

**SHARE**

Technology • Connections • Results


# z/OS Workload Management Update for z/OS V1.11 and V1.12

Stefan Wirag (stefan.wirag@de.ibm.com)  
IBM Corporation

Monday, February 28, 2011  
Session 8857



# Trademarks



**SHARE**

Technology • Connections • Results

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

**The following are trademarks or registered trademarks of other companies.**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* 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.

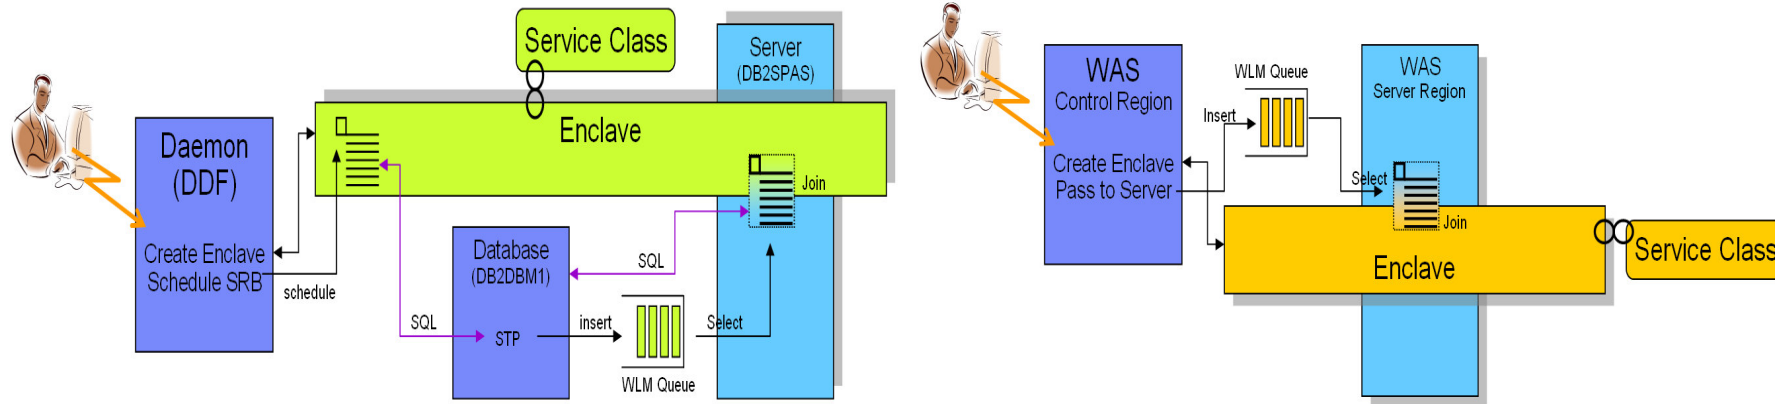
**SHARE**  
in Anaheim  
2011

# Agenda



- Enclave Enhancements
  - Enclave Server Management
  - Work-Dependent Enclaves
- WLM Management
  - LDAP Support
  - Resource Group Enhancements
  - Do not always honor Skip Clock in Policy Adjustment
- WLM Reporting
  - Extend Number of Report Classes
  - Response Time Distribution for Velocity Goals
- Externalized IEAOPT Information
- Hyperdispatch APAR
- WLM support for IBM zEnterprise 196
- z/OSMF Workload Management
- WLM support for zManager
- WLM Tools Overview

# WLM Enclaves – An Overview




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



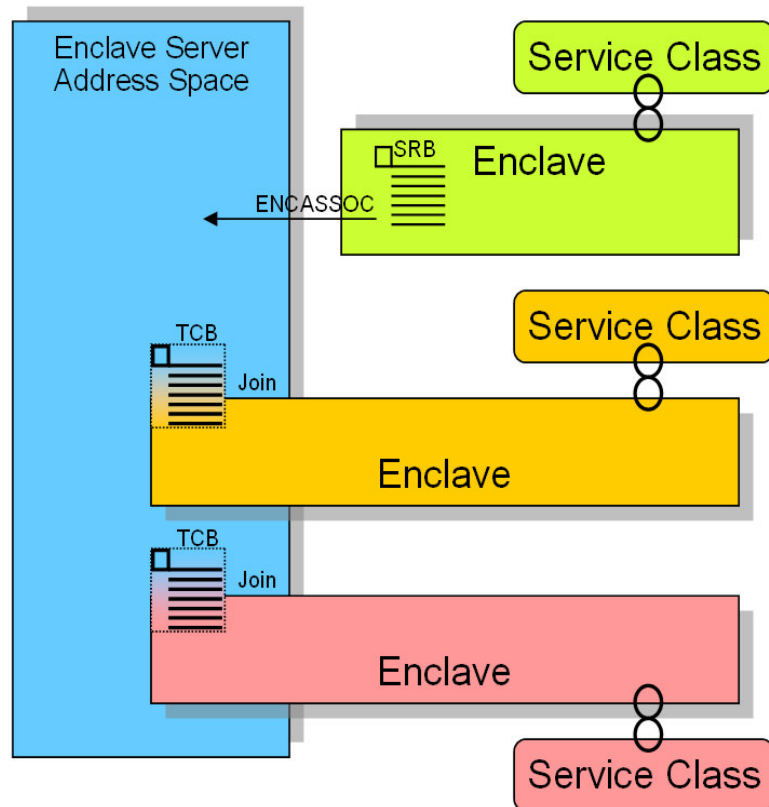
# WLM Enclave Server Address Spaces

## A Short Retrospective



**SHARE**

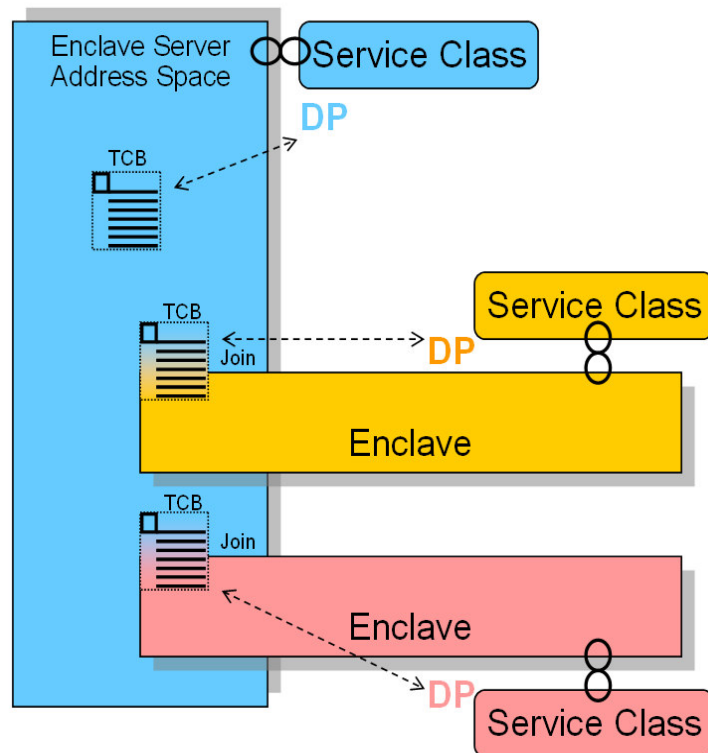
Technology • Connections • Results



- An address space becomes an enclave server when
  - An enclave SRB issues SYSEVENT ENCASSOC
  - A TCB of the address space joins an enclave, and does not specify ENCLAVESERVER=NO (which is typically not the case)
- Assumption (Programming Model)
  - All work being executed within the address space is related to enclaves
  - That means
    - There is no significant amount of work (TCBs) executing in such address spaces which is not related to enclaves
- Enclave Server Management
  - CPU and I/O DP is derived from service class of most important enclave
    - Meaning: No CPU and I/O management exists for these server address spaces
  - Storage management is done to meet the served enclave's goals

