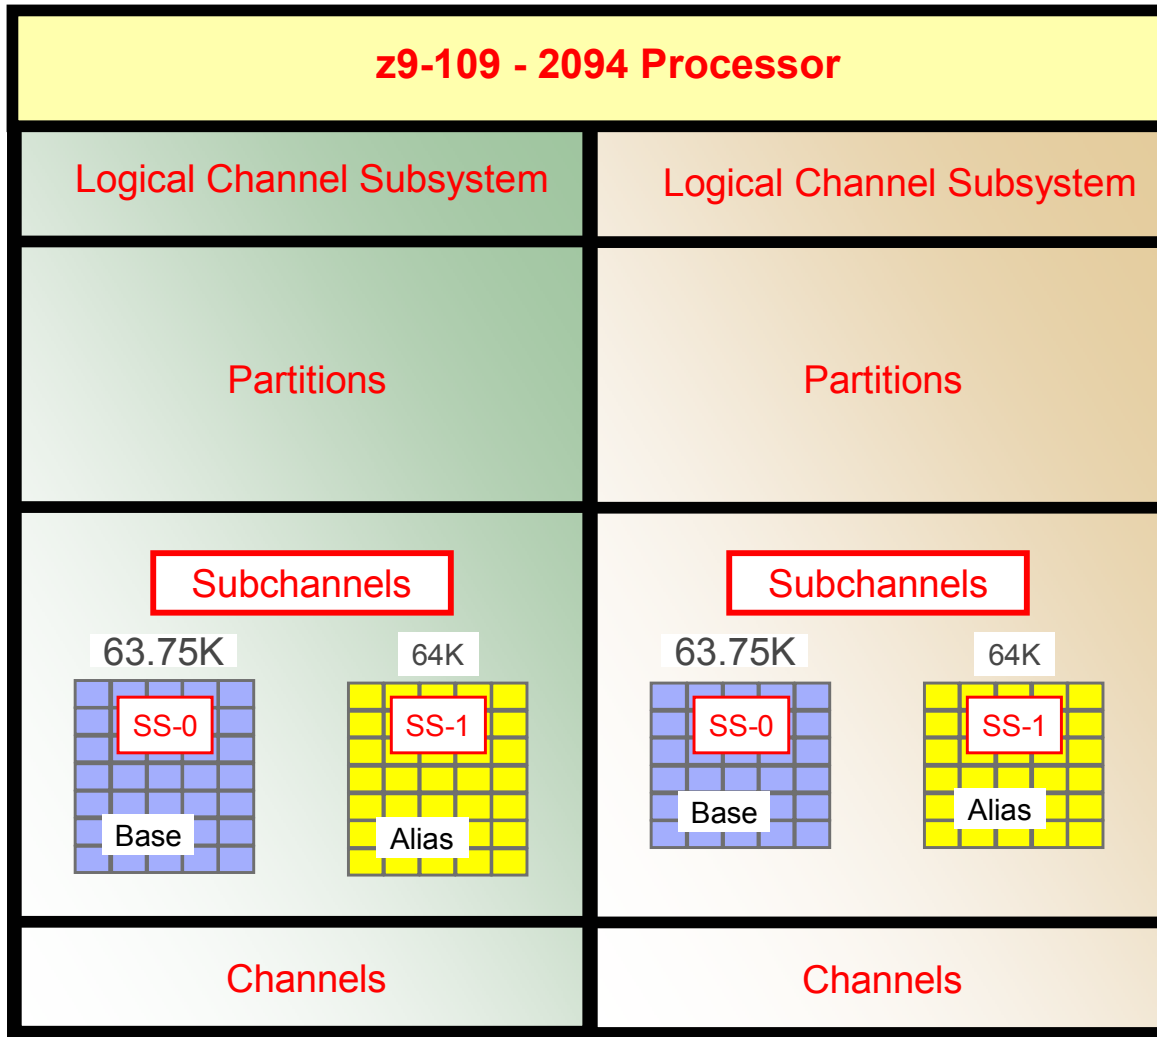
- New function overview
- **Dynamic handling of processor speed changes**
- **Multiple subchannel sets**
- **zSeries Application Assist Processors , WLM zAAP reporting support**
- Routing enhancements, server specific load balancing
- Contention Management Enhancements
- WLM Support for Enterprise Workload Management
- WLM-managed batch initiator enhancements
- Sub-capacity pricing support
- Service definition coefficients
- User Friendly Interface for WLM Administrative Application
- Summary

Dynamic Recognition of Speed Changes

- **Processor speeds are no longer constant**
 - E.g. z890 speed/model changes, On/Off Capacity on Demand
 - Speed changes only for general purpose processors
 - BTW: Small degradations also in case of z990 or z9-109 cooling problems
- **Hardware presents “Capability Change” SCLP event**
 - STSI information changes
 - ENF20 posted in z/OS
 - Significant speed changes call for WLM/SRM re-adjustment
- **WLM will update processor speed related constants and refresh its policy**
 - New message: IWM063I WLM POLICY WAS REFRESHED DUE TO A PROCESSOR SPEED CHANGE
 - RMF interval synchronized, SMF99 records written

z9-109 – Multiple Subchannel Sets per LCSS

Each 2094 Logical Channel Subsystem can have up to two sets of Subchannels



WLM Support for z9-109 Multiple Subchannel Sets

- **For z9-109, two subchannel sets are now available per LCSS,**
 - Enables a total of 63.75K subchannels in set-0 and 64K-1 subchannels in set-1.
 - Only Parallel Access Volume Alias (PAV-alias) devices are allowed in subchannel set-1.
- **Multiple subchannel sets provides growth for I/O device configuration**
- **WLM enables PAV alias management for alias' in subchannel set-1.**
- **z/OS 1.7 or later required**

zSeries Application Assist Processor (zAAP)

- **What is it?**
 - Powerful z/OS Java execution environment
 - Allows installations to offload Java work from general purpose CPs to zAAPs
- **Advantages**
 - Significantly less expensive than general purpose CPs
 - Possible reduction of demand on general purpose CPs
 - Upgrade of additional zAAP processing power does not affect total MSU rating or machine model designation
 - No IBM software charges on zAAP capacity



