SHARE in Atlanta, March 2012

z/OS WLM: The Basics Every Performance Analyst Should Know

Session 10888

Glenn Anderson, IBM Technical Training

© 2011 IBM Corporation

z/OS Workload Manager (WLM)

- A contract between the installation and the z/OS operating system

- Installation
  - Classifies work running on z/OS in distinct Service Classes
  - Defines goals that express the expectation of how work should perform
- WLM
  - Uses goal definitions to manage work across all systems of a sysplex through distribution of resources

© Copyright IBM Corporation 2012

Service definition hierarchy

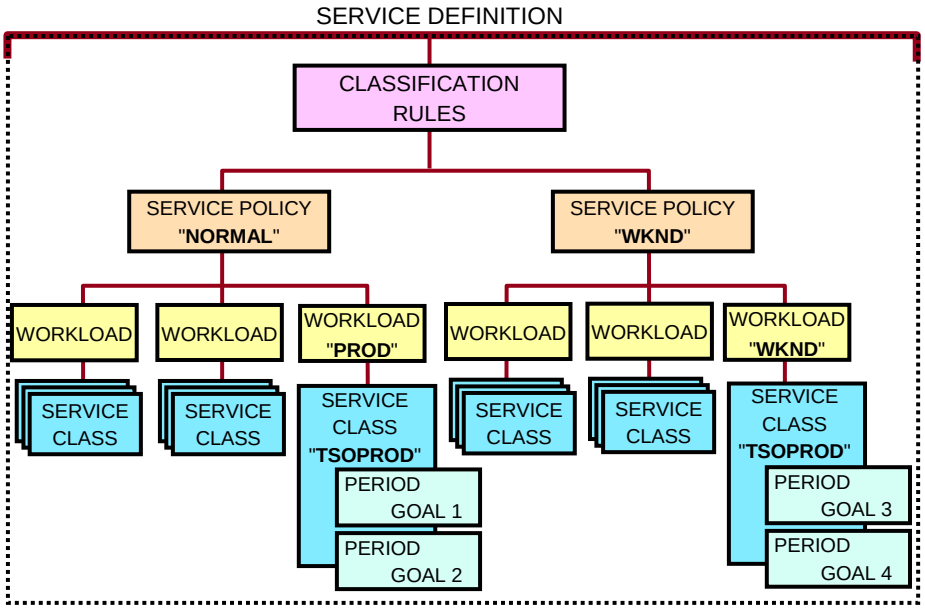

© Copyright IBM Corporation 2012

Parallel Sysplex plus WLM

What is needed?