**SHARE**  
in Anaheim  
2011

# WLM Enclave Server Management Changes with z/OS 1.12



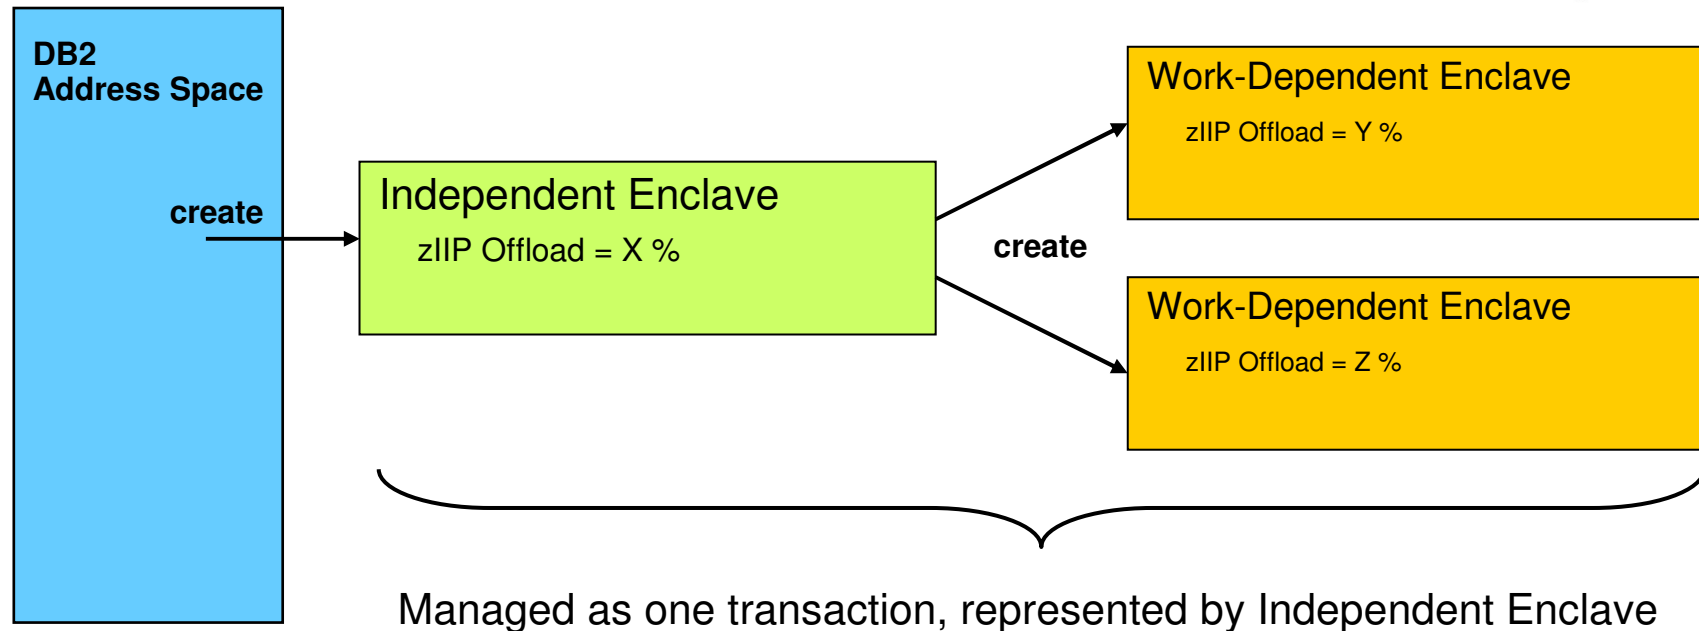
- New IEAOPT Parameter
  - ManageNonEnclaveWork = {**No**|Yes}
    - Default: No (no change to previous releases)
  - Causes everything in the address space, which is not associated to an enclave, to be managed towards the goals of the external Service Class to which the address space has been classified to
- Advantages
  - Enclave (Queue) server address spaces in which no enclave is running will be managed as usual address spaces
  - The importance and goal of the service class for the address space now has a meaning
- **Note: With ManageNonEnclaveWork = Yes the importance and goal of the service class for the address space is more important than it used to be**
  - Verify goal settings for server address spaces

# Work-Dependent Enclaves



- Background
  - zIIPs allow middleware components to run a certain percentage of their work “offloaded” from regular processors
  - The offload percentage is an attribute of the enclave under which the unit of work runs
  - The offload percentage is defined by the middleware component.
- Limitations
  - It is not possible to specify different offload percentages for different units of work running under the same enclave
- Intended Use Case
  - 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)

# Work-Dependent Enclaves




- Solution

Implement a new type of enclave named "Work-Dependent" as an extension of an Independent Enclave. A Work-Dependent enclave becomes part of the Independent Enclave's transaction but allows to have its own set of attributes (including zIIP offload percentage)

