

ibm.com



Powerful and secure infrastructures with WebSphere Application Server for z/OS

WAS 5 WLM considerations



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2003. All rights reserved.

Notices

This information was developed for products and services offered in the U.S.A.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of 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.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

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

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Unit Objectives

- Repetition: WLM Basics
- Understand the relationship between WebSphere and WLM
 - Application Environments
 - Workload Classification
 - Classification inheritance
- Learn about the performance problems that can occur if WLM is not tuned correctly for WebSphere
- Identify the configuration changes that should be in place for WLM to provide the desired levels of service



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Overview

Workload Manager (WLM) Overview

- Component of z/OS that allocates system resources
- Operates based on rules you define
 - You classify similar business applications into workloads
 - You define importance (priority) goals for each workload
- WLM allocates resources based on priority goals

WLM Service Classes

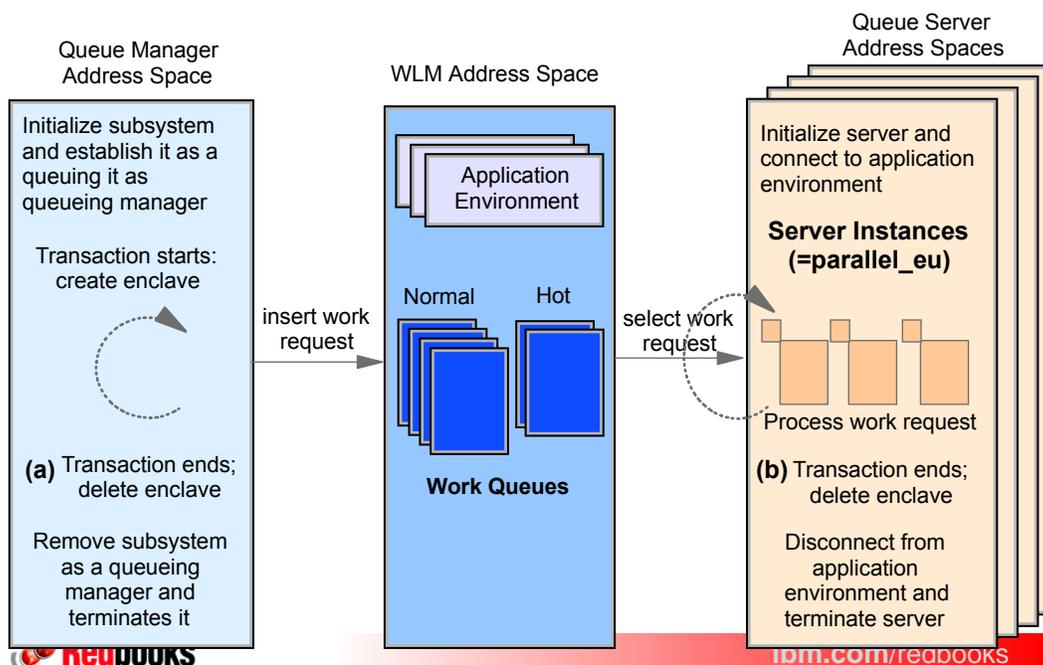
- Classification rules assign a workload to a Service Class
- Service Class priority can be assigned as:
 - Percentage Response Time (transaction-based work)
 - Velocity (non-transaction-based work)
 - Discretionary (low priority)
- Supports multiple performance periods with unique priorities



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM Application Environment Overview



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM Application Environments

- A feature of WLM used by WAS and other subsystems
- Composed of a Queuing Manager and Server Managers
- Queuing Manager
 - Usually only one
 - Receives work and passes it to WLM
 - WLM passes work on to a Queue Server
- Queue Server (servant)
 - Usually multiple; running in separate address spaces
 - Accepts work from Queuing Manager and WLM
 - WLM will start the first one when work arrives
 - WLM will start/stop others to meet Service Class goals
 - May be started across the Sysplex
- Implemented using Enclaves
- Two types of application environments:
 - static
 - dynamic