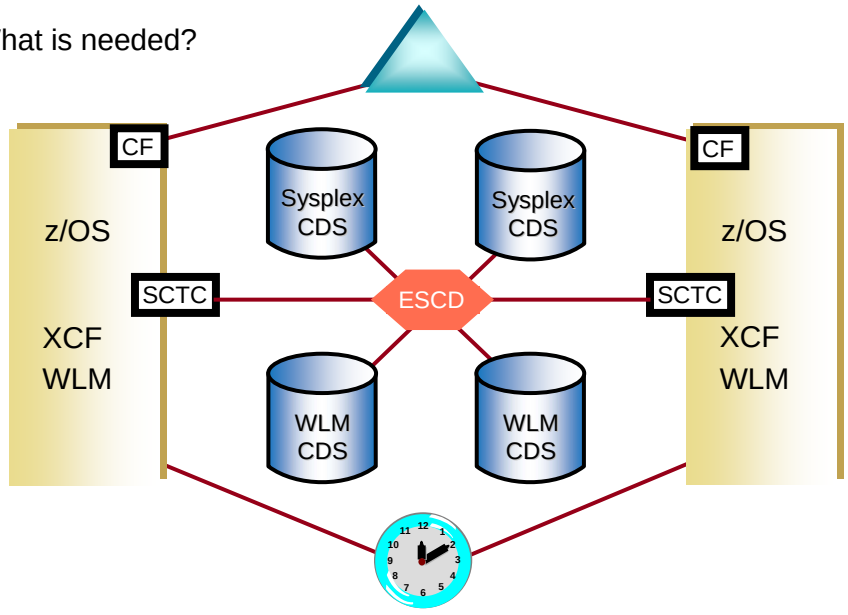

© Copyright IBM Corporation 2012

Service definition structure

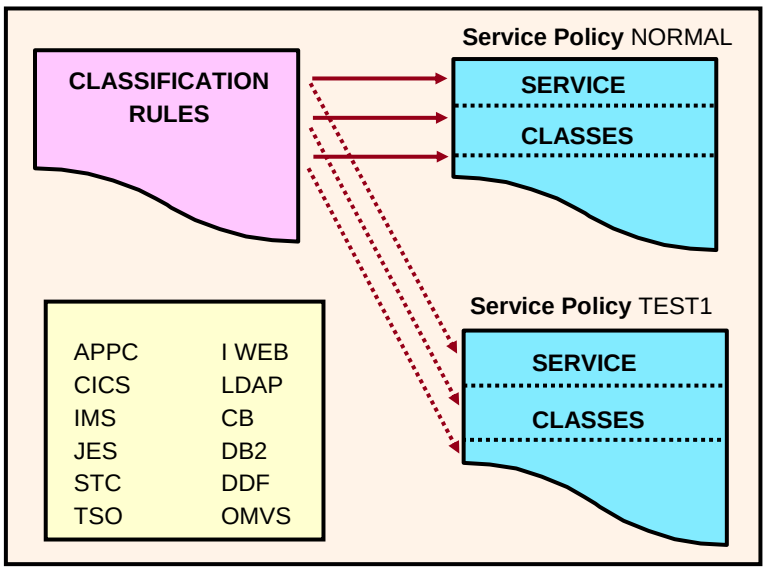

© Copyright IBM Corporation 2012

Service definition workload types

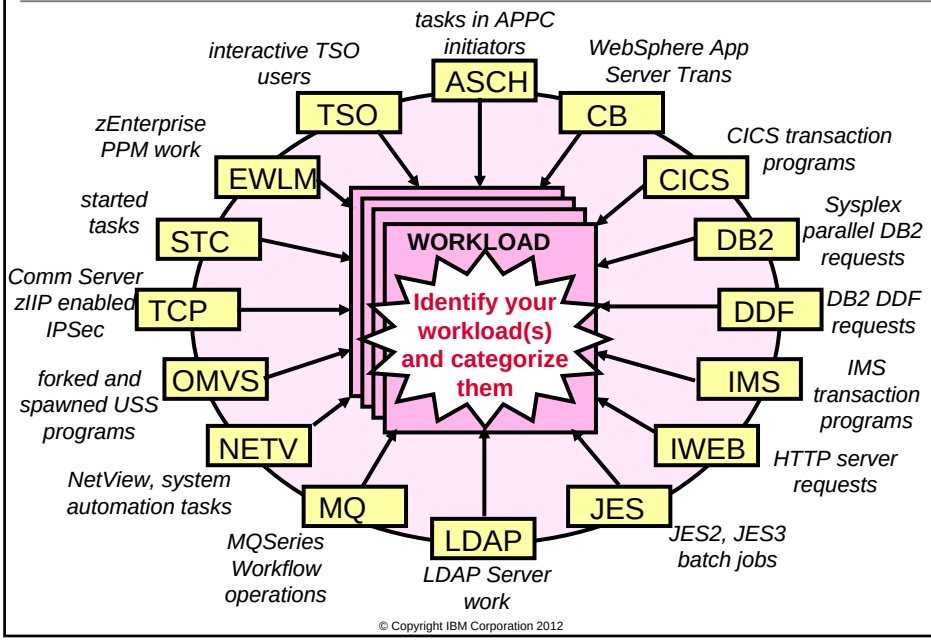

© Copyright IBM Corporation 2012

Classification rules

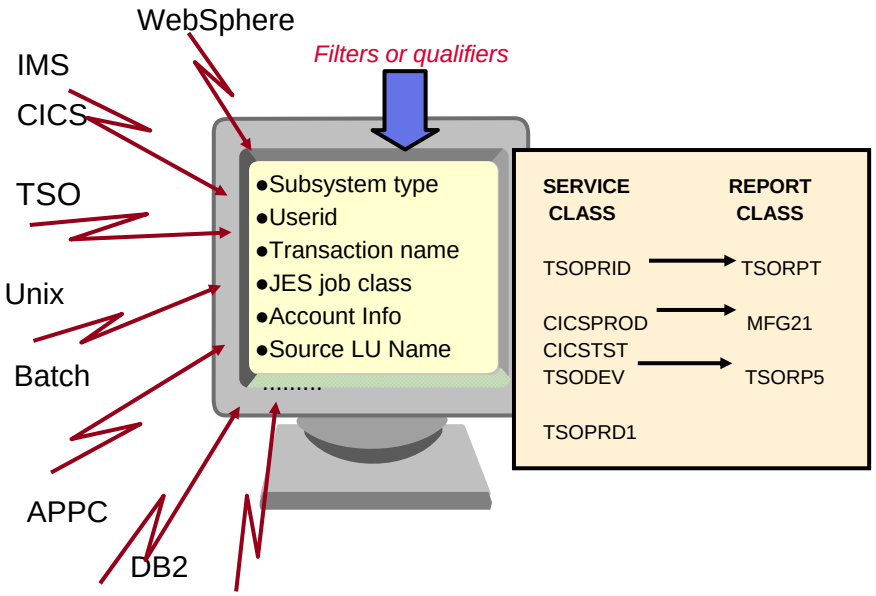

© Copyright IBM Corporation 2012

Example of a service class

```
Service Class: TSOPRD
Description: Normal TSO Work
Workload: TSO
Service Goal:
 Performance Period:1
   Response Time: 80% AT 1.0s
   Importance: 3
   Duration: 800
 Performance Period:2
   Discretionary
```

© Copyright IBM Corporation 2012

The WLM Original ISPF application

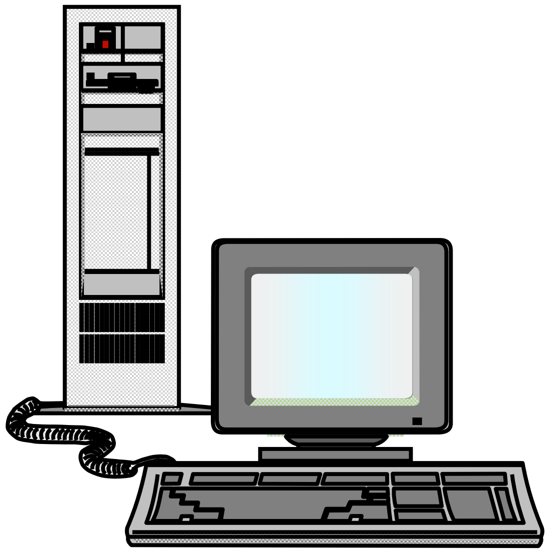

© Copyright IBM Corporation 2012

The WLM welcome screen

```
 File  Help
------------------------------------------------------------------------
Command ===> ___________________________________________________________



                          W   W   L      M   M
                          W   W   L      MM MM
                          W W W   L      M M M
                          WW WW   L      M   M
                          W   W   LLLLL  M   M

                 Licensed Materials - Property of IBM

                 5647-A01 (C) Copyright IBM Corp. 2001.
                 All rights reserved.

                           ENTER to continue
```

© Copyright IBM Corporation 2012

Choose service definition

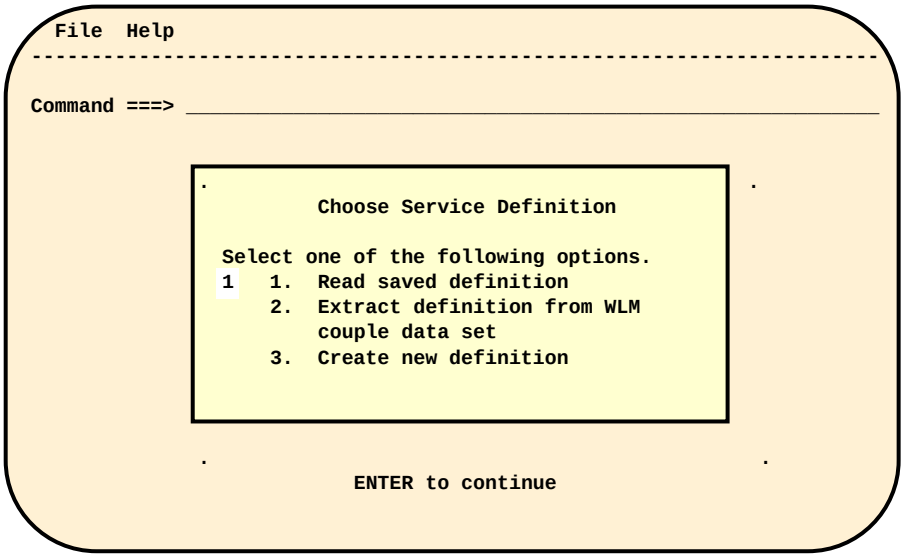

© Copyright IBM Corporation 2012

Service definition menu

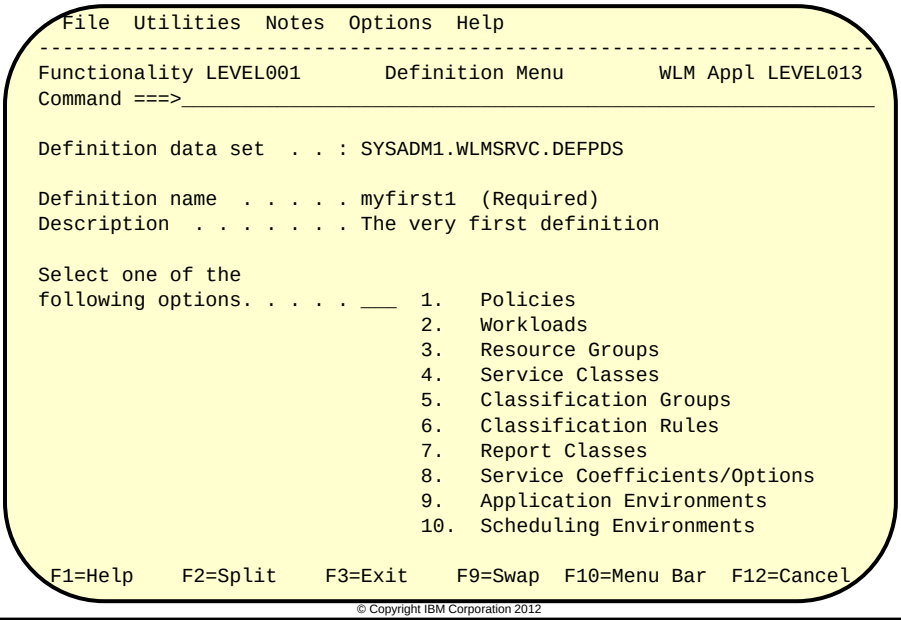

© Copyright IBM Corporation 2012

IBM z/OSMF Management Facility

z/OS application, browser access

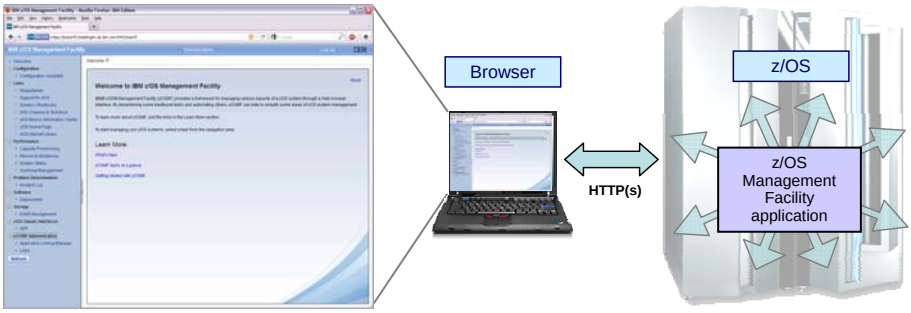

- *z/OS Management Facility is a Web 2.0 application on z/OS*
  - Manages z/OS from z/OS
  - Browser communicates with z/OSMF via secure connection, anywhere, anytime