# Work-Dependent Enclaves

## Reporting in SDSF Enclave Panel and RMF Monitor III



**SHARE**

Technology • Connections • Results

| Display              | Filter      | View     | Print        | Options      | Search | Help     |          |         |    |  |
|----------------------|-------------|----------|--------------|--------------|--------|----------|----------|---------|----|--|
| SDSF ENCLAVE DISPLAY |             | SYS1     | ALL          | LINE 1-8 (8) |        |          |          |         |    |  |
| COMMAND INPUT ==>    |             |          |              | SCROLL ==>   |        |          |          | CSR     |    |  |
| PREFIX=* DEST=(ALL)  |             | OWNER=*  | SYSNAME=SYS1 |              |        |          |          |         |    |  |
| NP                   | NAME        | Status   | Type         | SrvClass     | Per    | RptClass | CPU-Time | OwnerAS | Re |  |
|                      | 28000000006 | ACTIVE   | IND          | VEL_1        | 1      | RC_2     | 0.00     | 36      |    |  |
|                      | 2C000000008 | ACTIVE   | WDEP         | VEL_1        | 1      | RC_2     | 0.83     | 36      |    |  |
|                      | 30000000007 | ACTIVE   | WDEP         | VEL_1        | 1      | RC_2     | 0.83     | 36      |    |  |
|                      | 34000000009 | ACTIVE   | WDEP         | VEL_1        | 1      | RC_2     | 0.83     | 36      |    |  |
|                      | 3800000000A | ACTIVE   | WDEP         | VEL_1        | 1      | RC_2     | 0.83     | 36      |    |  |
|                      | 3C00000000B | ACTIVE   | WDEP         | VEL_1        | 1      | RC_2     | 0.83     | 36      |    |  |
|                      | 24000000002 | INACTIVE | DEP          | SYSSTC       | 1      | RC_0     | 0.00     | 22      |    |  |
|                      | 20000000001 | INACTIVE | DEP          | SYSTEM       | 1      | RC_0     | 0.00     | 7       |    |  |

```


RMF V1R12 Enclave Report
Command ==>
Samples: 100      System: SYS1  Date: 02/23/10  Time: 03.06.40
Current options:  Subsystem Type: ALL
                  Enclave Owner:
                  Class/Group:

Enclave  Attribute  CLS/GRP  P Goal  % D X  EAppl%  TCPU
*SUMMARY
ENC000006  VEL_1      1      5 W  0.812  0.163  2.530
ENC000002  VEL_1      1      5 W  0.163  0.163  2.532
ENC000004  VEL_1      1      5 W  0.162  0.162  2.528
ENC000005  VEL_1      1      5 W  0.162  0.162  2.519
ENC000003  VEL_1      1      5 W  0.162  0.162  2.518
ENC000001  VEL_1      1      5 W  0.000  0.000  0.007

```

F1=help F2=SPLIT F3=  
F7=UP F8=DOWN F9=  
MA c

# Enclave Enhancements: Availability



**SHARE**


Technology • Connections • Results

| Function                       | z/OS<br>V1.12 | z/OS<br>V1.11 | z/OS<br>V1.10 | Older<br>Releases     |
|--------------------------------|---------------|---------------|---------------|-----------------------|
| Non Shell Server<br>Management | +             |               |               |                       |
| Work-dependent Enclaves        | +             | +             | OA26104       | OA26104<br>→ z/OS 1.8 |

- Non Shell Server Management
  - New OPT Parameter ManageNonEnclaveWork=YES/NO.  
Default is NO, meaning the function is not yet enabled
- Work-Dependent Enclaves
  - New function available with WLM APAR OA26104
  - DB2 exploitation with APAR PK76676
  - SDSF support with APAR PK74125
  - RMF support with z/OS 1.11

# Agenda



- Enclave Enhancements
  - Enclave Server Management
  - Work-Dependent Enclaves
-  • WLM Management
  - LDAP Support
  - Resource Group Enhancements
  - Do not always honor Skip Clock in Policy Adjustment
- WLM Reporting
  - Extend Number of Report Classes
  - Response Time Distribution for Velocity Goals
- Externalized IEAOPT Information
- Hyperdispatch APAR
- WLM support for IBM zEnterprise 196
- z/OSMF Workload Management
- WLM support for zManager
- WLM Tools Overview

# WLM Management: LDAP Subsystem is supported




|                             | L<br>D<br>A<br>P |
|-----------------------------|------------------|
| Accounting Information      |                  |
| Collection Name             |                  |
| Connection Type             |                  |
| Correlation Information     |                  |
| EWLM Service Class          |                  |
| EWLM Transaction Class      |                  |
| LU Name                     |                  |
| Netid                       |                  |
| Package Name                |                  |
| Perform                     |                  |
| Plan Name                   |                  |
| Priority                    |                  |
| Procedure Name              |                  |
| Process Name                |                  |
| Scheduling Environment Name |                  |
| Subsystem Collection Name   |                  |
| Subsystem Instance          | ●                |
| Subsystem Parameter         |                  |
| Sysplex Name                | ●                |
| System Name                 |                  |
| Transaction Class/Job Class |                  |
| Transaction Name/Job Name   | ●                |
| Userid                      |                  |

- Work requests include all work processed by the z/OS LDAP server
- Supported Work Qualifiers
  - Subsystem Instance (SI)  
The z/OS LDAP server's job name. Needed to distinguish between different LDAP servers
  - Transaction Name/Job Name (TN)  
The z/OS LDAP server's enclave transaction name.  
"GENERAL" for all LDAP work that is not assigned a user-defined exception class. Any transaction name that is also defined in the configuration file of the directory server
- For further information see  
*z/OS IBM Tivoli Directory Server Administration and Use for z/OS (SC23-5191-XX)*



# WLM Management: Resource Group Enhancements



**SHARE**

Technology • Connections • Results

```
vamp - wc3270
Resource-Group Xref Notes Options Help
-----
Modify a Resource Group
Command ==> _____
Enter or change the following information:
Resource Group Name . . . . . : RGROUP1
Description . . . . . : Res_Group_No._1
Define Capacity:
1. In Service Units (Sysplex Scope)
2. As Percentage of the LPAR share (System Scope)
3. As a Number of CPs times 100 (System Scope)
Minimum Capacity . . . . . : 0000
Maximum Capacity . . . . . : 15000000
```

```
vamp - wc3270
File Utilities Notes Options Help
-----
Resource Group Compatibility Check

One or several Resource Group capacity values in this WLM Service
Definition exceed 999,999. Please make sure to have APAR OA29704
applied on all systems in your Sysplex that are currently at level
z/OS V1R11 (FMID HBB7760) or previous releases. If this APAR is not
installed, the Service Definition can get corrupted when it is
opened, modified, or installed via such a system and will not be
extractable from there.

Have you applied the necessary maintenance (OA29704) to all systems
running with z/OS V1R11 and previous releases in your Sysplex ?
1. YES, WLM APAR OA29704 has been installed on all systems
   running with z/OS V1R11 or previous releases.
2. NO, I am not sure if the necessary service is installed and
   I will have to check first. Terminate the installation of
   the Service Definition.

F1=Help      F2=Split    F5=KeysHelp  F9=Swap      F12=Cancel
```

- OA29704 for z/OS 1.10 and z/OS 1.11 allows you to specify new limits of up to 8 digits.
- Because this is a PTF (APAR) a warning message is shown when a min/max capacity value greater than 6 digits is entered.

Make sure that all systems are at least at z/OS 1.10 with OA29704 applied before installing and activating such a service definition

- On systems w/o this support the WLM Administrative Application is not able to extract the service definition from the Couple Data Set.
- On systems w/o this support the WLM Administrative Application would truncate the resource group capacity values to 6 digits if it is attempted to read the Service Definition from ISPF tables.

Regardless of whether or not the APAR has been applied, systems w/o the support honor the definition during runtime.