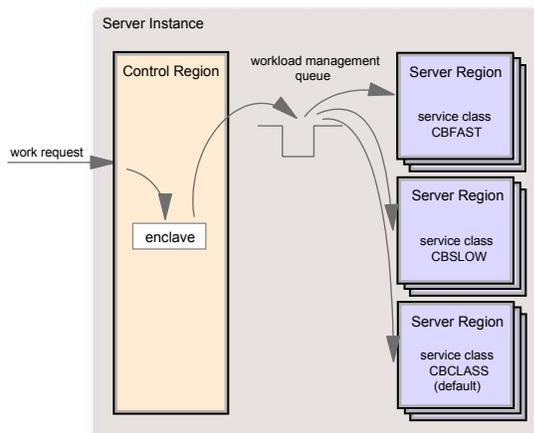
ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM Application Environments

- ▶ The WLM application environment definition tells WLM how it can start servant address spaces in order to manage the number of servants for the environment. The administrator defines a name, the subsystem type of the application, the name of the start procedure for the server address spaces and optionally parameters which are passed to the application server.
- ▶ WLM manages the Servants within the environment. (i.e., starts new Servants when the workload gets heavy; distributes work evenly across available Servants; etc.)

Queue Manager Queue Server



Application Environment Definition consists of a Name, Subsystem Type, a JCL Start procedure, and any parms for the Start procedure.



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Static Application Environments

■ Defining static ApplEnvs for WAS:

- Defined in WebSphere and WLM
- Application Environment name = WAS Server name
- The subsystem type will always be CB
- Must specify the PROC name to be started by WLM
- Set number of servers to 'No limit'
 - May want to specify '1' just during new application testing
 - Required if you need more than one server region per system
 - Required if you have multiple Service Classes for that J2EE Server defined

Service Class Recommendation

- Within the type CB classification rules...
- Define a default Service Class with high priority
 - Avoids problems during classification
 - Otherwise, default is SYSOTHER (low priority)

```

Application Environment Name . . : HAOASR2
Description . . . . . WAS HA J2EE Application Server
Subsystem Type . . . . . CB
Procedure Name . . . . . HAOASR2S
Start Parameters . . . . . IWMSSNM=&IWMSSNM

Limit on starting server address spaces for a subsystem instance:
1 1. No limit
   2. Single address space per system
   3. Single address space per sysplex
    
```



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Dynamic Application Environments (DAE)

- Provided with fix for APAR OW54622
- Definitions made only in WebSphere
- Changes SERVER class profiles from

- ▶ <subsys_type>.<subsys_name>.<appl_envir_name> to ..
- ▶ <subsys_type>.<subsys_name>.<appl_envir_name>.<cell_name>

■ WAS Environment needs additional variables, defined automatically

- ▶ server_region_dynapplenv_jclproc
- ▶ server_region_dynapplenv_jclparms
- ▶ daemon_group_name (== cell_short_name)
 - ♦ BBOM0001I server_region_dynapplenv_jclparms: JOBNAME=&IWMSSNM.S,ENV=CELDF.NODDF.&IWMSSNM
 - ♦ BBOM0001I server_region_dynapplenv_jclproc: WAS5DSF .

■ does it work?

- ▶ server_generic_sh
- ▶ MVS Commands:

```

DISPLAY WLM,DYNAPPL=CLUDF
IWM029I 17.49.41 WLM DISPLAY 457
DYNAMIC APPL. ENVIRON. NAME STATE STATE DATA
CLUDF AVAILABLE
ATTRIBUTES: PROC=WAS5DSF SUBSYSTEM TYPE: CB
SUBSYSTEM NAME: WASDF NODENAME: CELDF
    
```



ibm.com/redbooks

- ▶ check log

© Copyright IBM Corp. 2003. All rights reserved.

static vs dynamic

Characteristic	Static AE definitions	Dynamic AE definitions
Definition	by the user through the WLM administrative application. available on the system of a sysplex after policy activation.	by the queue manager address space of the application.
Scope	all systems of a sysplex	is the queue manager which defined it subsystem_type, node_name, subsystem_name
Stored	WLM Couple Dataset after Policy Install WLM ISPF dataset of the end user	(responsibility of the application, not saved by WLM.)
Existence	until the user deletes it from the WLM service definition by activating a service definition without the definition	as long as the queue manager instance is up and running.