Workload Management Task Overview

- **Manage Service Definitions:** Create, modify, import, export, print, install service definitions
- **Manage Service Policies for Sysplex:** Activate or view the service policies in the service definition that is currently installed in the WLM couple data set
- **Manage Settings:** Specify history length, codepage, user preferences
- **View Status for Sysplex:** Displays information about the service definition installed in the WLM couple data set and the service policy active in the sysplex.

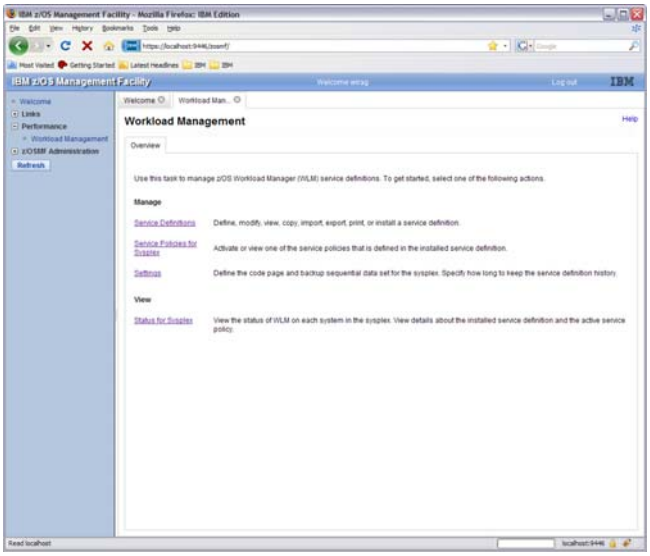

© Copyright IBM Corporation 2012

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

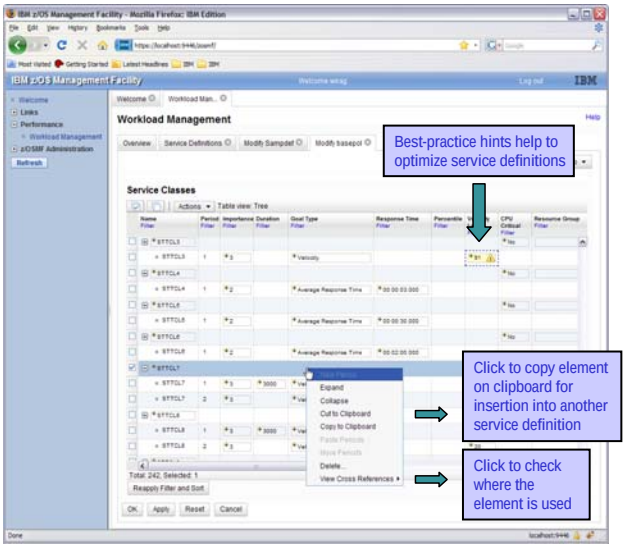

WLM Component Environment Overview

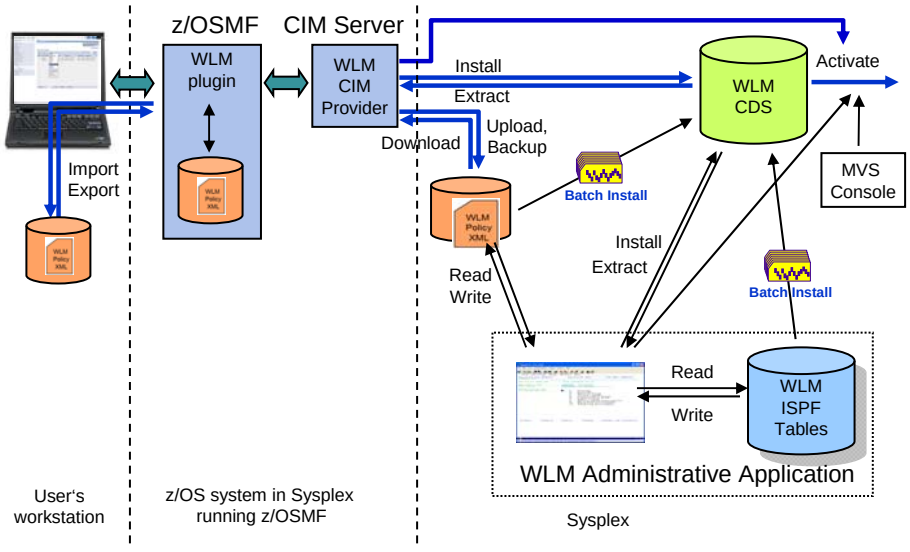

© Copyright IBM Corporation 2012

The structure of a service class

A service class is build from the following characteristics:

- Performance Period
- Duration
- Goal Type and Setting
  - Average Response Time
  - Response Time and Percentile
  - Velocity
  - Discretionary
- Importance

© Copyright IBM Corporation 2012

What is a WLM transaction?

- A WLM transaction represents a WLM "unit of work"
  - Basic workload entity for which WLM collects a resource usage value
  - Foundation for statistics presented in workload activity report
  - Represents a single subsystem "work request"
- Subsystems can implement one of three transaction types
  - Address Space:
    - WLM transaction measures all resource used by a subsystem request in a single address space
    - Used by **JES** (a batch job), **TSO** (a TSO command), **OMVS** (a process), **STC** (a started task) and **ASCH** (single APPC program)
  - Enclave:
    - Enclave created and destroyed by subsystem for each work request
    - WLM transaction measures resources used by a single subsystem request across multiple address spaces and systems
    - Exploited by "new workload" subsystems - **C**omponent **B**roker (WebSphere), **DB2**, **DDF**, **IWEB**, **MQ** (Workflow), **LDAP**, **NETV**, **TCP**
  - CICS/IMS Transactions
    - Neither address space or enclave oriented - special type
    - WLM transaction measures resource used by a single CICS/IMS transaction program request

© Copyright IBM Corporation 2012

Response time goals

- Average Response Time Goal
  - Defines the average transaction response time for all ended transactions
  - Example: Average response time = 1 second

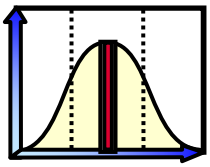

$$\text{Average Response Time} = \frac{\text{Sum of elapsed time for ended transactions*}}{\text{Number of ended transactions*}}$$

- Percentile Response Time Goal
  - Defines the number of transactions ending with a response time lower than or equal to the time value
  - Example: Goal = 90% < 1 sec

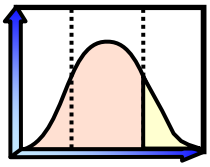

$$\text{Percentile Response Time} = \frac{\text{No. of transactions ended with time<= goal*}}{\text{Number of ended transactions*}}$$

*Measured in a given interval

© Copyright IBM Corporation 2012

Velocity goals

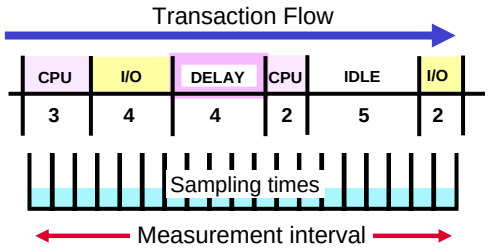

$$\text{Velocity} = \frac{\text{(CPU Using + I/O Using) x 100}}{\text{CPU Using + I/O Using + WLM Delay*}}$$

$$= \frac{11 \times 100}{11 + 4} = 73\%$$

- Delayed I/O requests queued by goal achievement, not DP
- I/O Using
  - Includes non-paging DASD I/O only
  - Device connect
- I/O Delay
  - IOS queue
  - Subchannel pending
  - CU queue

*Delay = CPU Delay + I/O Delay + Paging Delay + MPL Delay + A/S Delay

© Copyright IBM Corporation 2012

Goal type: Discretionary

WLM-defined Goal Type to run the work and apply resources only when there are resources left over.