IBM eServer zSeries 990, 890,
System z9-109 or later

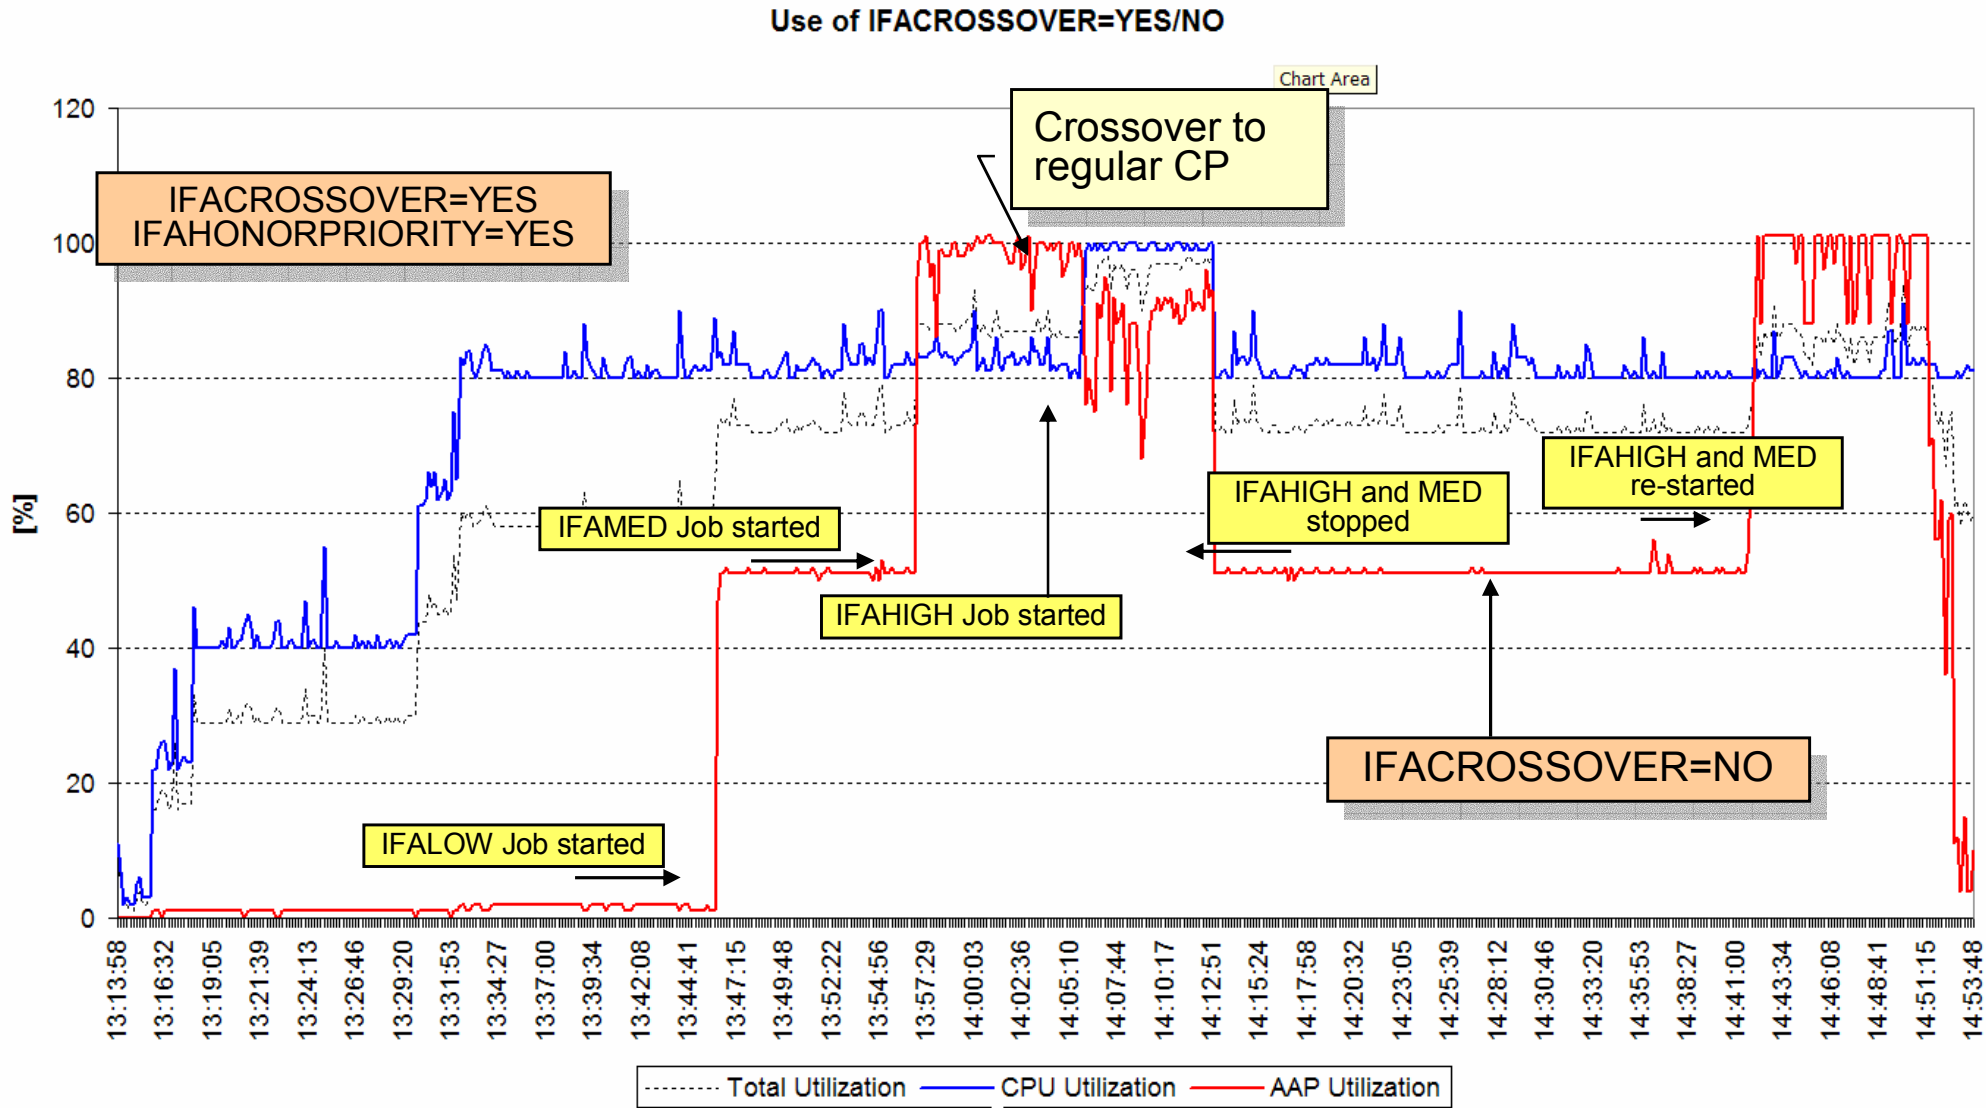
zAAP Exploitation

- **zSeries Application Assist Processors (zAAP) can only be deployed on IBM eServer zSeries 890 or 990 or later with z/OS V1.6 or later and the IBM SDK for z/OS Java 2 Technology Edition, V1.4 with PTF.**
- **The following subsystems provide zAAP execution potential:**
 - WAS 5.1
 - IMS V7, V8, V9
 - CICS 2.3
 - DB2 V7, V8
 - WBI Brokers V5
- **Additional information can be found on:**
 - zAAP home: <http://www.ibm.com/servers/eserver/zseries/zaap/>
 - FAQ: <http://www.ibm.com/servers/eserver/zseries/zaap/faq/>
 - Redbook: <http://www.redbooks.ibm.com/redbooks/pdfs/sg246386.pdf>
 - Whitepaper: <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10337>

Current Handling of zAAP-Related IEAOPTxx Parameters

IFACROSSOVER	IFAHONORPRIORITY	Meaning
<u>YES</u>	<u>YES</u>	<ul style="list-style-type: none"> ▪ CP's help zAAP's honoring priorities until discretionary work is reached. ▪ CP's always help zAAP discretionary work after all CP work. <ul style="list-style-type: none"> – With soft-capping applied: See next row
YES	NO	<ul style="list-style-type: none"> ▪ CPs always help zAAP's after all CP work
NO	YES	
NO	NO	<ul style="list-style-type: none"> ▪ CPs never help zAAPs, except when no zAAPs are operational

zAAPs: Crossover Example to Regular CPs



Planned Enhanced Handling of zAAP-Related IEAOPTxx Parameters

IFACROSSOVER	IFAHONORPRIORITY	Meaning
<u>YES</u>	<u>YES</u>	<ul style="list-style-type: none"> ▪ CPs help zAAP's when they become busy, honoring priorities until discretionary work is reached. ▪ CP's always help zAAP discretionary work after all CP work.
YES	NO	<ul style="list-style-type: none"> ▪ CPs always help zAAPs after all CP work.
NO	YES	<ul style="list-style-type: none"> ▪ Like YES/YES
NO	NO	<ul style="list-style-type: none"> ▪ CPs never help zAAPs , except when no zAAPs are operational

WLM Treatment of the zAAP Resource

- **A zSeries Application Assist Processor is a new resource type WLM is aware of**
 - Contributes using and delay samples
 - Contributes service times
 - zAAP utilization is reported by RMF (SPE OA05731)
- **zAAPs are managed by WLM as extension of CPs**
 - Java work executing on zAAPs inherits the dispatch priority from its execution on regular CPs
 - Execution is accounted for in execution velocity and goal achievement (PI)
- **zAAPs are not varied by IRD Vary CPU Management**
- **On z890 and z990: zAAPs inherit weight from CP pool but are part of the ICF/IFL pool**
- **On z9-109 zAAPs are in a separate pool.**