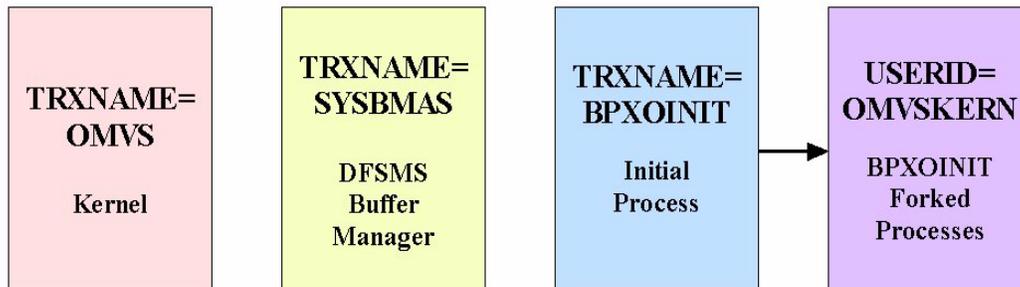
ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM definitions for z/OS UNIX

Assign a High-Priority Started Task (STC) Service Class

- Kernel (TRXNAME=OMVS)
- DFSMS Buffer Manager (TRXNAME=SYSBMAS)
- Initial process (TRXNAME=BPXOINIT)
- BPXOINIT forked processes (USERID=OMVSKERN)



All of these should be assigned high-priority, velocity-based Service Classes



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM Tuning Considerations

Service Class Assignments for WAS Controllers

- Use a high-velocity Service Class (such as SYSSTC)
- Assign Service Class to all WebSphere related address spaces as
 - Run-time servers
 - Daemon
 - Deployment MGR
 - JMS...
 - Application servers
 - J2EE Server Controllers



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Service Class Assignments for WAS Servants

- Use a high-velocity Service Class for startup
- Classify work into Service Class based on
 - Application Server name (CN) – This is the “server_generic_short_name” or “Cluster Transition Name” (which is also the application environment name.)
 - The Server instance name (SI) – This is the “server_specific_short_name”. This is usually not very useful because you cannot control which server instance runs a transaction within a cluster.
 - User ID assigned to the transaction (UI)
 - Transaction class (TC) - This can be assigned to a transaction using the transaction class mapping file for the server.
- Transaction Service Class should be response-time based
- Work running under server regions
 - Subsystem type=CB
 - Percentage response time goal is recommended
 - Ex. 80% of transactions < 0.5 seconds
- SUBSYS IWEB is gone



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Default Service Class Assignments for WAS Servants

- ▶ Server Startup
- ▶ Java garbage collection runs under this classification. Java GC is a CPU and storage intensive process, so if you set the velocity goal too high GC could consume more of the system resources than desired. On the other hand, if your Java heap is correctly tuned, GC for each server region should run no more than 5% of the time. Also, providing proper priority to GC processing is necessary since other work in the server region is stopped during much of the time GC is running.
- ▶ JSP compiles run under this classification. If your system is configured to do these compiles at run-time, setting the velocity goal too low could result in longer delays waiting for JSP compiles to complete.

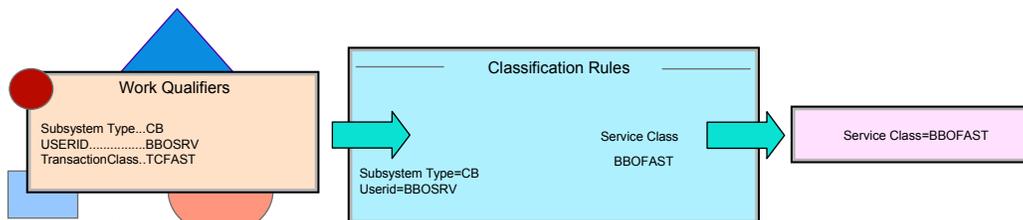


ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Workload Classification

- CB (IWEB is gone) subsystem classification rules (option 6 WLM panels)
- WLM CB rules use three ways to classify
 - by generic server name (J2EE server)
 - by USERID
 - Transaction class (protocol(s)_http_transaction_class; HTTP Transport Handler)
- Enclave
 - A "business transaction" without address space boundaries