# WLM Management: Do Not Always Honor “Skip Clock”




- What is the skip clock ?
  - If WLM cannot help a service class it sets a skip clock to not assess it in the next 3 policy adjustment cycles
  - This is done for efficiency reasons and to help other work
- Is this always a good thing to do ?
  - Usually yes!
  - But if only very few service classes miss their goals it is not beneficial to no longer assess a service class for 3 consecutive policy adjustment cycles
    - *Especially when it might be possible to help the work with IRD Weight Changes. In this event the situation on another LPAR can change and might make it possible to help a service class in the next policy adjustment cycle*
- Solution introduced with z/OS 1.11



The skip clock will no longer be honored if 5 or less service class periods do not meet their performance objectives.


# WLM Management Availability



**SHARE**  
Technology • Connections • Results

| Function  | z/OS V1.12 | z/OS V1.11 | z/OS V1.10 | Older Releases |
|---|------------|------------|------------|----------------|
| New Resource Groups (Type 2 and Type 3)                 | +          | +          | +          | z/OS 1.8       |
| 8 digit resource group minimum and maximum (for Type 1) | +          | OA29704    | OA29704    |                |
| Change in skip clock processing                         | +          | +          |            |                |
| LDAP Support  | +          | +          |            |                |

# Agenda



**SHARE**  
Technology • Connections • Results

- Enclave Enhancements
  - Enclave Server Management
  - Work-Dependent Enclaves
- WLM Management
  - LDAP Support
  - Resource Group Enhancements
  - Do not always honor Skip Clock in Policy Adjustment
- WLM Reporting
  - Extend Number of Report Classes
  - Response Time Distribution for Velocity Goals
- Externalized IEAOPT Information
- Hyperdispatch APAR
- WLM support for IBM zEnterprise 196
- z/OSMF Workload Management
- WLM support for zManager
- WLM Tools Overview




**SHARE**  
in Anaheim  
2011

# WLM Reporting: Extend Number of Report Classes



- The WLM-supported maximum on the number of defined report classes (999) has become insufficient for large installations or consolidations
- Solution
  - Extend number of report classes in multiple steps:
    - First Step (z/OS 1.11):
      - *Extend to 2047 Report Classes*
      - *Expand internal data structures to be able to deal with 4095 report classes*
- Remarks
  - New WLM functionality level in z/OS 1.11: LEVEL023
  - For Service Definitions in XML format, the corresponding XML namespace is <http://www.ibm.com/xmlns/prod/zwlm/2009/09/ServiceDefinition.xsd>

# Extended Number of Report Classes Availability



**SHARE**  
Technology • Connections • Results

| Function            | z/OS V1.12 | z/OS V1.11 | z/OS V1.10 | Older Releases |
|---------------------|------------|------------|------------|----------------|
| 2047 Report Classes | +          | +          |            |                |

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



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

# Response Time Distribution for Velocity Goals



- IWMWRCAA enhancements

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




# Response Time Distribution for Velocity Goals



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

# Response Time Distribution for Velocity Goals – RMF WLMGL Enhancement



**SHARE**

Technology • Connections • Results

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

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

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


| SYSTEM | RESPONSE TIME EX | PERF | AVG   | --EXEC USING%-- | ----- EXEC DELAYS % ----- | -USING%- | --- DELAY % ---       | %   |
|--------|------------------|------|-------|-----------------|---------------------------|----------|-----------------------|-----|
|        | VEL%             | INDX | ADRSP | CPU AAP IIP I/O | TOT                       | CRY CNT  | UNK IDL CRY CNT       | QUI |
| *ALL   | --N/A--          | 88.2 | 0.2   | 47.0            | 0.0 0.0 0.0 0.2           | 0.0      | 0.0 0.0 38 62 0.0 0.0 | 0.0 |
| SYSD   |                  | 88.2 | 0.2   | 15.0            | 0.0 0.0 0.0 0.3           | 0.0      | 0.0 0.0 40 60 0.0 0.0 | 0.0 |
| SYSE   |                  | 88.6 | 0.2   | 17.0            | 0.0 0.0 0.0 0.3           | 0.0      | 0.0 0.0 35 64 0.0 0.0 | 0.0 |

## -----RESPONSE TIME DISTRIBUTIONS-----

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


# Response Time Distribution for Velocity Goals

## The Mid-Point Change Algorithm



**SHARE**  
Technology • Connections • Results

- Rationale
  - Velocity goals do not have a „reference“ response time
  - The mid-point (MP) should be set to values which accurately reflect the current workload conditions
  - These conditions may drastically change from time to time, but WLM expects them to be consistent for a time long enough, so that it can compute sensible mid-point (MP) values
- Algorithm
  - The model behind the algorithm is a gaussian RT distribution, with  $d = M/3$  (M is the mean, d is the standard deviation). Ideally is  $MP = M$
  - If a trx response time fulfills  $MP/3 \leq RT \leq MP*3$  counter C is decreased by a value that reflects the distance of RT to the MP
  - If a trx response time does not fulfill  $MP/3 \leq RT \leq MP*3$  counter C is increased by a value that reflects the distance of RT to the MP
  - If counter C becomes greater than a threshold, a new mid-point is calculated




This is the ideal case:

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

# Response Time Distribution for Velocity Goals



## The Mid-Point Change Algorithm



**SHARE**



Technology • Connections • Results

$M = 2.791$ ,  $d = 0.93$



With a mid-point 2.791x smaller than M, WLM will change the mid-point after about 68 minutes.

("moderate" deviation)



M is 4.176 x bigger, so the # transactions in buckets 24 to 28 (> 3x mid-point) is more than 80% of the total.


Change after  $\approx 21$  min.

("strong" deviation)