Agenda


- New function overview
- Multiple subchannel sets
- Dynamic handling of processor speed changes
- zSeries Application Assist Processors , WLM zAAP reporting support
- **Routing enhancements, server specific load balancing**
- **Contention Management Enhancements**
- WLM Support for Enterprise Workload Management
- WLM-managed batch initiator enhancements
- Sub-capacity pricing support
- Service definition coefficients
- User Friendly Interface for WLM Administrative Application
- Summary

Routing Services: Before z/OS 1.7

Step 1

– Registration of servers by IWMSRSRG services

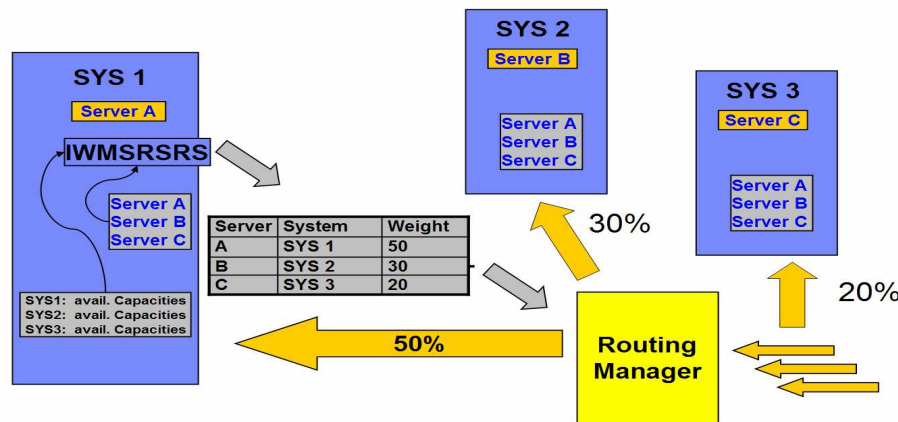
- Each server registers on the local system as a routing candidate
- WLM communicates registered servers and capacity data between systems



Step 2

– Router asks WLM for recommendations by using IWMSRSRS service

- WLM calculates a weight for each registered server based on capacity
 - Servers on systems with serious constraints are excluded (weight 0) unless all systems have serious constraints
- The returned weight reflects the available or replaceable system capacity



Sysplex Routing: Observations Before z/OS 1.7

- **Current recommendation is purely capacity based**
 - Weight reflects system capacity and is not server specific
 - Some observed situations:
 - If the LPAR is loaded with low important work, but the server has high importance
→ weight too low
 - If the LPAR is lowly utilized, but the server has a bad performance index or is waiting for resources
→ weight too high
- **TCP/IP stack registration does not reflect how work is processed on the system**
 - TCP/IP address space is usually running in a high important service class
 - The registered address space does not consume the work, it just passes the work to the real consumer on the system
- **Abnormal conditions of the work receiver or consumer are not factored into the routing recommendations**
 - This can lead to *storm drain* scenarios where a system seems to be eligible for processing many requests while in fact the work receiver or consumer is ill
 - In many cases this requires that the work receiver consumer provides additional information to WLM
- **LPAR Wait Time includes wall-clock time**
 - In addition soft capping and LPAR weights are not correctly reflected
 - This can result in an over estimation of the available capacity of a logical partition
 - As a result to many requests may be routed to certain systems (partitions) in a sysplex while other systems receive to few requests

Sysplex Routing: Enhancements

- **Current recommendation is purely capacity based**
 - New option on routing services will allow a router to receive weights which consider the goal achievement of the work
 - **z/OS 1.7 (IWMSRSRS FUNCTION=SPECIFIC)**
 - May be exploited by DB2 in the future
- **TCP/IP only registers the stack and doesn't know how the work is being processed**
 - Introduce new routing service e.g. for TCP/IP which allows to provide routing recommendations for TCP/IP backend applications
 - **z/OS 1.7 (IWM4SRSC)**
- **Abnormal conditions of the work receiver or consumer are not factored into the routing recommendations**
 - In a first step the queue time of the work is accounted for DDF regions and its proportion to the execution time is factored into the routing recommendations
 - **z/OS 1.7 (incorporated in IWMSRSRS FUNCTION=SPECIFIC)**
- Current activities to provide broader solutions
 - Include abnormal terminations which are recorded by CICS
 - New interface/option which allows a server to provide its state to WLM
- **Measured Waittime is put in perspective with real available system and partition capacity**
 - Over estimation is reduced
 - This part is also available as **OA10006 on z/OS 1.4 and above**

Routing Services: Summary

Functions in grey/non-bold font are considered for a future release. Plans may be subject to change without notification.

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. Considers: <ul style="list-style-type: none"> ▪ Goal Achievement (PI) ▪ Queue Time for Enclaves ▪ Health Indicator
IWMSRSRG	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. IWMSRSRG 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

Sysplex Routing: Enhancements for IWMSRSRS Service

- **New Function code „SPECIFIC“**
- **Weight calculation: Product of four factors**
 - System Utilization Factor:
 - Same as the resulting system weight for old „SELECT“ function (as described in the earlier foils)
 - PI Factor
 - This gives an indication of how good this server, respective the work that is related to this server, is achieving its goals as defined in the active WLM policy.
 - Queue Time Ratio:
 - If the server owns independent enclaves, the ratio of queue time to elapsed time of those enclaves
 - Health Indicator
 - A registered server can use the IWMSRSRG HEALTH= parameter to inform WLM about additional states which are unknown to WLM but should be factored into the returned weight
 - Remark
 - If multiple servers are registered on the same system, the weight is divided by the number of the servers.

Sysplex Routing: New IWM4SRSC Service

- **Characteristics**
 - Provides routing recommendation for a given server
 - System scope
 - No registration of servers necessary
 - The intention of this service is to obtain routing recommendations for work which is passed through the registered server (e.g. TCP/IP) to a server which consumes/processes the work
 - Input: STOKEN that identifies the address space of a server
 - Output: Recommendation of how good this server is suitable to receive work → the weight
 - A value between 1 and 64

- **Weight calculation: Product of two (four) factors, scaled by 64**
 - The PI Factor (Performance Indicator Factor)
 - This gives an indication of how good this server, respective the work that is related to this server, is achieving its goals as defined in the active WLM policy.
 - The Importance factor
 - This is a measurement of how much CPU Capacity is displaceable by work of the server's importance, respective the work that is related to this server
 - The health factor
 - This is an indicator an address space can set via the new IWM4HLTH service
 - Abnormal termination rate