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM Address Space Management & Replication

- WLM manages performance at Address Space level
- SRs are assigned to "internal queues" based on the SC
 - For each internal queue, WLM will assign a number of SRs equivalent to MIN_SRS (ex wlm_minimumSRCount=2)
 - Each of these SRs will have the same SC
- wlm_maximumSRCount=2 may interfere with WLM management
 - MAX has to be at least MIN*#SCs
 - #SCs - number of expected Service Classes
 - If MAX less than MIN*#SCs, situation illustrated in example 2 may occur
 - Consider having MAX=multiple of (MIN_SRS*#SCs) to give WLM more room to work



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

GUI: Setting Min/Max for WAS5

[Application Servers](#) > [wdfndcfdsc59](#) >

Server Instance

Configuration settings for servers which may dynamically have more than one servant process (such as on z/OS).

Configuration		
General Properties		
Multiple Instances Enabled	<input checked="" type="checkbox"/>	Specifies whether multiple server instance feature is enabled for this server. This field is ignored on platforms that do not include multiple server instance support. An example of where this is typically enabled is z/OS
Minimum Number of Instances	<input type="text" value="1"/>	The minimum number of server instances that may be created dynamically.
Maximum Number of Instances	<input type="text" value="3"/>	The maximum number of server instances that may be created dynamically.
<input type="button" value="Apply"/> <input type="button" value="OK"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/>		

The Maximum should be at least as large as the number of different service classes that might be used by transactions run in the server. Don't forget to account for the "default CB-type service class" and enclaves that may originate outside WebSphere servers and are classified by other classification rules such as the "HS"



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

Transport Handler Classification

Application Servers > wdfndfcdsc59 > Web Container >

Advanced Settings

Advanced settings for the Web Container.

Configuration

General Properties

Network QoS: NONE

Transaction Class Mapping: */u/wunder/tclass.conf

Apply OK Reset Cancel

Note: If you enable TC mapping via this file WAS will disable your default TC mapping.

Transaction Class Mapping file, EBCDIC

protocol_http_transport_class_mapping_file=u/wunder/tclass.conf

protocol_https_transport_class_mapping_file=u/wunder/tclass.conf

Syntax

- TransClassMap <host>:<port> <uritemplate> <tclass>

TransClassMap *.ibm.com:*	/SuperSnoop*	TCLASS
TransClassMap *:80*	/SuperSnoop*	TCLASS
TransClassMap *:*	/SuperSnoop*	TCLASS
TransClassMap *:80	/SuperSnoop*	TCLASS
TransClassMap *	*	BBTESTU2



defaults: protocol_http(s) transactionClass=

oks

© Copyright IBM Corp. 2003. All rights reserved.

Transaction Class - > Service Class

TransClassMap * * BBTESTU1

Subsystem Type CB - WAS V4 Series classification

Created by user ANDRE on 2001/10/02 at 14:54:45
Last updated by user ANDRE on 2002/11/18 at 01:54:03

Classification:

Default service class is WASDFLTH
There is no default report class.

#	Qualifier type	Qualifier name	Starting position	Service Class	Report Class
1	CN	BBPRD		WASPRODH	
1	CN	BBTEST		WASTESTH	OZ85BBT
1	TC	BBTESTU1		WASTESTH	BBTESTU1
1	TC	BBTESTU2		WASTESTH	BBTESTU2

Results reported in RMF monitor I workload activity report

- Transactions per second (not always the same as client tran rate)
- Response times (Goal achieved ?)
- How much of the machine you are using



OVERVIEW(REPORT)
SYSRPTS(WLMGL(SCPER))
SYSRPTS(WLMGL(RCPER))

© Copyright IBM Corp. 2003. All rights reserved.