- Discretionary Workload receives resources from:
  - Higher importance work
  - Overachieving its goal if giving resources will not cause goals to be missed
- Runs in lowest MTTW dispatching priority
- Always last period in a service class

© Copyright IBM Corporation 2012

Multiple periods and velocity goals

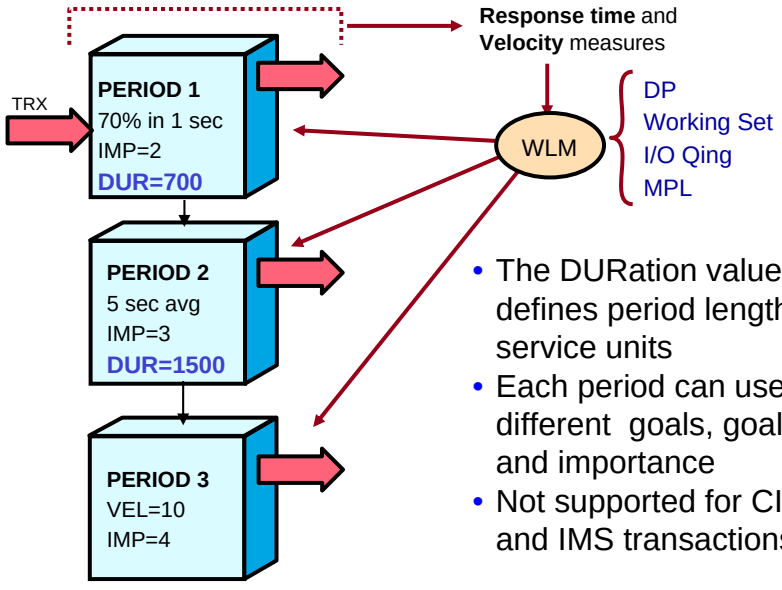

- The DURation value defines period length in service units
- Each period can use different goals, goal types, and importance
- Not supported for CICS and IMS transactions

© Copyright IBM Corporation 2012

Importance

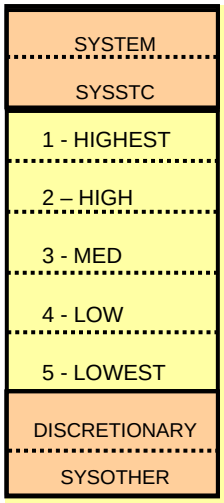

Workload Importance

- Importance is relevant when system is overloaded
- WLM uses it to decide which workload goals are most important to satisfy
- Generally importance decreases across multiple periods

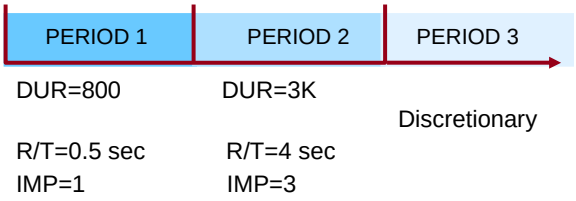

© Copyright IBM Corporation 2012

System goals and dispatching priority

WLM-defined Goal Types that are automatically assigned to certain types of workload recognized by WLM.

- SYSTEM
  - z/OS system address spaces created at IPL.
  - Highest dispatching priority.
- SYSSTC
  - You assign important STCs
  - Second highest dispatching priority.
- SYSOTHER
  - Catcher for forgotten subsystem definitions.
  - Same as discretionary. Lowest dispatching priority.

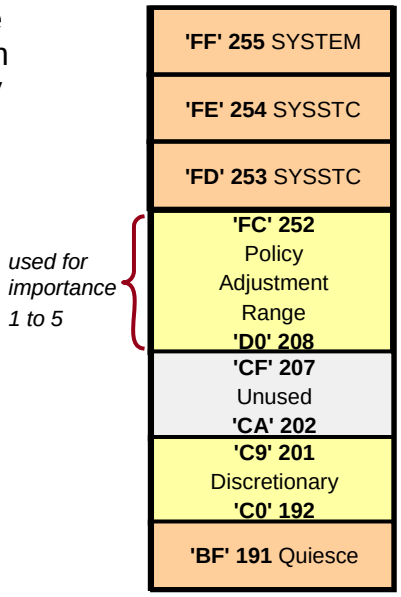

© Copyright IBM Corporation 2012

Goals: System-wide view

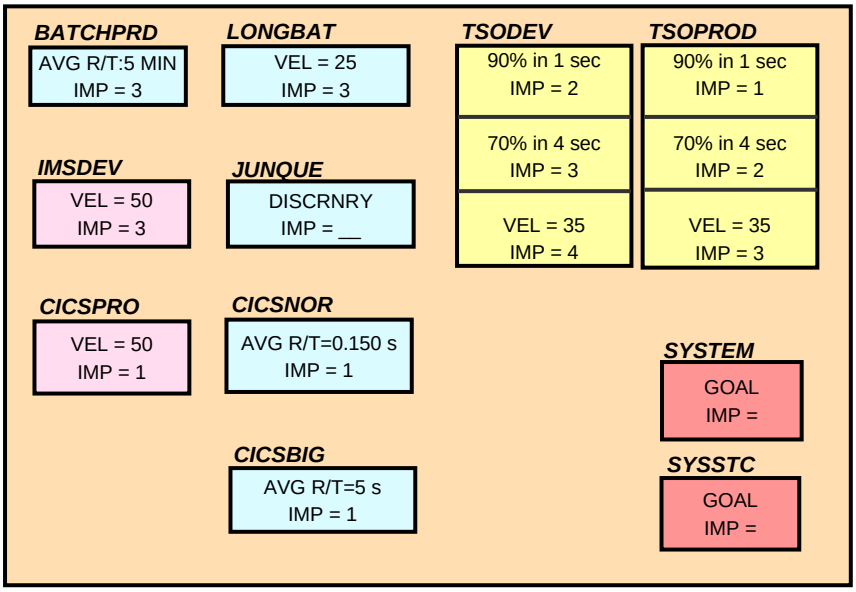

© Copyright IBM Corporation 2012

Typical goals (1 of 3)

Typical Goals (1 of 3)

Workload	Service Class	Goal	Period	Duration	Importance
TSO	TSOPROD	80% within 1 second	1	500	2
		80% within 10 seconds	2	10000	2
		VEL=5	3		3
	TSODEV	80% within 1 second	1	400	2
		VEL=20	2	10000	2
		VEL=20	3		4
BATCH	BATP	VEL=15	1		3
	BATDEV	VEL=10	1		4/D
	BATNOR	AVG=1 min	1	200K	3
		VEL=10	2		5

© Copyright IBM Corporation 2012

Typical goals (2 of 3)