$M = 4.176$ ,  $d = 1.392$

# New Programming Interface for Monitors

## Control Block: IRARMCTZ



**SHARE**

Technology • Connections • Results

- New extension to SRM Control Table (PI) for information which is of interest for externalization
  - For example all information related to RMF's Monitor II OPT report is included in this table

```

boewlm1 - wc3270
RMF - OPT Settings
Line 1 of 29
Scroll ==> PAGE

CPU= 4    UIC= 65K PR= 0    System= WLM1 Total

OPT: 00    Time: N/A
-- Parameter -- - Default - -- Value -- Unit ----- Description -----
ABNORMALTERM      Yes      Yes Y/N  Abnormal terminations in routing
BLWLINTHD         20      20 sec  Time blocked work waits for help
BLWLTRPCT         5        5 0/00  CPU cap. to promote blocked work
CCCAWMT           12000    12000 usec Alternate wait management time
ZAAPAWMT           12000    12000 usec AWM time value for zAAPs
ZIIPAWMT           12000    12000 usec AWM time value for zIIPs
CNTCLIST          No      No Y/N  Clist commands count individually
CPENABLE          10,30|0,0  10,30 %  Threshold for TPI (low,high)
DVIO              Yes      Yes Y/N  Directed VIO is active
ERV              500      500/CB SU Enqueue residency CPU Service/DP
HIPERDISPATCH    No      No/No Y/N Hiperdispatch is desired/active
IFAHONORPRIORITY  Yes      Yes Y/N  Allows CPs to help zAAPs
IIPHONORPRIORITY  Yes      Yes Y/N  Allows CPs to help zIIPs
INITIMP           0        0/FE #   INITIMP value/DP for initiators
IRA40SI           70,50,50  70,50,50 % Fixed storage of <16M,16M-2G,tot
MAXPROMOTETIME    6        6 *10s  Holder allowed to run promoted
MCCAFCTH          400,800    400,800 # Threshold for storage (low,ok)
MCCFXEPR           92      92 %    Fixed storage threshold < 16 MB
MCCFXTPR           80      80 %    Fixed online storage threshold
PROJECTCPU         No      No Y/N  CPU projection for zAAPs, zIIPs
RCCFXET           82,88    82,88 %  Physical MPL threshold (low,high)
RCCFXTT           66,72    66,72 %  Logical MPL threshold (low,high)
RMPITOM           1000|3000  3000 msec SRM invocation interval
RTPIFACTOR         100      100 %   PI affects server routing weights
STORAGENSWD        Yes      Yes Y/N  Sets non-swap, ASID non-dispatch.
STORAGEWTOR        Yes      Yes Y/N  WTOR to cancel AS in shortage
VARYCPU            Yes      Yes Y/N  VARYCPU is enabled
VARYCPUMIN         1        1 #     VARYCPUMIN value
WASROUTINGLEVEL    0        0 #     WebSphere routing level

F1=HELP  F2=SPLIT  F3=END  F4=RETURN  F5=RFIND  F6=SORT
F7=UP    F8=DOWN   F9=SWAP  F10=LEFT  F11=RIGHT F12=RETRIEVE
4B      X
    
```

CVTOPCTP ... +254


IRARMCT  
...  
RMCTX3 ... +178

**IRARMCTZ Map**

| Offsets |     | Type/Value  | Len | Name (Dim)                 | Description  |
|---------|-----|-------------|-----|----------------------------|--|
| Dec     | Hex |             |     |                            |  |
| 0       | (0) | STRUCTURE   | 0   | RMCTZ                      |  |
| 0       | (0) | CHARACTER   | 8   | RMCTZ_NAME                 | control block acronym > IRARMCTZ <   |
| 8       | (8) | BITSTRING   | 1   | RMCTZ_VERSION              |  |
|         |     |             |     |                            | Rmctz version  |
| 9       | (9) | BITSTRING   | 1   |                            | Reserved   |
| 10      | (A) | SIGNED      | 2   | RMCTZ_LENGTH               | Size of RMCTZ  |
| 12      | (C) | BITSTRING   | 1   | RMCTZ_LPAR_FLAGS           |  |
|         |     |             |     | (0)                        | LPAR Management flags updated by SRM.  |
|         |     | 1... ....   |     | RMCTZ_LPARMGMT_ENABLED     | "X80" ON if WLM LPAR Management Processing is enabled                                  |
|         |     | ..1... .... |     | RMCTZ_LPAR_VARYCPU_ENABLED | "X40" ON if VARYCPU option is turned on either by default or is explicitly set to 'on' |
| 13      | (D) | CHARACTER   | 1   | RMCTZ_FLAG1                | RMCTZ Flag 1   |
|         |     |             |     | (0)                        |  |
|         |     | 1... ....   |     | RMCTZ_ABN_OPT              | "X80" ABNORMALTERM option set  |
|         |     | ..111 1111  |     | RMCTZ_FLAG1_RSVD1          | "X7F" reserved   |
| 14      | (E) | CHARACTER   | 1   | RMCTZ_FLAG2                | RMCTZ Flag 2   |
|         |     |             |     | (0)                        |  |
|         |     | 1... ....   |     | RMCTZ_VCM_OPT              | "X80" 1:=VCM specified   |
|         |     | ..11 1111   |     | RMCTZ_VCM                  | "X40" 1:=Running in vertical CP management mode  |
|         |     |             |     | RMCTZ_FLAG2_RSVD1          | "X3F" reserved   |

# New Programming Interface for Monitors: Availability

Control Block: IRARMCTZ




**SHARE**  
Technology • Connections • Results

| Function                                  | z/OS<br>V1.12<br><small>as previewed 2/2010</small>                                   | z/OS<br>V1.11  | z/OS<br>V1.10  | Earlier<br>Releases |
|---|---|----------------|----------------|---------------------|
| RMF Monitor II OPT Display                | +   | +              |                |                     |
| WLMOPT Tool (bundled with<br>WLMQUE Tool) | No longer extended<br><small>Still bundled with WLMQUE but on z/OS 1.10 level</small> |                | +              | Since z/OS 1.8      |
| IRARMCTZ                                  | +   | <b>OA31201</b> | <b>OA31201</b> |                     |

- RMF Monitor II OPT Display
  - Replaces WLMOPT Tool
    - Bundled with WLMQUE Tool but no longer extended (remains on z/OS 1.10 level)
    - WLMQUE Tool is still valid (see also WLM Tools summary)
- New data interface for Monitors
  - Introduced with z/OS 1.12, Rollback to z/OS 1.10

**SHARE**  
in Anaheim  
2011

# Hiperdispatch: WLM APARs



**SHARE**  
Technology • Connections • Results

| APAR    | Description  | Close Date | Remark                        |
|---------|--|------------|-------------------------------|
| OA31733 | Corrects calculation of capacity for medium processors | 04/2010    | Affects larger configurations |

**SHARE**  
in Anaheim  
2011

# Agenda



- Enclave Enhancements
  - Enclave Server Management
  - Work-Dependent Enclaves
- WLM Management
  - LDAP Support
  - Resource Group Enhancements
  - Do not always honor Skip Clock in Policy Adjustment
- WLM Reporting
  - Extend Number of Report Classes
  - Response Time Distribution for Velocity Goals
- Externalized IEAOPT Information
- Hyperdispatch APAR
- WLM support for IBM zEnterprise 196
- z/OSMF Workload Management
- WLM support for zManager
- WLM Tools Overview





# WLM Support for IBM zEnterprise 196



- IBM zEnterprise 196 (z196)
  - STSI instruction no longer returns the alternate CPU capability
    - CPU adjustment factors are now calculated based on the Model Capacity Ratings by the machine
  - Supplies additional information about speed change
  - Speed changes may occur due to model changes (*capacity level*), or to physical processor tact (*cycle steering*)
- WLM
  - uses the new MSU values to calculate pricing adjustment factors
  - introduces message IWM064I to explain the reason for a processor speed change
  - makes new HW information available via public data areas IRARCT, IRARMCT, IRARMCTZ and via SYSEVENT QVS