RMF WLM Report

REPORT BY: POLICY=WEB_RAL

REPORT CLASS=BBTESTU2
HETEROGENEOUS

TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT	--DASD I/O--	---SERVICE---	--SERVICE RATES--	PAGE-IN RATES	---STORAGE---
AVG 0.00	ACTUAL	39.495	SSCHRT 0.0	IOC 0	ABSRPTN 50618	SINGLE 0.0	AVG 0.00
MPL 0.00	EXECUTION	210	RESP 10.7	CPU 21310	TRX SERV 50618	BLOCK 0.0	TOTAL 0.00
ENDED 2	QUEUED	39.284	CONN 7.1	MSO 0	TCB 0.2	SHARED 0.0	CENTRAL 0.00
END/S 0.00	R/S AFFINITY	0	DISC 3.3	SRB 0	SRB 0.0	HSP 0.0	EXPAND 0.00
#SWAPS 0	INELIGIBLE	0	Q+PEND 0.3	TOT 21310	RCT 0.0	HSP MISS 0.0	
EXCTD 0	CONVERSION	0	IOSQ 0.0	/SEC 24	IIT 0.0	EXP SNGL 0.0	SHARED 0.00
AVG ENC 0.00	STD DEV	54.482			HST 0.0	EXP BLK 0.0	
REM ENC 0.00					APPL % 0.0	EXP SHR 0.0	
MS ENC 0.00							

SDSF Enclave Report

```
SDSF ENCLAVE DISPLAY SC59 ALL LINE 1-1 (1)
COMMAND INPUT ==> - SCROLL ==
NP TOKEN SStype Status SrvClass Per PGN RptClass ResGroup
2800000D1D CB ACTIVE WASTESTH 1 BBTESTU2
```

Enclave 2400000D3B on System SC59

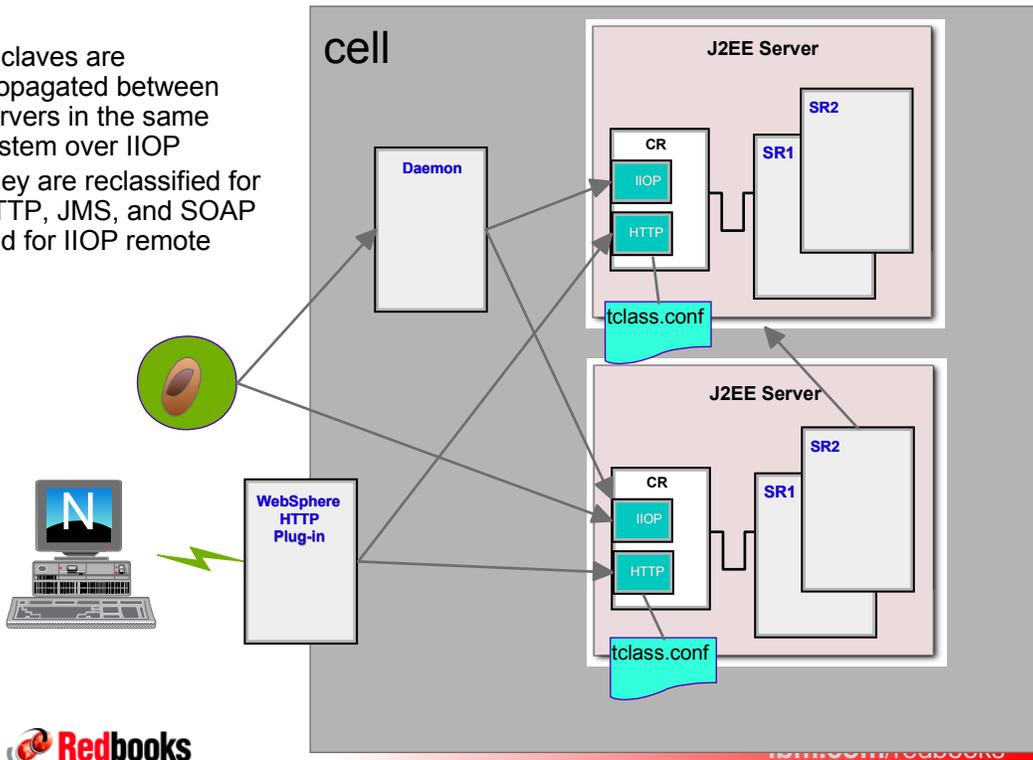
```
Subsystem type CB Plan name
Subsystem name WASDF Package name
Priority Connection type
Userid WDCFSTU Collection name CLUDF
Transaction name Correlation
Transaction class BBTESTU2 Procedure name
Netid Function name BOSS
Logical unit name Performance group
Subsys collection Scheduling env
Process name
```



© Copyright IBM Corp.