Typical Goals (2 of 3)					
Workload	Service Class	Goal	Period	Duration	Importance
CICS	CICSNOR	90% within 0.2 second	1		1
	CICSSPCL	70% within 0.1 second	1		1
	CICSLOW	Average 5 seconds	1		1
IMS	IMSNOR	90% within 0.4 second	1		1
	IMSHOT	80% within 0.2 secs	1		1
	IMSLOW	Average 5 seconds	1		5
DDF	DB1A	70% within 0.5 second	1	2000	3
		VEL=10	2		3
	DB1B	VEL=10	1		3
APPC	APP1	70% within 0.5 second	1	1000	2
		VEL=20	2		3

© Copyright IBM Corporation 2012

Typical goals (3 of 3)

Typical Goals (3 of 3)					
Workload	Service Class	Goal	Period	Duration	Importance
OMVS	UNIX1	80% within 0.5 second	1	500	2
		VEL=20	2		3
TRNMGR	VEL50I1	VEL=50	1		1
STC	STCHI	VEL=40	1		2
	STCMED	VEL=15	1		3
	STCLOW	VEL=5	1		5/D

© Copyright IBM Corporation 2012

The classification process

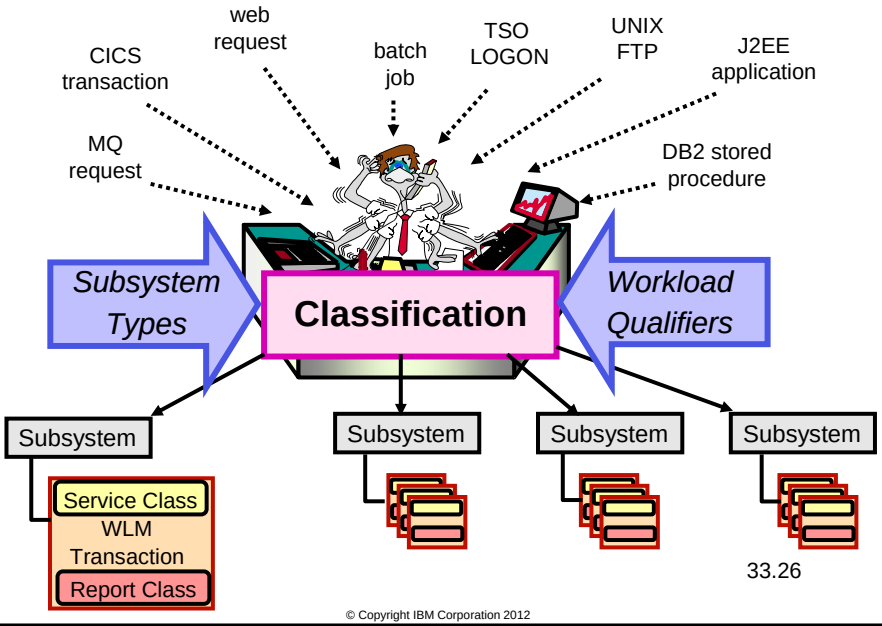

© Copyright IBM Corporation 2012

Getting to the classification rules

```
File   Utilities            Notes  Options       Help
------------------------------------------------------------
                      Definitions Menu
Definition data set . . :    'SYSADM1.WLMSRVC.DEFPDS'

Definition name . . . . .    myfirst1      Required)
Description . . . . . . . . My first service definition

Select one of the
following options . . .  __  1. Policies
                            2. Workloads
                            3. Resource Groups
                            4. Service Classes
                            5. Classification Rules
                            6. Classification Groups
                            7. Report Classes
                            8. Service Coefficients
                            9. Application Environments
                           10. Scheduling Environments
Command ===> _____________________________________________________
F1=Help F2=Split F3=Exit F9=Swap F10=Menu Bar      F12=Cancel
```

© Copyright IBM Corporation 2012

Subsystem types used by classification

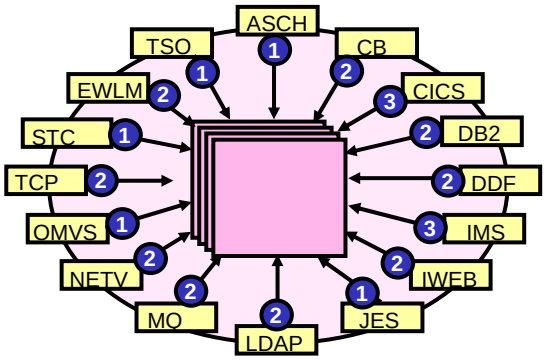

- Subsystems follow one of three transaction type models
- Need to understand how this affects the value of figures shown in workload activity report

Transaction Type	Allowable Goal Types	Allowable # Periods
Address space oriented (1)	Response Time Execution Velocity Discretionary	Multiple
Enclave (2)	Response Time Execution Velocity Discretionary	Multiple
CICS/IMS (3)	Response Time	1

© Copyright IBM Corporation 2012

Workload qualifiers supported by WLM (1 of 2)

	ASCH	CB	CICS	DB2	DDF	IMS	IWEB	JES	LSFM	MQ	NETV	OMVS	SOM	STC	TSO	SYSH
Accounting Information (AI)	*			*	*			*				*		*	*	
Collection Name (CN)		*		*	*								*			
Connection Type (CT)				*	*											
Correlation Information (CI)				*	*											
LU Name (LU)			*	*	*	*					*					
Netid (NET)				*	*	*										
Package Name (PK)				*	*											
Perform (PF)				*				*						*	*	
Plan Name (PN)				*	*											
Priority (PRI)				*				*		*	*					
Procedure Name (PR)				*	*											
Process Name (PC)				*	*					*						

© Copyright IBM Corporation 2012

Workload qualifiers supported by WLM (2 of 2)

	ASCH	CB	CICS	DB2	DDF	IMS	IWEB	JES	LSFM	MQ	NETV	OMVS	SOM	STC	TSO	SYSH
Scheduling Environment Name (SE)				*				*								
Subsystem Collection Name (SSC)				*	*			*								
Subsystem Instance (SI)		*	*	*	*	*	*	*	*	*	*					
Subsystem Parameter (SPM)				*			*			*			*	*		
Sysplex Name (PX)	*	*	*	*	*	*	*	*	*	*		*	*	*	*	*
System Name (SY)	*											*		*	*	*
Transaction Class/Job Class (TC)	*	*		*		*	*	*		*	*					
Transaction Name/Job Name (TN)	*	*	*	*		*	*	*	*	*	*	*		*		
Userid (UI)	*	*	*	*	*	*	*	*		*	*	*	*	*	*	

© Copyright IBM Corporation 2012

Example of batch classification rules

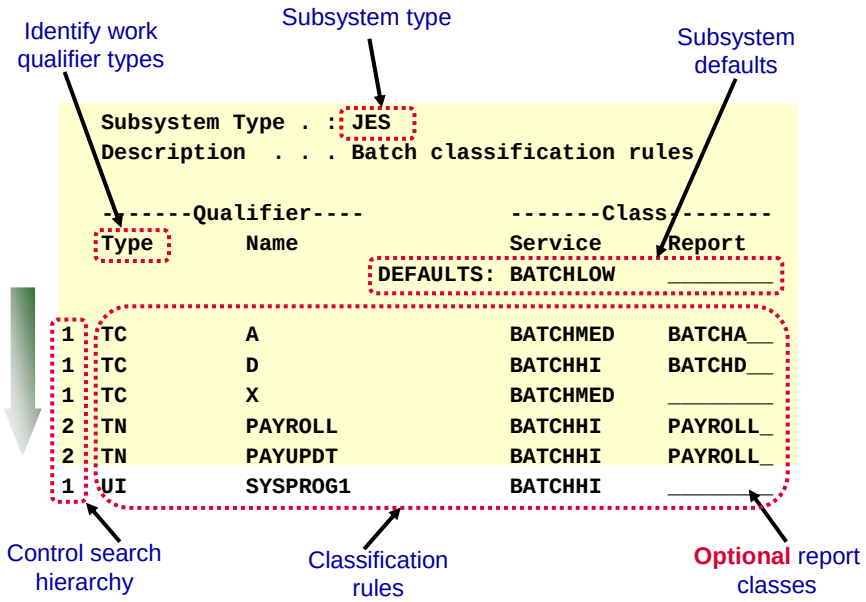

© Copyright IBM Corporation 2012

When Do the "OMVS" Classification Rules Apply?