# WLM Support for IBM zEnterprise 196

## New Pricing Adjustment Factors



- z196 provides MSU values instead of the alternate CPU capability together with MP factors table as base for pricing factors
- WLM uses the new MSU values to calculate pricing adjustment factors on z196
  - Add more granular new pricing adjustment factors RCTPCPUA\_actual and RCTPCPUA\_nominal plus a scaling factor RCTCPCPUA\_scaling\_factor
  - Maintain existing RCTPCPUA for compatibility
- Values are also available in SMF30 and SMF89
  - SMF30\_RCTPCPUA\_xxxxxx, SMF30\_Capacity\_...
  - SMF89\_RCTPCPUA\_xxxxxx

# WLM Support for IBM zEnterprise 196

## New Message IWM064I



- Depending on the reason for the speed change one of the following message accompanies

**IWM063I WLM POLICY WAS REFRESHED DUE TO A PROCESSOR SPEED CHANGE**

- IWM064I THE SYSTEM IS RUNNING AT NOMINAL CAPACITY.
- IWM064I THE SYSTEM IS RUNNING AT NOMINAL CAPACITY; MODEL CONVERSION OCCURRED.
- IWM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF A MANUAL CONTROL SETTING.
- IMM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF A MACHINE EXCEPTION CONDITION.
- IWM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF A NON-EXCEPTION MACHINE CONDITION.
- IWM064I THE SYSTEM IS RUNNING WITH REDUCED CAPACITY BECAUSE OF AN EXCEPTION CONDITION EXTERNAL TO THE MACHINE.

# WLM Support for IBM zEnterprise 196 Enhanced API



- Query Virtual Server interface can be used to query a server's capacity status
  - Called by Sysevent QVS (assembler interface) or IWMQVS (C interface)
- The Output is mapped by IRAQVS (assembler mapping) or IWMQVS.H (C header file)
  - New data field QvsCecCapacityStatus has been added to indicate if the machine is running at nominal or reduced capacity
  - Version QvsVer has been incremented to QvsVer2

# WLM Support for IBM zEnterprise 196 Extended Data Areas



- IRARMCTZ
  - RMCTZ\_Capacity\_Change\_Time
    - Time when the capacity was last changed
  - RMCTZ\_Capacity\_Adjustment\_Indication
    - When zero, the indication is not reported. When in the range 1-99, some amount of reduction is indicated. When 100, the machine is operating at its normal capacity. Primary CPUs and all secondary-type CPUs are similarly affected
  - RMCTZ\_Capacity\_Change\_Reason
    - Indicates the reason which is associated with the present value contained in RMCTZ\_Capacity\_Adjustment\_Indication
  - RMCTZ\_CAI\_IPL
    - Capacity adjustment indication at IPL
  - RMCTZ\_CCR\_IPL
    - Capacity change reason at IPL
  - RMCTZ\_nominal\_CPMP
    - Nominal CPU adjustment factor (similar to RMCTCPMP but for nominal speed)
- IRARCT
- IWMWRCAA
  - RCAAADJCCPU                      CPU adjustment factor
  - RCAAADJCCPUNOM                nominal CPU adjustment factor
  - RCAAADJCCEC                      CEC adjustment factor

# WLM Support for IBM zEnterprise 196 Availability



| Function                                     | z/OS<br>V1.12 | z/OS<br>V1.11 | z/OS<br>V1.10 | z/OS<br>V1.9 |
|--|---------------|---------------|---------------|--------------|
| New message, API<br>enhancements             | OA30968       | OA30968       | OA30968       |              |
| Support for new MSU<br>computation/reporting | OA30968       | OA30968       | OA30968       | OA30968      |

# Temporary Capacity Reporting via SYSEVENT REQLPDAT



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

# Temporary Capacity Reporting via SYSEVENT REQLPDAT



- **IRALPDAT new data fields**

- **LPDATMODELCAPIIDENT**

- The 16-character EBCDIC model-capacity identifier of the configuration.

← Permanent+OOCoD + (CBU+PE)

- **LPDATMODEL**

- The 16-character EBCDIC model identifier of the configuration. If not valid, field LPDatModelCapIdent represents both the model-capacity identifier and the model.

← H/W model

- **LPDATMODELPERMCAPIIDENT**

- The 16-character EBCDIC model-permanent capacity identifier of the configuration.

← Permanent configuration

- **LPDATMODELTEMPCAPIIDENT**

- The 16-character EBCDIC model-temporary capacity identifier of the configuration.

← Permanent + OOCoD



# Temporary Capacity Reporting via SYSEVENT REQLPDAT



- **IRALPDAT new data fields**

- **LPDATMODELCAPRATING**

- When non-zero, an unsigned integer whose value is associated with the model capacity as identified by the model-capacity identifier. There is no formal description of the algorithm used to generate this integer.

Permanent+OOCoD + (CUB+PE)

- **LPDATMODELPERMCAPRATING**

- When non-zero, an unsigned integer whose value is associated with the model-permanent capacity as identified by the model-permanent- capacity identifier

Permanent configuration

- **LPDATMODELTEMPCAPRATING**

- When non-zero, an unsigned integer whose value is associated with the model-temporary capacity as identified by the model-temporary- capacity identifier.

Permanent + OOCoD

# REQLPDAT Sample with Active Temporary Capacity



- **Example output**

```
LPDatModelCapIdent      : 714
LPDatModel              : E26
LPDatModelPermCapIdent  : 709
LPDatModelTempCapIdent  : 711
LPDatModelCapRating     : 00000473
LPDatModelPermCapRating : 00000324
LPDatModelTempCapRating : 000003B0
```

- **Meaning**

- The base model is 709
- This model has active OOCOD capacity
  - It temporarily has a capacity like a model 711
- This model also has active CBU capacity
  - In total it temporarily has a capacity like a model 714

# Agenda



- Enclave Enhancements
  - Enclave Server Management
  - Work-Dependent Enclaves
- WLM Management
  - LDAP Support
  - Resource Group Enhancements
  - Do not always honor Skip Clock in Policy Adjustment
- WLM Reporting
  - Extend Number of Report Classes
  - Response Time Distribution for Velocity Goals
- Externalized IEAOPT Information
- Hyperdispatch APAR
- WLM support for IBM zEnterprise 196
- z/OSMF Workload Management
- WLM support for zManager
- WLM Tools Overview