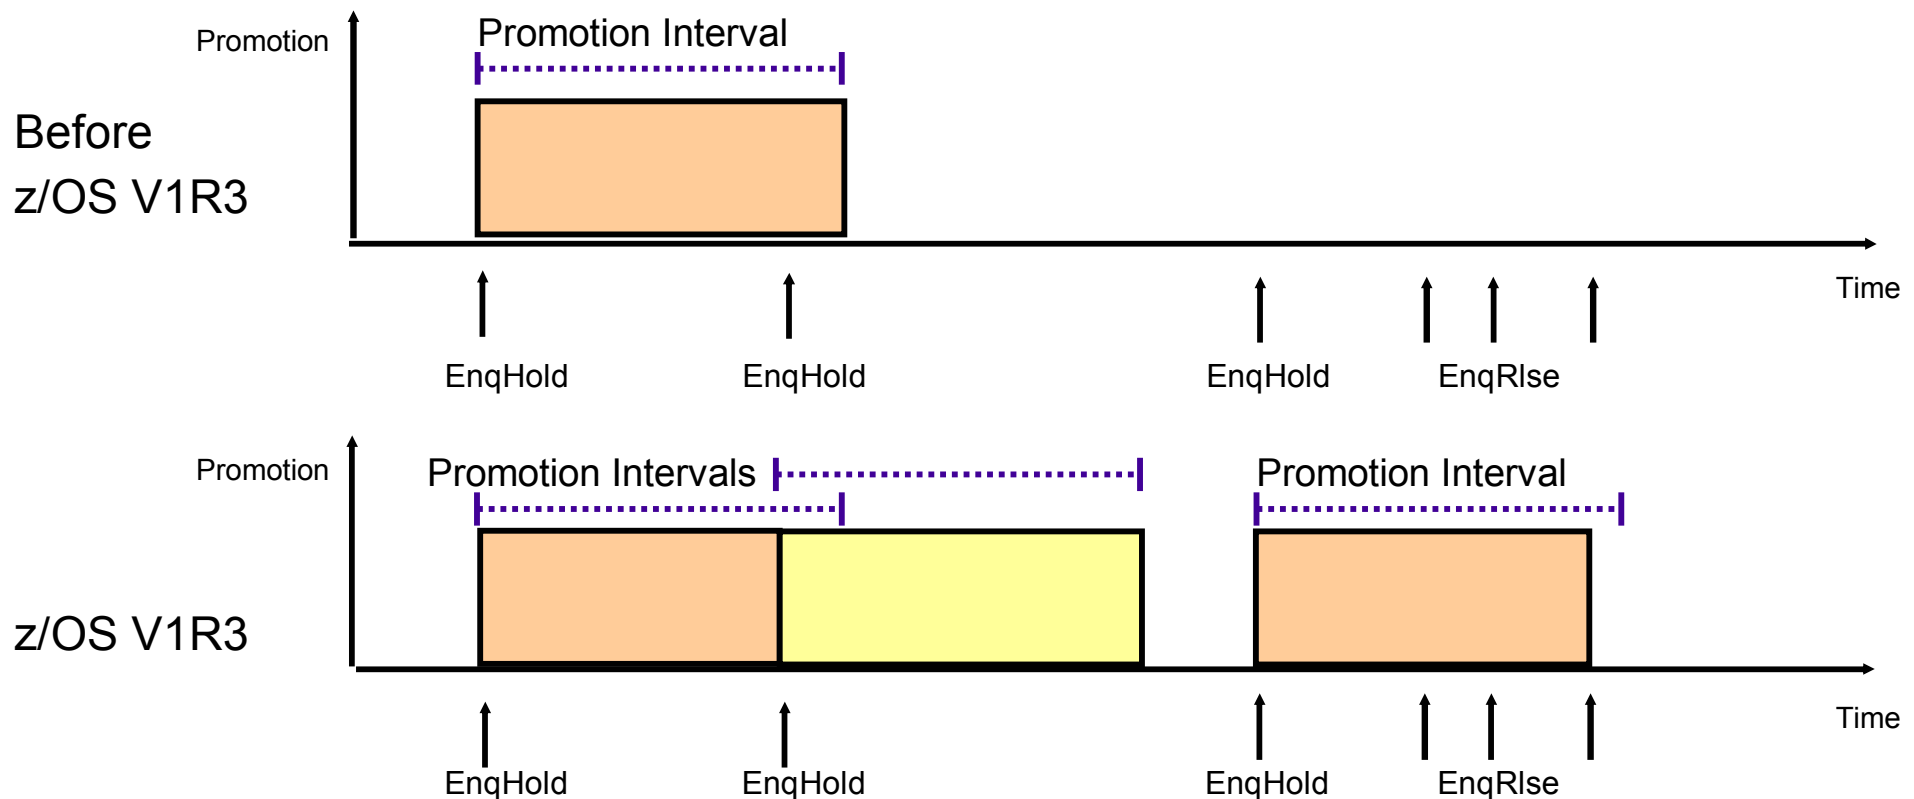
WLM Contention Management

Overview

- **WLM provides two interfaces to allow resource managers to signal contention situations which need special treatment**
 - ENQHOLD:
 - A unit of work waits for a resource which is in use by another unit of work
 - ENQRLSE:
 - The situation has been cleared
- **WLM/SRM**
 - Promotes holder (increases DP) for a specified duration to assist in resolving the contention situation
 - Period is specified through ERV parameter in IEAOPT member of SYS1.PARMLIB
- **With z/OS 1.3**
 - WLM introduced some significant enhancements to the interfaces
 - Serviceability reasons
 - To extend the promotion if the queue for the resource builds up
- **Further required enhancements**
 - Provide a short term promotion
 - which does not require to signal the release
 - Solved with DB2 Latch Contention Item for z/OS 1.7 and as APAR on z/OS 1.6 (OA08949)

Enhanced Enqueue Promotion (z/OS V1R3)

- Promotion interval (ERV in IEAOPTxx) restarts at every contention notification
- Quiesced address spaces are kept swapped-in as long as there is ready work to run



- With z/OS V1R3, enqueue management has been improved by a more sophisticated enqueue promotion algorithm.
- An address space or enclave is promoted in terms of dispatch priority when it holds a resource that another address space or enclave wants to have. The resource manager indicates this situation to SRM through an ENQHOLD-sysevent. By promoting the address space or enclave for a limited amount of time, the system hopes that the holder gives up the resource faster than it usually would with a lower DP. Also, while being promoted, the system ensures that the address space or address space associated with the enclave is not swapped out.
- When a contention disappears, the resource manager notifies SRM through an ENQRLSE-sysevent.
- The enqueue promotion interval can be set by the installation through the ERV-option in IEAOPTxx-parmlib member. The ERV-option specifies the CPU service units, that an address space or enclave can absorb while it is promoted before the promotion is taken back.
- In goal mode, the enqueue promotion dispatch priority is determined dynamically at every policy adjustment interval (10 sec). It can change based on available processing capacity and amount of high dispatch work in the system. Address spaces are promoted to that priority, if their current dispatch priority is not already higher than that.
- Before z/OS V1R3, an address space or enclave was only promoted once for a given resource. If the resource wasn't returned before the expiration of the promotion interval, it was just a nice try.
- With z/OS V1R3, an address space or enclave is promoted anew by SRM for every contention indication as told by the resource manager. Also, while there is an outstanding contention release notification, quiesced work is kept swapped-in as long as there is ready work to run. This includes address spaces associated with enclaves.

WLM Contention Management

Enhancements for DB2 Latch Contention: z/OS 1.7 and z/OS 1.6 and 1.5 (OA08949)

Observation

- Latch contention in DB2 can cause significant system degradations
- Latch contentions only require a short boost and no elongated promotion technique

WLM


- Introduces extension to ENQHOLD
 - Short Term Promotion
 - Does not require enqrise, expires automatically
 - Provides a short boost
 - Much smaller than regular enqhold/rise
 - Length of boost is fixed