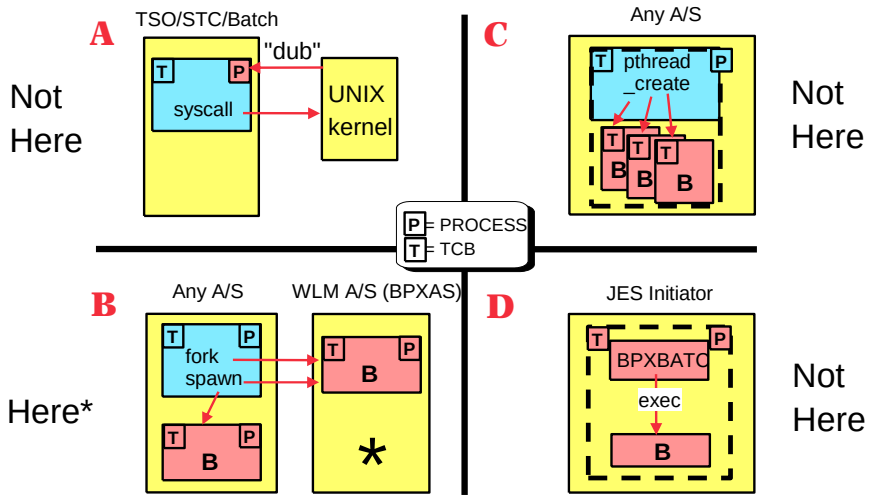

© Copyright IBM Corporation 2012

Enclave Services: A Dispatching Unit

- **Standard dispatching**
  - dispatchable units (DUs) are the TCB and the SRB
    - TCB runs at dispatching priority of address space and is pre-emptible
    - SRB runs at supervisory priority and is non-pre-emptible
- **Advanced dispatching units**
  - **Enclave**
    - Anchor for an address space-independent transaction managed by WLM
    - Can comprise multiple DUs (TCBs and Enclave SRBs) executing across multiple address spaces
  - **Enclave SRB**
    - Created and executed like an ordinary SRB but runs with Enclave dispatching priority and is pre-emptible
- **Enclave Services enable a workload manager to create and control enclaves**

© Copyright IBM Corporation 2012

Enclave Characteristics

- Created by an address space (the "owner")
- One address space can own many enclaves
- One enclave can include multiple dispatchable units (SRBs/tasks) executing concurrently in multiple address spaces (the "participants")
  - Enclave SRBs are preemptible, like tasks
  - All its dispatchable units are managed as a group
- Many enclaves can have dispatchable units running in one participant address space concurrently
- RMF produces separate T72 SMF records for independent enclaves

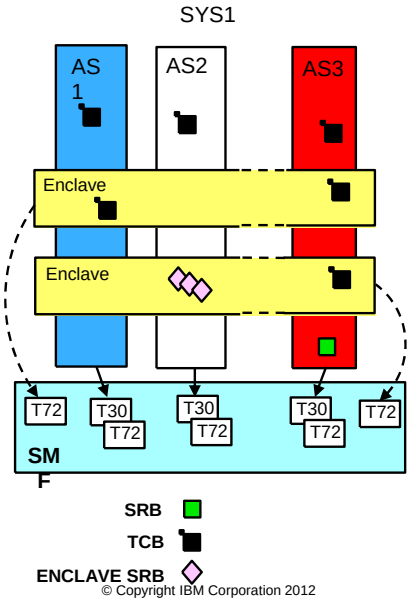

© Copyright IBM Corporation 2012

DDF and Enclave SRBs

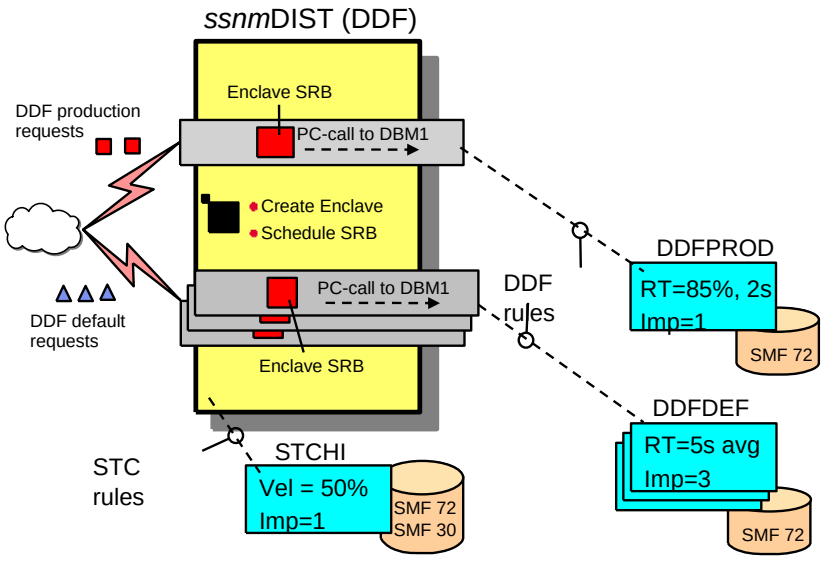

© Copyright IBM Corporation 2012

WebSphere App Server Use of Enclaves

WAS on z/OS transactions, arriving at the Control Region, each run in an enclave that is classified under the "CB" rules.

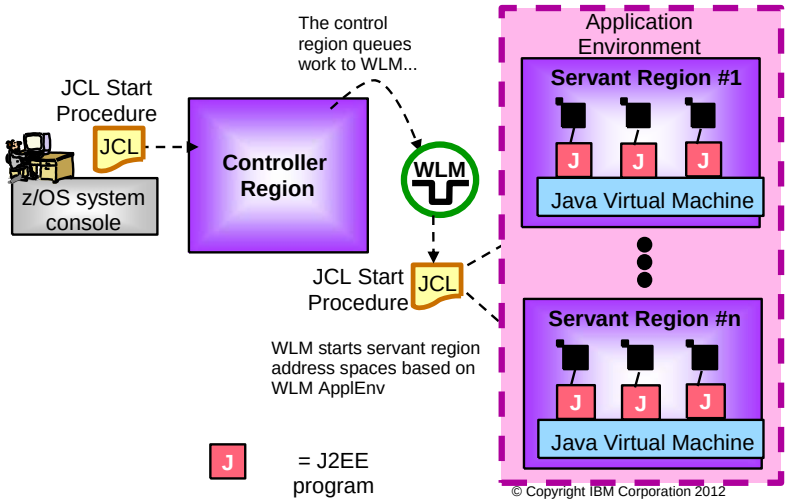

© Copyright IBM Corporation 2012

WLM Managed DB2 Stored Procedures

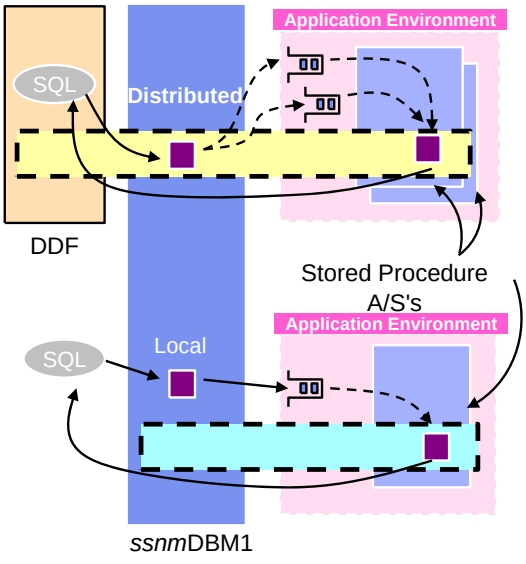

- **Stored Procedures run in WLM managed server regions in an application environment**
- **Distributed requests (DDF)**
  - DBM1 processes SQL request under existing Enclave
- **Local requests (CICS, batch)**
  - DBM1 creates a dependent Enclave
- **DBM1 inserts work request into WLM work queue**
- **Available task in server region selects the request and processes it under the Enclave**