CB Classification

- ▶ enclaves are propagated between servers in the same system over IIOp
- ▶ they are reclassified for HTTP, JMS, and SOAP and for IIOp remote



© Copyright IBM Corp. 2003. All rights reserved.

WLM - Table of Possibilities 1

	coming	WLM Enclave	Service Class chosen (example)	Server Region
1	from IHS/WAS plugin in	created under CB subsys rules (a)	CB01	SR1
2	HTTP Transport Handler (b)	created under CB subsys rules (a)	CB02	SR3
3	HTTP Transport Handler (c)	created using value from variable CB03 (c)	CB03	SR4
4	IIOp local/sysplex	created under callers token or CB subsys rules	CB04	SR5
5	IIOp remote	created under CB subsys rules (a)	CB05	SR6

- MIN=1
- MAX not set
- Multiple Instances enabled

- (a) generic J2EE server name or USERID
- (b) TRANSACTION_CLASS not set
- (c) HTTP_TRANSACTION_CLASS=CB03

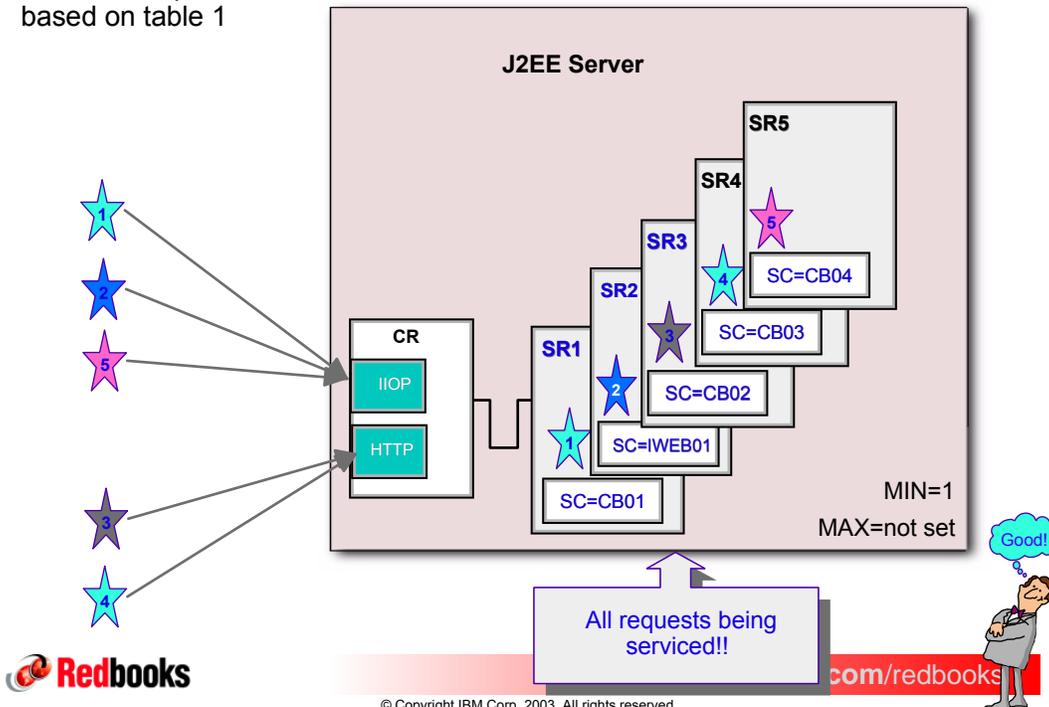


ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM - Classification - example 1

obs.: example based on table 1



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM - Table of Possibilities 2

	coming	WLM Enclave	Service Class chosen (example)	Server Region
1	from IHS/WAS plugin in	created under CB subsys rules (a)	CB01	SR1
2	HTTP Transport Handler (b)	created under CB subsys rules (a)	CB02	SR2
3	HTTP Transport Handler (c)	created using value from variable CB03 (c)	CB03	SR3
4	IIOp local/sysplex	created under callers token or CB subsys rules	CB04	SR4
5	IIOp remote	created under CB subsys rules (a)	CB05	?