DB2

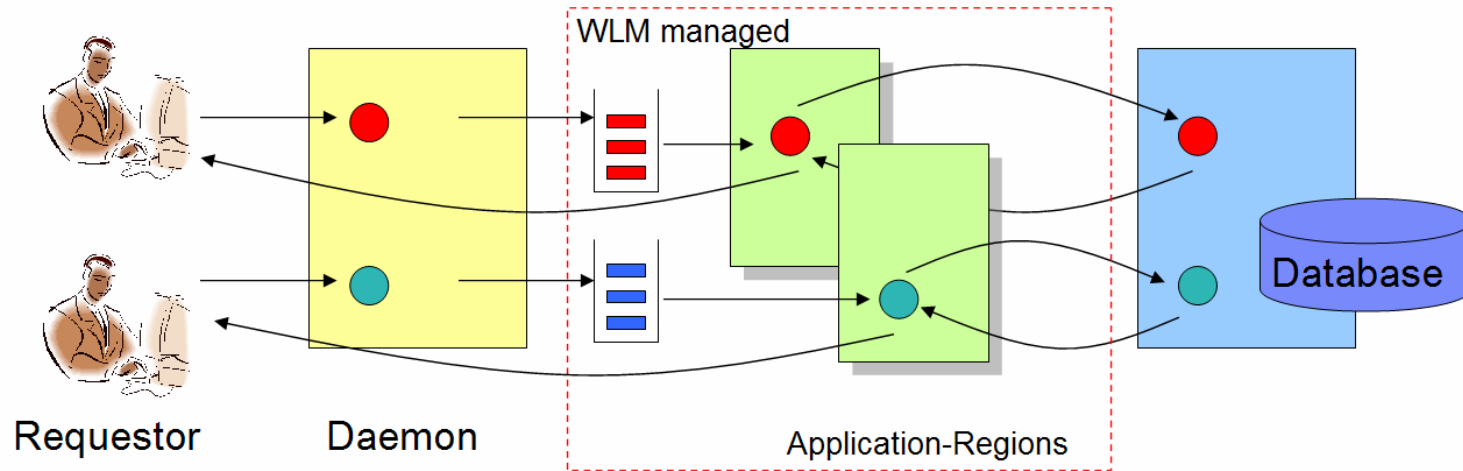
- Will be using new latch contention mechanism in some situations
- Issues Enqhold (short term promotion) to allow WLM to raise the dispatch priority for the holder and complete its work

GRS

- Exploits new ENGHOLD format with STOKEN



WLM Application Environments: Review



- **Named group of similar server programs which execute in a server address space (region)**
- **Transactions are processed in the server address space under an enclave task in their own service class**
- **WLM maintains and monitors transaction queues**
 - For each application environment there is one queue per transaction service class
 - A variable number of server address spaces is associated with each queue to process transactions
 - WLM starts / stops server address spaces based on goal fulfillment and demand

OA04555: Enhancements for WLM-Established Stored Procedures (z/OS V1R5)

■ Current Assumption

- All work requests (STPs) inserted by DB2 are independent of each other


■ Problem

- Assumption does not always hold true
- STPs can recursively call other STPs
- UDFs can insert multiple dependent STPs
- Processing of these sets is all or nothing

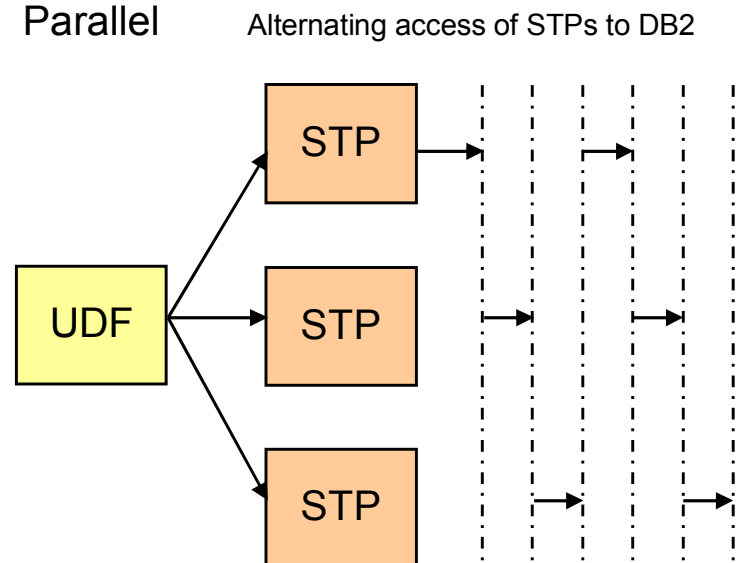
■ Solution

- DB2 tells WLM about dependent STPs
- WLM gives dependent requests priority over other requests
- WLM attempts to start server regions more aggressively if dependent requests have to wait as long as sufficient system resources are available

Recursive



Parallel

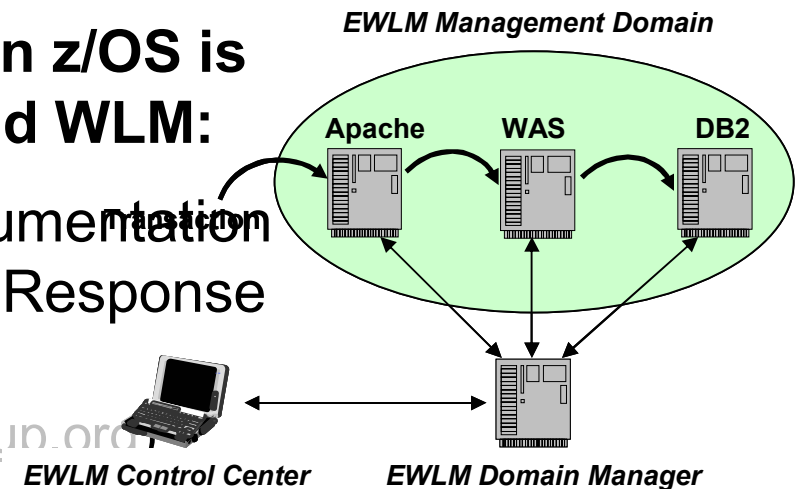


Agenda

- New function overview
- Multiple subchannel sets
- Dynamic handling of processor speed changes
- zSeries Application Assist Processors , WLM zAAP reporting support
- Routing enhancements, server specific load balancing
- Contention Management Enhancements
- **WLM Support for Enterprise Workload Management**
- WLM-managed batch initiator enhancements
- Sub-capacity pricing support
- Service definition coefficients
- User Friendly Interface for WLM Administrative Application
- Summary

OA07196 – Support for Enterprise Workload Management (EWLM)

- **EWLM provides end-to-end transaction reporting and will, in the future, provide end-to-end transaction management**
- **Enterprise workload management on z/OS is a co-operative function of EWLM and WLM:**
 - WLM provides new application instrumentation services adhering to the Application Response Measurement (ARM 4.0) standard from The Open Group (www.opengroup.org)
 - WLM implements the EWLM extensions of the WLM Enclave services to allow applications using these services to participate in EWLM.



Agenda

- New function overview
- Multiple subchannel sets
- Dynamic handling of processor speed changes
- zSeries Application Assist Processors , WLM zAAP reporting support
- Routing enhancements, server specific load balancing
- Contention Management Enhancements
- WLM Support for Enterprise Workload Management
- **WLM-managed batch initiator enhancements**
- **Sub-capacity pricing support**
- **Service definition coefficients**
- **User Friendly Interface for WLM Administrative Application**
- Summary

OA10814

- **Prior to OA10814:**
New initiators started only, if

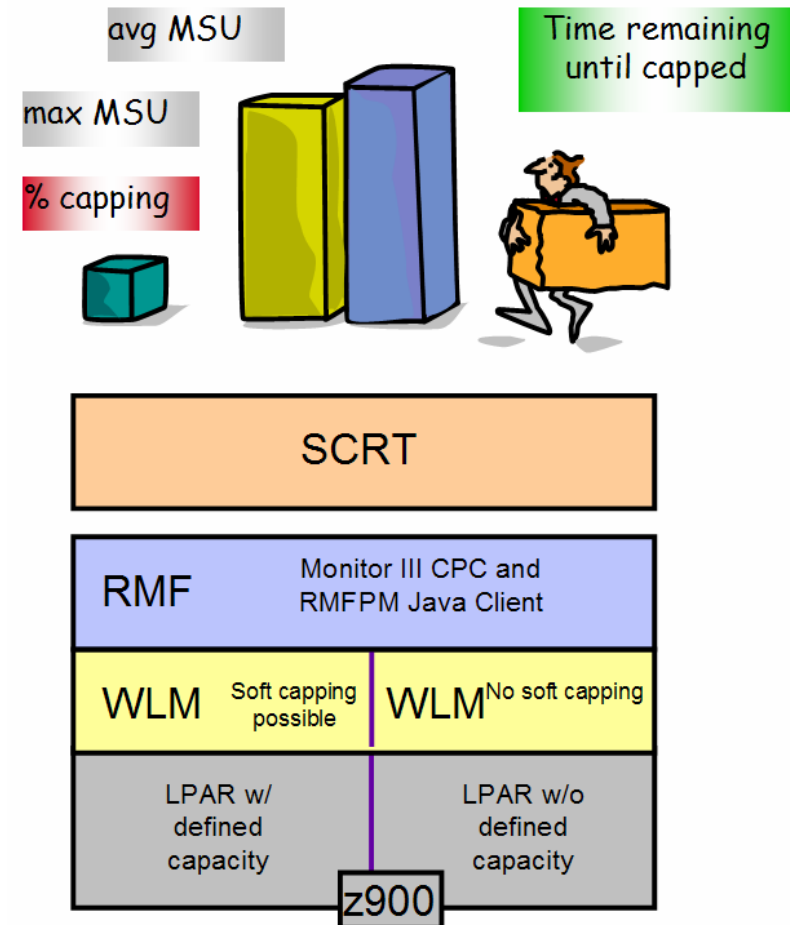
$$n_{newInits} + \sum_{allSystems} n_{oldInits} \leq 1.5 \cdot \sum_{allSystems} (J_{activeRequests} + J_{queuedRequests})$$