© Copyright IBM Corporation 2012

Execution Delay Monitoring Services

Allows you to assigns goals to the transaction and let the system determine which work managers need the resources to meet these goals

- Infrastructure that allows WLM to assign resources for address spaces based on the behavior of the transactions being serviced by them
- Exploited by complex work managers, that do not allow WLM to individually manage resource consumption of the transactions
- Work managers can report their own view of transaction states
- Exploiting subsystems
  - CICS
  - IMS

© Copyright IBM Corporation 2012

CICS / IMS Transactional Goal Management

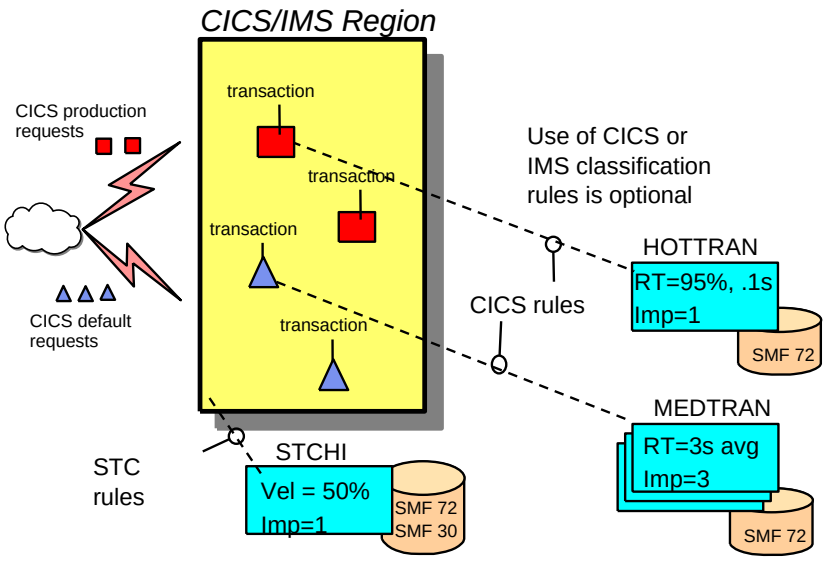

© Copyright IBM Corporation 2012

Manage Region Using Goals Of.....

STC Classification Rules

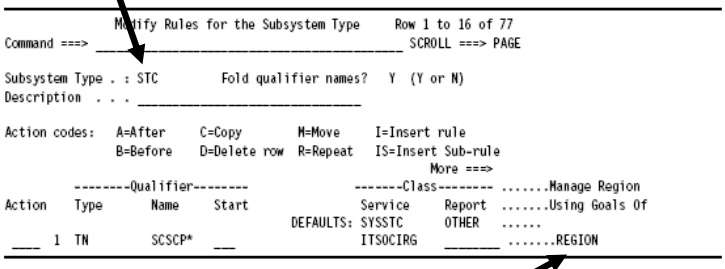

Manage Region Using Goals Of
- REGION
- TRANSACTION

© Copyright IBM Corporation 2012

Performance index

- Separate PI is calculated for each service class period, across all sysplex hosts
- Different calculations produce a comparable value from different goal types

© Copyright IBM Corporation 2012

Policy adjustment cycle

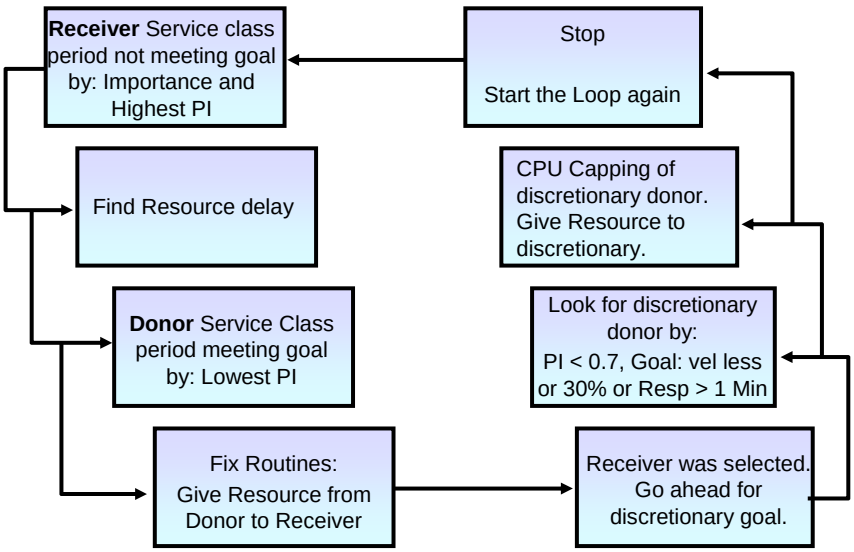

- Adjustment is made on the basis of selected service class **periods**

© Copyright IBM Corporation 2012

Blocked Workload Support: IEAOPTxx

BLWLTRPCT	**Percentage of the CPU capacity of the LPAR to be used for promotion** • **Specified in units of 0.1%** • **Default is 5 (=0.5%)** • **Maximum is 200 (=20%)** • **Would only be spent when enough units of work exist which need promotion**
BLWLINTHD	**Specifies threshold time interval for which a blocked address space or enclave must wait before being considered for promotion.** **•Minimum is 5 seconds. Maximum is 65535 seconds.** **•Default is 20 seconds.**