# z/OSMF Workload Management

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



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

# z/OSMF Workload Management Service Definition Repository



**SHARE**

Technology • Connections • Results


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

**Store all service definitions in one repository**

| Name                         | Description                      | Activity | Sysplex  | Messages    | Last Modified (GMT)     | Modified By |
|------------------------------|----------------------------------|----------|----------|-------------|-------------------------|-------------|
| R12RGRP2                     | D10.WLM.ZOSMF.POLICY.R12RGRF     |          |          | Error       | Feb 23 2010 2:30:24 PM  | bmor        |
| RTDST3                       | Copy of RTDST3                   |          |          |             | Mar 21 2001 8:23:19 PM  | bmali       |
| RTDST3                       | SDS1 copy 5                      |          |          | Warning     | Jan 31 2010 10:49:38 PM | wirag       |
| SampdeF                      | Sample WLM Service Definition 62 |          |          |             | Sep 24 2007 8:48:22 AM  | tblatt      |
| SampdeF (Installed & Active) | Sample WLM Service Definition 57 |          | WLMMPLEX | Information | Feb 1 2010 8:52:56 PM   | wirag       |
| SPMinTst                     |                                  |          |          | Information | Jan 26 2010 3:50:46 PM  | wirag       |
| T13DEC07                     | add/remove SAP DB2s              |          |          |             | Dec 13 2007 9:01:59 PM  | ks50551     |
| TEST15                       |                                  |          |          | Information | Jan 12 2010 12:43:29 PM | wirag       |
| TESTFIX1                     |                                  |          |          |             | Oct 3 2006 11:40:35 AM  | sig011      |
| TESTSD1                      |                                  |          |          | Error       | Dec 30 2009 6:42:37 PM  | wirag       |
| WLM_BOF1                     | Large                            |          |          | Error       | Feb 19 2010 5:12:06 PM  | debug22     |
| WLM_BOF2                     |                                  |          |          |             |                         |             |
| WLM_DESC                     | WL De                            |          |          |             |                         |             |
| WLM001                       | Service                          |          |          |             |                         |             |
| WLM000                       |                                  |          |          |             |                         |             |
| WLM700                       |                                  |          |          |             |                         |             |
| wlmpol01                     | policy                           |          |          |             |                         |             |
| WLMPOL03                     |                                  |          |          | Warning     | Jan 13 2010 9:19:00 AM  | wirag       |
| WLMPOL04                     |                                  |          |          | Warning     | Feb 2 2010 12:09:54 AM  | wirag       |
| WLMSTT                       | AVT R10+R11RAS                   |          |          |             | Jul 8 2008 10:38:57 AM  | bmor        |
| WSCVLMDE                     | WSC Sample WLMServiceDefinition  |          |          | Error       | Jan 27 2010 4:05:01 AM  | p3asru      |

**Click to view, edit, print, install a service definition**

# z/OSMF Workload Management Editing Service Definitions



**SHARE**

Technology • Connections • Results

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

The screenshot displays the IBM z/OS Management Facility - Workload Management interface. The left sidebar shows navigation links: Welcome, Links, Performance, Workload Management (selected), and z/OSMF Administration. The main area is titled 'Workload Management' and contains tabs for Overview, Service Definitions (selected), Modify Sampdef, and Modify basepol. A 'Best-practice hints help to optimize service definitions' box points to a warning icon in the 'Velocity' column of the 'Service Classes' table. The table lists various service classes (STTCL3 through STTCL8) with columns for Name, Period, Importance, Duration, Goal Type, Response Time, Percentile, Velocity, CPU Critical Filter, and Resource Group. A context menu is open over the 'STTCL7' row, showing options like 'New Period', 'Expand', 'Collapse', 'Cut to Clipboard', 'Copy to Clipboard', 'Paste Periods', 'Move Periods', 'Delete...', and 'View Cross References'. A 'Click to copy element on clipboard for insertion into another service definition' box points to the 'Copy to Clipboard' option. Another box, 'Click to check where the element is used', points to the 'View Cross References' option. The bottom of the interface shows 'Total: 242, Selected: 1' and buttons for 'Reapply Filter and Sort', 'OK', 'Apply', 'Reset', and 'Cancel'.

# **z/OSMF Workload Management**

## **The new WLM Control Center in z/OSMF V1.12**



A complete overview is presented in session

*Manage your Workloads and Performance with z/OSMF*

*Friday, 11:00 AM-12:00 PM*

# Agenda




- Enclave Enhancements
  - Enclave Server Management
  - Work-Dependent Enclaves
- WLM Management
  - LDAP Support
  - Resource Group Enhancements
  - Do not always honor Skip Clock in Policy Adjustment
- WLM Reporting
  - Extend Number of Report Classes
  - Response Time Distribution for Velocity Goals
- Externalized IEAOPT Information
- Hyperdispatch APAR
- WLM support for IBM zEnterprise 196
- z/OSMF Workload Management
- WLM support for zManager
- WLM Tools Overview





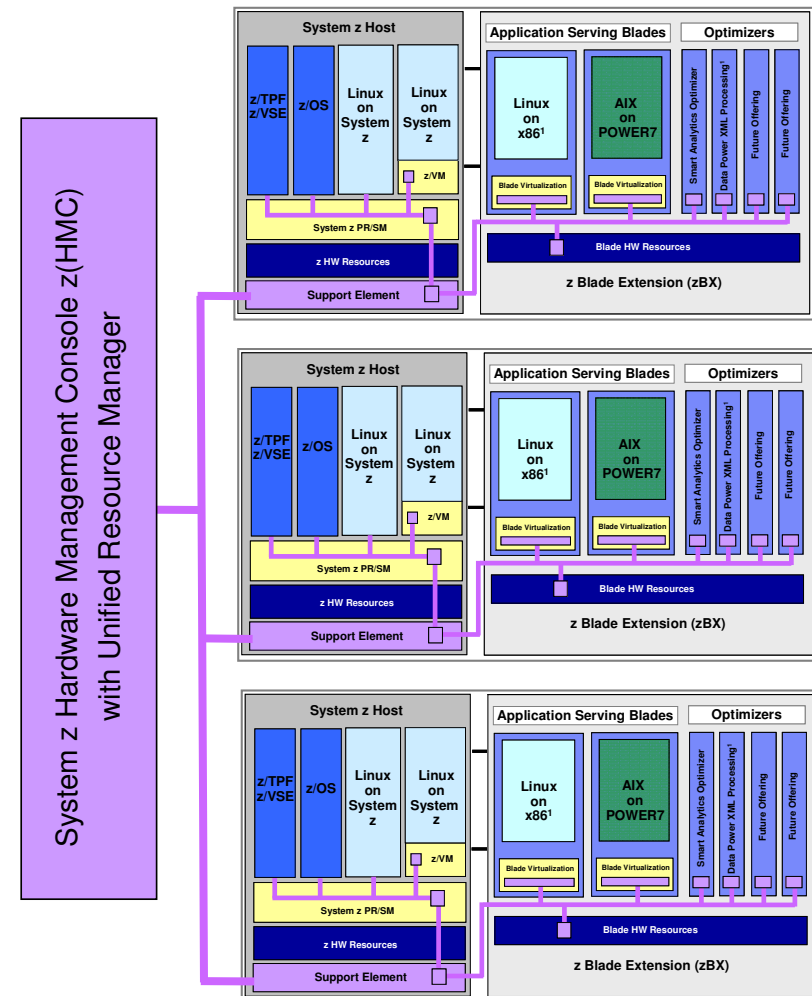
# zEnterprise Ensembles



**SHARE**


Technology • Connections • Results

- Ensemble
  - A zEnterprise Ensemble is a collection of zEnterprise Nodes managed as a single virtualized pool of server resources
    - Native LPAR and z/VM Virtual Images
    - Power VM Virtual images
    - IBM Smart Analytics Optimizer for DB2
  - A zEnterprise Node can be a member of at most one Ensemble
- zEnterprise Unified Resource Manager
  - allows for the management and optimization of a zEnterprise Ensemble as a single resource pool
  - System z Hardware Management Console (HMC) is management console
  - Ensemble-wide scope of responsibility