- **But initiators could be idle due to system affinity or scheduling environment**
- **With OA10814:**
 - If “remote” system has idle initiators but zero eligible requests this indicates that queued requests are not eligible on that system
 - Continue assessment in that case, i.e. allow for initiators to be added on the local system

SYS	Active Inits	Total Inits	Idle Inits	Queued elig. Requests
SYS1	1	1	0	1
SYS2	0	1	1	0
SYS3	0	1	1	0
SYS4	0	0	0	0
SYS5	3	5	2	0
SYS6	0	1	1	0
SYS7	0	1	1	0
SYS8	1	1	0	0
SYS9	1	1	0	0
SYSA	0	1	1	0
SYSB	0	1	1	0
SYSC	0	1	1	0
	Σ=6	Σ=15	Σ=9	

Review: OW55509 Sub-Capacity Pricing Before z/OS R7

- **Workload license charge (WLC)**
 - New pricing model applicable with zSeries servers running z/OS in 64-bit mode in an LPAR
 - Defined capacity can be set to enable WLM soft capping in order to control costs
- **WLM support with OW55509**
 - Supports changes in WLC-pricing and the Sub Capacity Reporting Tool SCRT V3
 - Calculates 4h rolling average capacity regardless of defined capacity specification
 - Uses 48 SUs as 4h rolling average when system IPLs
 - The 4h rolling average has a lower limit of 1MSU if calculated
 - Turns soft capping on when
 - Defined capacity specified AND
 - 4h rolling average > defined capacity
- **RMF support**
 - Monitor III CPC Capacity report
 - Additional metrics in RMF PM
- **Prior to z/OS R7:**
 - Some restrictions apply. See next charts.



Sub-Capacity Pricing Enhancements for z/OS R7

■ Prior to z/OS 1.7

- WLM provides the 4-hour rolling average for systems running in LPAR mode
- No support for systems running as guests under z/VM
- 4 hour rolling average is calculated using the total dispatch time of all logical CPUs of the running LPAR
 - I.e., dispatch times for dedicated processors and LPARS using Waitcompletion = YES includes total wait time for all logical CPUs
 - Support for dedicated and Waitcompletion= YES configurations in SCRT tool

■ On z/OS 1.7 and above:

- Support for systems running as guests under z/VM
 - All z/OS guests need to be at z/OS 1.7 level or above
 - SCRT 11.1 required
- 4 hour rolling average is calculated using the actual CPU time consumed, i.e. without wait time
 - Support for dedicated and Wait completion= YES configurations now in z/OS

OA09081/OA09118 (z/OS R6): Provide Service Definitions Coefficients

- **OA09081 (SRM): Provide Service Definitions Coefficients (MSO, CPU, IOC, SRB) to SMF**
 - SYSEVENT REQSVDAT
 - **OA09118 (SMF): Externalize coefficients in SMF30**
 - Macro IFASMFR3:
 - SMF30MSC: fullword for the MSO SDC, scaled by 10000
 - SMF30CPC: halfword for the CPU SDC, scaled by 10
 - SMF30LOC: halfword for the IOC SDC, scaled by 10
 - SMF30SRC: halfword for the SRB SDC, scaled by 10
- Can be used to calculate the zAAP normalization factor
- **OA10901 adds zAAP Normalization factor to SMF30:**
 - SMF30ZNF: halfword for zAAP normalization factor, scaled by 256

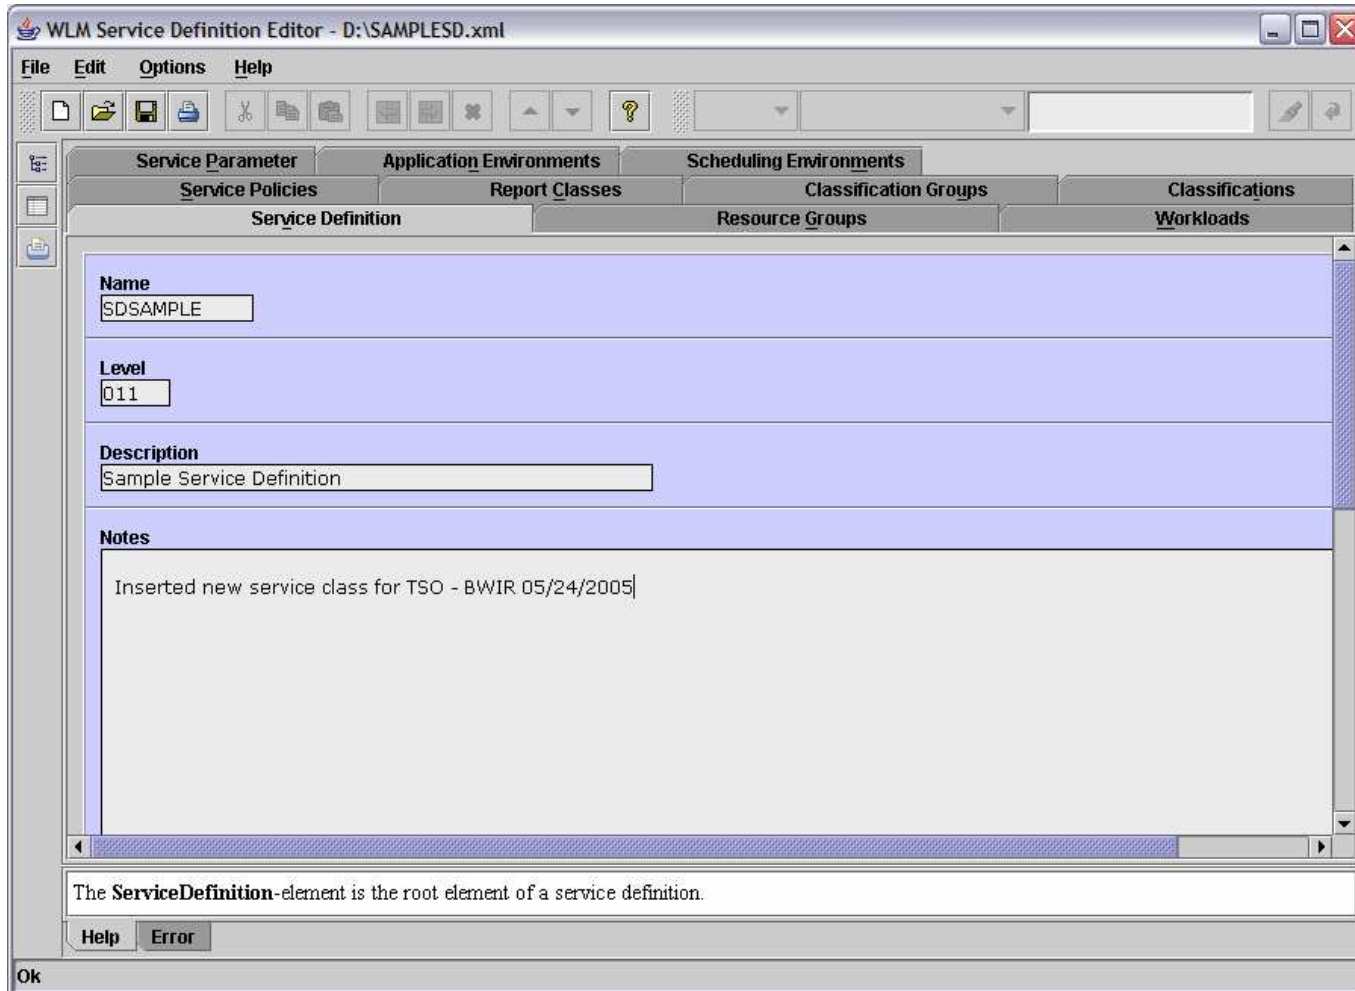
OA10641: How to use MSO Service Definition Coefficient for Calculations