© Copyright IBM Corporation 2012

RMF Workload Activity Report – Batch

```
  REPORT BY: POLICY=WLMPOL01   WORKLOAD=BATCH      SERVICE CLASS=BATI2V50   RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=2
                                                  CRITICAL     =NONE

-TRANSACTIONS-  TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE---   SERVICE TIME  ---APPL %---  --PROMOTED--  ----STORAGE----
AVG    1025.81  ACTUAL                 0  SSCHRT  1279  IOC       9537K  CPU 20893.24  CP     2313.1 BLK     0.000  AVG     1454.83
MPL    1025.81  EXECUTION              0  RESP     0.8  CPU       4357M  SRB    76.975  AAPCP   8.24  ENQ     1.427  TOTAL   1492382
ENDED        0  QUEUED                 0  CONN     0.4  MSO       1344M  RCT     0.000  IIPCP   0.00  CRM     0.000  SHARED   595.92
END/S     0.00  R/S AFFIN              0  DISC     0.2  SRB      17583K  IIT     7.672                 LCK    59.319
#SWAPS       0  INELIGIBLE             0  Q+PEND   0.2  TOT       5729M  HST     0.029  AAP    17.53                 -PAGE-IN RATES-
EXCTD        0  CONVERSION             0  IOSQ     0.0  /SEC      6366K  AAP   157.808  IIP     0.00                 SINGLE       0.0
AVG ENC   0.00  STD DEV                0                                  IIP     0.000                               BLOCK        0.0
REM ENC   0.00                                           ABSRPTN  6205                                              SHARED       0.0
MS ENC    0.00                                           TRX SERV 6205                                              HSP          0.0

GOAL: EXECUTION VELOCITY 50.0%     VELOCITY MIGRATION:   I/O MGMT  54.3%     INIT MGMT 54.3%

         RESPONSE TIME EX   PERF  AVG   --EXEC USING%--  -------------- EXEC DELAYS % -----------  -USING%-  --- DELAY % ---    %
SYSTEM                 VEL% INDX ADRSP  CPU AAP IIP I/O  TOT CPU                                    CRY CNT  UNK IDL CRY CNT  QUI

*ALL        --N/A--    54.3  0.9  1026  2.3 0.0 0.0 0.1  2.1 2.0                                    0.0 0.0   95 0.6 0.0 0.0  0.0
JA0                    68.9  0.7 256.0  2.1 0.0 N/A 0.1  1.0 1.0                                    0.0 0.0   96 0.4 0.0 0.0  0.0
JB0                    53.8  0.9 256.0  3.8 0.0 0.0 0.3  3.5 3.5                                    0.0 0.0   92 0.4 0.0 0.0  0.0
J80                    51.3  1.0 256.0  1.1 0.0 0.0 0.1  1.2 1.1                                    0.0 0.0   97 0.4 0.0 0.0  0.0
J90                    46.8  1.1 255.0  2.2 0.0 0.0 0.1  2.7 2.6                                    0.0 0.0   95 0.0 0.0 0.0  0.0
TPN                    75.8  0.7   3.0  0.5 0.0 0.0 1.0  0.5 0.1                                    0.0 0.0  0.0  98 0.0 0.0  0.0
```

© Copyright IBM Corporation 2012

RMF Workload Activity Report – TSO (1 of 2)

```
REPORT BY: POLICY=WLMPOL01   WORKLOAD=TSO        SERVICE CLASS=TSO        RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=2
                                                 CRITICAL     =NONE

-TRANSACTIONS-  TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE---   SERVICE TIME  ---APPL %---  --PROMOTED--  ----STORAGE----
AVG     293.25  ACTUAL              2.335  SSCHRT 187.4  IOC      81821K  CPU  951.417  CP     114.26  BLK     0.000  AVG      2616.59
MPL     293.23  EXECUTION           2.335  RESP     0.4  CPU     197358K  SRB   33.108  AAPCP   0.00  ENQ     0.000  TOTAL  767254.5
ENDED    80683  QUEUED                 0  CONN     0.3  MSO     106510K  RCT   42.824  IIPCP   0.00  CRM     0.000  SHARED   563.70
END/S    89.66  R/S AFFIN              0  DISC     0.0  SRB       7339K  IIT    1.079                 LCK    62.436
#SWAPS   46123  INELIGIBLE             0  Q+PEND   0.2  TOT     393027K  HST    0.009  AAP     0.00                 -PAGE-IN RATES-
EXCTD        0  CONVERSION             0  IOSQ     0.0  /SEC    436732   AAP    0.000  IIP     0.00                 SINGLE       0.0
AVG ENC   0.00  STD DEV             5.296                                  IIP    0.000                               BLOCK        0.0
REM ENC   0.00                                           ABSRPTN  1489                                              SHARED       0.0
MS ENC    0.00                                           TRX SERV 1489                                              HSP          0.0

GOAL: RESPONSE TIME 000.00.02.000 AVG

         RESPONSE TIME EX   PERF  AVG   --EXEC USING%--  -------------- EXEC DELAYS % -----------  -USING%-  --- DELAY % ---    %
SYSTEM   HHH.MM.SS.TTT VEL% INDX ADRSP  CPU AAP IIP I/O  TOT CPU                                    CRY CNT  UNK IDL CRY CNT  QUI

*ALL     000.00.02.335 76.4  1.2 561.0  0.2 0.0 0.0 0.0  0.1 0.1                                    0.0 0.0  2.6  97 0.0 0.0  0.0
JA0      000.00.01.654 65.5  0.8 117.1  0.1 0.0 N/A 0.1  0.1 0.1                                    0.0 0.0  2.0  98 0.0 0.0  0.0
JB0      000.00.02.169 65.2  1.1 254.7  0.2 0.0 0.0 0.0  0.1 0.1                                    0.0 0.0  2.2  98 0.0 0.0  0.0
JC0      000.00.01.612 79.2  0.8  61.3  0.1 0.0 0.0 0.0  0.0 0.0                                    0.0 0.0  5.2  95 0.0 0.0  0.0
JE0      000.00.02.685 95.0  1.3  44.0  0.8 0.0 N/A 0.0  0.0 0.0                                    0.0 0.0  3.0  96 0.0 0.0  0.0
JF0      000.00.02.923 92.6  1.5  13.0  0.3 0.0 0.0 0.0  0.0 0.0                                    0.0 0.0  2.5  97 0.0 0.0  0.0
J80      000.00.11.747 66.7  5.9  30.0  0.0 0.0 0.0 0.0  0.0 0.0                                    0.0 0.0  0.4  99 0.0 0.0  0.0
J90      000.00.02.655 92.9  1.3  21.0  1.0 0.0 0.0 0.0  0.1 0.1                                    0.0 0.0  2.1  97 0.0 0.0  0.0
TPN      000.00.04.188 82.4  2.1   8.0  0.1 0.0 0.0 0.1  0.0 0.0                                    0.0 0.0  7.4  92 0.0 0.0  0.0
Z0       000.00.01.689 76.0  0.8  12.0  0.2 0.0 0.0 0.0  0.1 0.0                                    0.0 0.0  4.0  96 0.0 0.0  0.0
```

© Copyright IBM Corporation 2012

RMF Workload Activity Report – TSO (2 of 2)

```
----------RESPONSE TIME DISTRIBUTION----------
    ----TIME----    --NUMBER OF TRANSACTIONS--    -------PERCENT-------  0    10   20   30   40   50   60
 70   80   90   100
    HH.MM.SS.TTT    CUM TOTAL        IN BUCKET    CUM TOTAL   IN BUCKET
 |....|....|....|....|....|....|....|....|....|....|
 <  00.00.01.000        40076            40076         49.7        49.7  >>>>>>>>>>>>>>>>>>>>>>>>>
 <= 00.00.01.200        43459             3383         53.9         4.2  >>>
 <= 00.00.01.400        43654              195         54.1         0.2  >
 <= 00.00.01.600        47904             4250         59.4         5.3  >>>
 <= 00.00.01.800        48135              231         59.7         0.3  >
 <= 00.00.02.000        48450              315         60.0         0.4  >
 <= 00.00.02.200        51009             2559         63.2         3.2  >>
 <= 00.00.02.400        51254              245         63.5         0.3  >
 <= 00.00.02.600        55364             4110         68.6         5.1  >>>
 <= 00.00.02.800        55563              199         68.9         0.2  >
 <= 00.00.03.000        55819              256         69.2         0.3  >
 <= 00.00.04.000        63291             7472         78.4         9.3  >>>>>
 <= 00.00.08.000        76667            13376         95.0        16.6  >>>>>>>>>
 >  00.00.08.000        80683             4016          100         5.0  >>>
```

© Copyright IBM Corporation 2012

IBM Technical Training

- Basic z/OS Tuning Using the Workload Manager (ES545) – 4.5 days, hands-on labs
- Advanced z/OS Performance: WLM, Sysplex, Unix Services, and Web - 4.5 days
- ibm.com/training

© Copyright IBM Corporation 2012