- MIN=1
- MAX=4
- Multiple Instances enabled "

- (a) generic J2EE server name or USERID
- (b) HTTP_TRANSACTION_CLASS not set
- (c) HTTP_TRANSACTION_CLASS=CB03

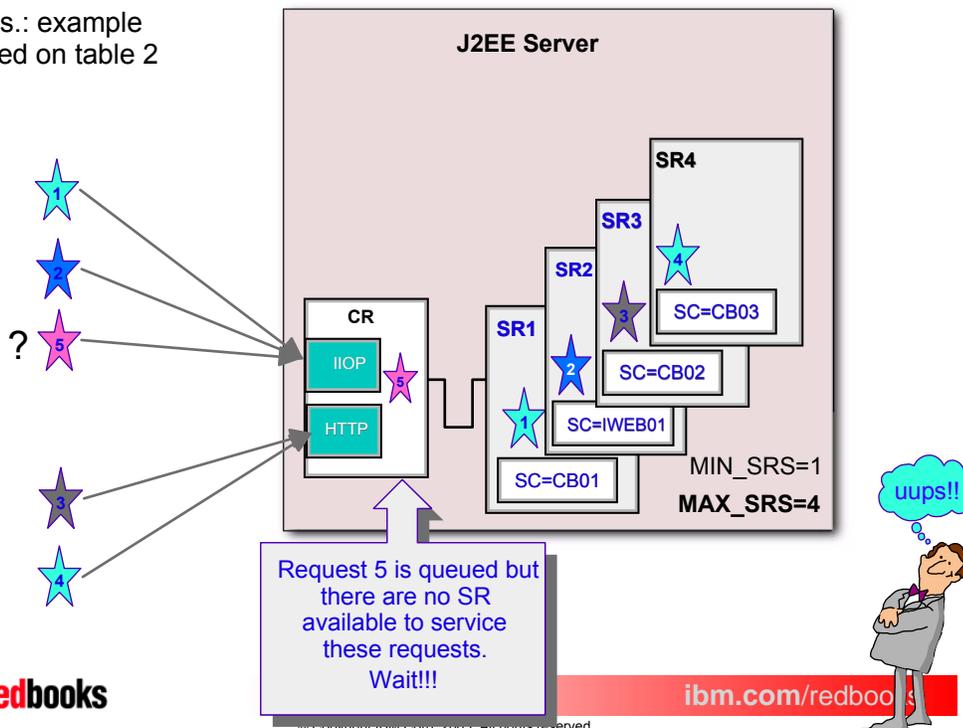


ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

WLM - Classification - example 2

obs.: example based on table 2



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

How many Server Regions ?

Only one:

- ▶ Cannot recycle servers for leaking memory

More than one:

- ▶ session externalization or session affinity

To many:

- ▶ high warm-up and JITting times
- ▶ High storage consumption (1 JVM heap per SR)
- ▶ Unnecessary resource consumption

Not enough:

- ▶ Queuing at a high arrival rate (see also # threads)
- ▶ Not enough heap for keeping sessions in memory
- ▶ Cannot connect address space to WLM queue



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

One tran or multiple trans?

You also have a choice of running server regions with an isolation policy of one tran per server region or multiple trans per server region. From a performance perspective, running more threads in a server region will consume less memory but at the cost of thread contention. This contention is application-dependent. We generally recommend the use of multiple trans unless you run into contention problems.

Specify the threads setting using the `server_region_workload_profile`. The variables include:

ISOLATE - sets the value to 1 thread.

CPUBOUND

IOBOUND - default

LONGWAIT - 40



ibm.com/redbooks

© Copyright IBM Corp. 2003. All rights reserved.