- **MSO Service Definition Coefficient is defined as part of the service definition (SVPOLMSO)**
 - Externalized via
 - IWMDEXTR WLM service
 - SMF Record 72 Subtype 3: R723MMSO (for years)
 - SMF 30 Performance Section: SMF30MSC
 - Value Range: 0.0001 to 99.999
 - The externalized values are scaled by 10000, so 0.0001 equals 1
 - Range: 1 to 999999
- **What to do if you want to use this value in calculations?**
 - SRM uses a rounded version of defined and externalized value
 - The basic idea of the rounding was that 1 MSO SU equals to keep **50 pages for 1 second** on a base machine. This was changed in 1988 to provide more precision for calculation purposes
 - *Unfortunately this extension is not precise:*
 - *The input value is scaled like:*

$$WLM _ InternalMSO _ Coeff = \frac{DefinedMSO _ CoeffScale \ dBy10000}{10000} \cdot \frac{4096}{50} + 1$$

- *As a result all values between 0.0001 and 0.0122 are the same and represent the minimum value that WLM uses for calculating MSO Service Units.*
All internal MSO Coefficients are multiples of 0.0122
- **If not already using a zero MSO coefficient reconsider whether that might be more appropriate...**

User Friendly Interface for WLM Administrative Application



- **Java application**
- **Interface to WLM ISPF datasets**
- **Different views to show relations between policy elements**
- **Real-time error checking**
- **Real-time help**
- **No popup menus**
- **Tabular input processing**

User Friendly Interface for WLM Administrative Application

Example: Workload and Service Class Panel/View

WLM Service Definition Editor - D:\SAMPLESD.xml

File Edit Options Help

local

Service Definition		Resource Groups		Workloads		Service Policies		Report Classes	
Name	ServiceClasses	Period	Goal	Im...	Duration	ResponseTime	Perce...	Level	Description
WKLDASC									ALL APPC Transaction
WKLDASC	A3V30STD							GBATCH20	ASCH default Service Class
WKLDASC	A3V30STD	1	Velocity	2	500			10	
WKLDASC	A3V30STD	2	Velocity	2	-			8	
WKLTK									All Batch Jobs
WKLTK	B4V10STD							-	Batch Standard VEL 10 IMP 4
WKLTK	B4V10STD		Velocity	4	-			10	
WKLTK	B4V20STD							-	Batch Standard VEL 20 IMP 4
WKLTK	B4V20STD	1	Velocity	4	1000			20	
WKLTK	B4V20STD	2	Velocity					10	
WKLDTSO									ALL TSO USERIDS
WKLDTSO	T2335DEV							-	Developer (Standard) TSO
WKLDTSO	T23		PrecentileResponseTime	2	2500	00:00:02.000	98		
WKLDTSO	T23		AverageResponseTime	3	300000	00:00:20.000	95		
WKLDTSO	T23		PrecentileResponseTime	5	-			10	
WKLDTSO	T23							-	No Production TSO Helpers
WKLDTSO	T23		PrecentileResponseTime	2	2000	00:00:01.000	99		
WKLDTSO	T23		PrecentileResponseTime	2	10000	00:00:00.000	90		

used by:
Classification: JES

Insert
Insert Before
Insert After
Replace by
Copy
Cut
Delete

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

Help Error

Ok

- Information is represented as tables
- Direct manipulation of tables
- Display of relations between policy elements
- Real-time error checking
- Context-sensitive help

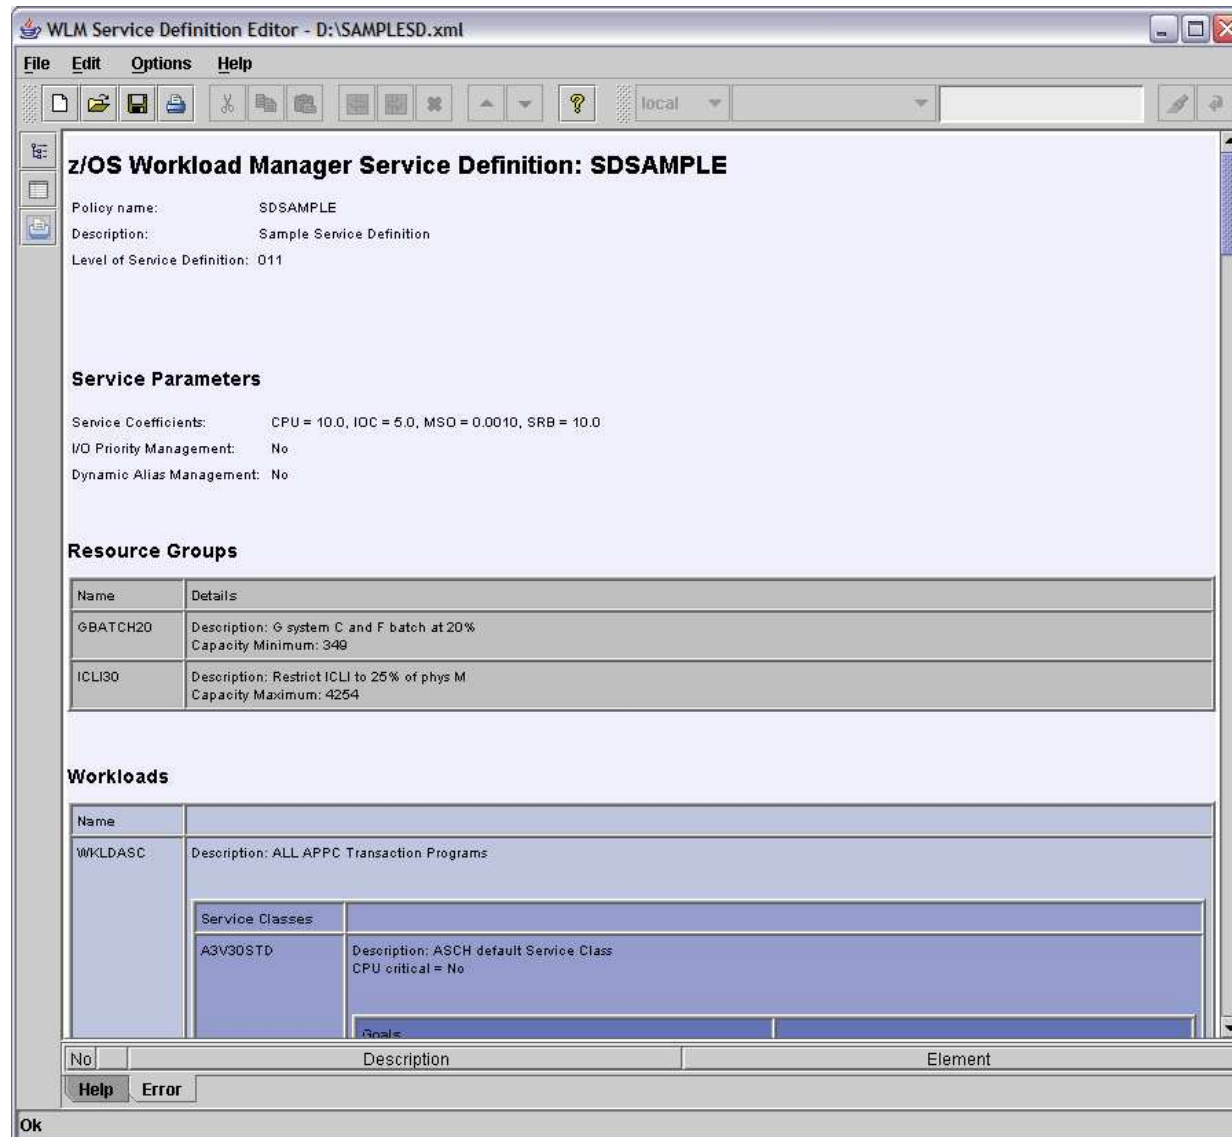
User Friendly Interface for WLM Administrative Application

Example: Easy generation of new views

The image displays two screenshots of the WLM Service Definition Editor interface, demonstrating table manipulation capabilities. The top screenshot shows a table with columns: Name, ServiceClasses, Period, Goal, Im..., Duration, ResponseTime, Perc..., Level, ResourceGro..., CPU..., and Description. The bottom screenshot shows a rearranged table with columns: Goal, ServiceClass..., P..., Im..., Level, Duration, Name, ResponseTime, Perc..., ResourceGr..., CPU..., and Description. Red arrows indicate the movement of columns and rows between the two views.

- Table columns can be moved as needed
- Table rows can be sorted as needed
- Rearranged tables can be edited
- Search function for arbitrary text within tables

WLM Administrative Application Example: Print View

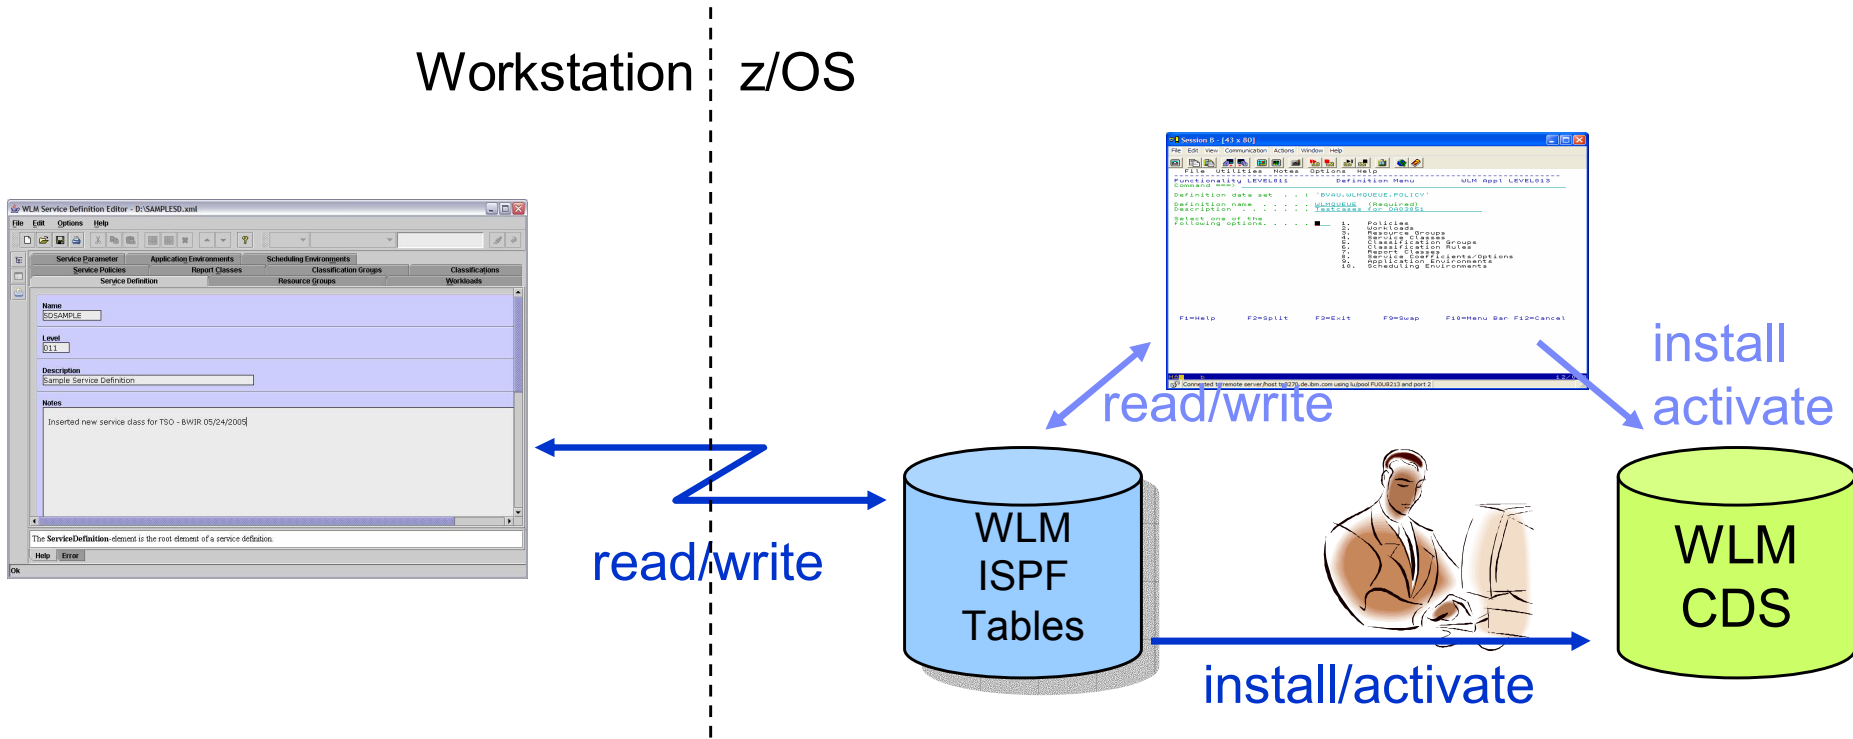


- Service Definition is converted to HTML-Document
- In order to print HTML-Document the configured WWW-Browser of the Operating System is launched

WLM Administrative Application

Link to existing WLM Administrative Application

Workstation z/OS



- **New workstation front-end to WLM ISPF Tables**
 - FTP connection with automatic upload and download
 - FTP connection profiles
- **Policy install and activate via operator command**

Delivery

- As a tool that can be downloaded via the WLM homepage

New 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>

Summary

- **Hardware Support**
 - zSeries z9 Support
 - zAAP Support
- **z/OS Themes**
- **Middleware support**
- **Improved Systems Management Capabilities**
- **Outlook**

Documentation

Documentation



z/OS MVS Planning: Workload Management (SA22-7602)

<http://publibz.boulder.ibm.com/epubs/pdf/iea2w151.pdf>

z/OS MVS Programming: Workload Manager Services (SA22-7619)

<http://publibz.boulder.ibm.com/epubs/pdf/iea2w252.pdf>

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

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

Redbook – z/OS Intelligent Resource Director (SG24-5952)

Redbook – Effective zSeries Performance Monitoring Using RMF (SG24-6645)

In German

Das Betriebssystem z/OS und die zSeries – Die Darstellung eines modernen Großrechnersystems, M.Teuffel/R.Vaupel, ISBN 3-486-27528-3

Internet Links



WLM

<http://www.ibm.com/s390/wlm>

SRM

<http://www.ibm.com/s390/srm>

RMF

<http://www.ibm.com/s390/rmf>

WSC

<http://www.ibm.com/support/techdocs>

Software Pricing

<http://www-1.ibm.com/servers/eserver/zseries/swprice/>

IRD

<http://www-1.ibm.com/servers/eserver/zseries/ird/>

Redbooks

<http://www.redbooks.ibm.com/>