**SHARE**  
in Anaheim  
2011

# zEnterprise Platform Performance Manager



**SHARE**

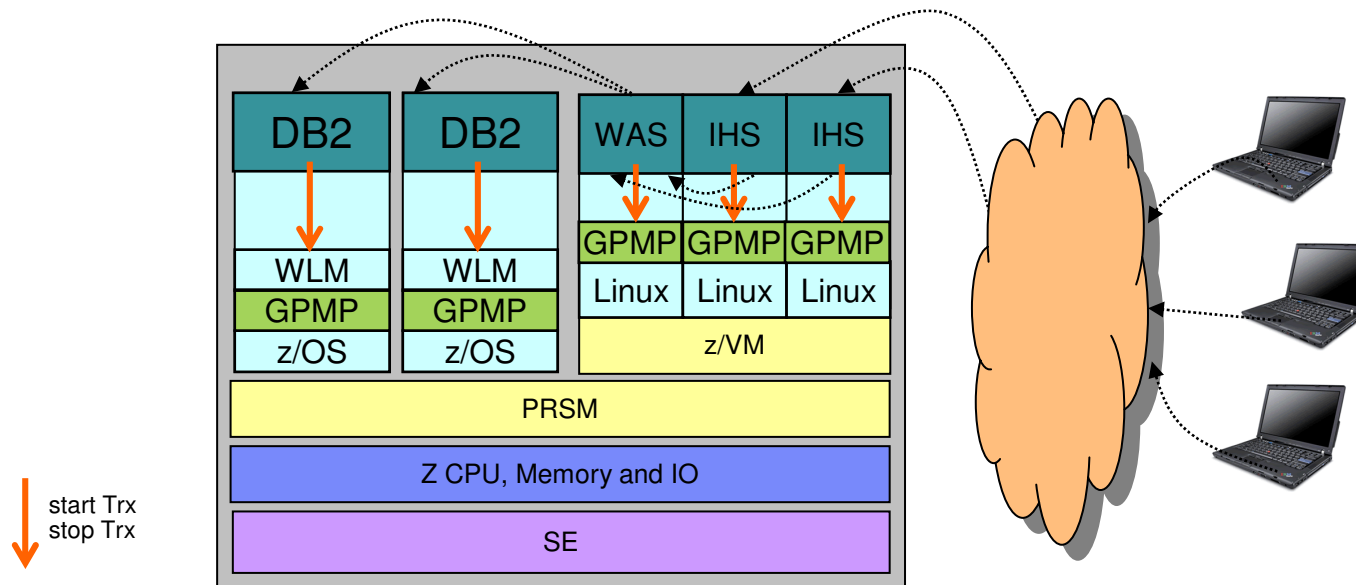
Technology • Connections • Results

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

**SHARE**  
in Anaheim  
2011


# zEnterprise Platform Performance Manager

## Resource management based on understanding of overall workload flow




- Applications / middleware has to be instrumented with ARM – Application Response Measurement (Open Group Standard) to collect transaction statistics
  - Enables to monitor the flow of transactions
  - Enables to monitor transaction response times and processing statistics
- OS Agent – guest platform management provider (GPMP)
  - is required to identify individual units of work
  - collects data about processes / address spaces and transactions
  - passes data to Platform Performance Manager
  - On z/OS the data is collected by WLM

# zManager Platform Workload Definition




**SHARE**  
Technology • Connections • Results

- A Platform Workload is a grouping mechanism and “management view” of virtual servers and optimizers supporting a business application
- Provides the context within which associated platform resources are presented, monitored, reported, and managed
- Management policies are associated to Platform Workload
  - Currently supports Performance Policy



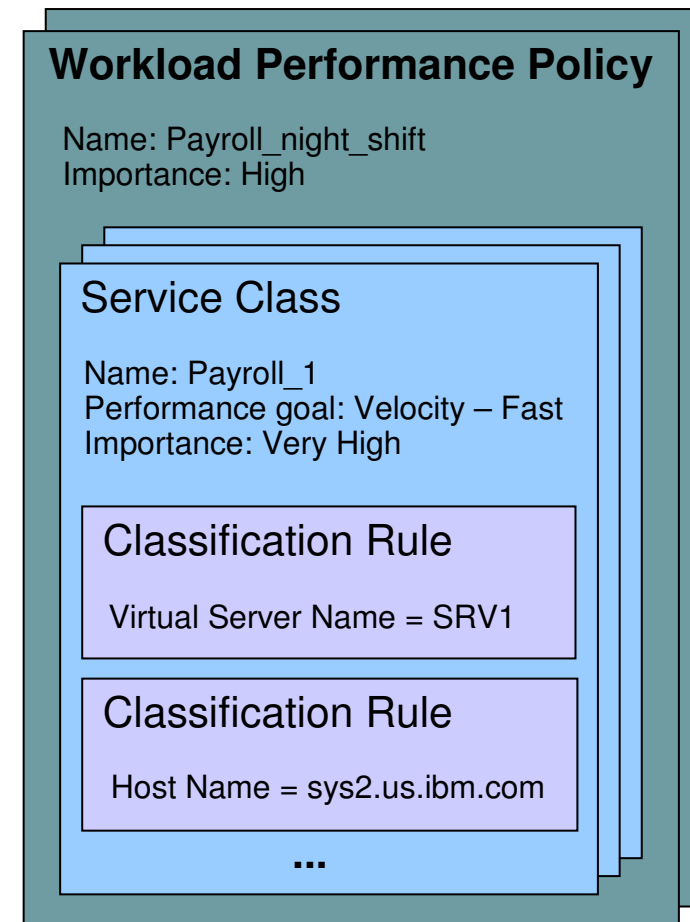
**SHARE**  
in Anaheim  
2011

# zManager Workload Performance Policy




**SHARE**  
Technology • Connections • Results

- Defines performance goals for virtual servers in a workload
  - Conceptually similar to simplified z/OS WLM Policy
- Provides basis for monitoring and management of platform resources used by virtual servers in a Workload
- Workload to performance policy relationship:
  - A Workload can have multiple performance policies associated with it
  - Single policy is active at a given time
  - Can dynamically change the policy that is active
    - Through the UI
    - Through a timed based schedule
      - *Example: Day shift policy / night shift policy*



**SHARE**  
in Anaheim  
2011

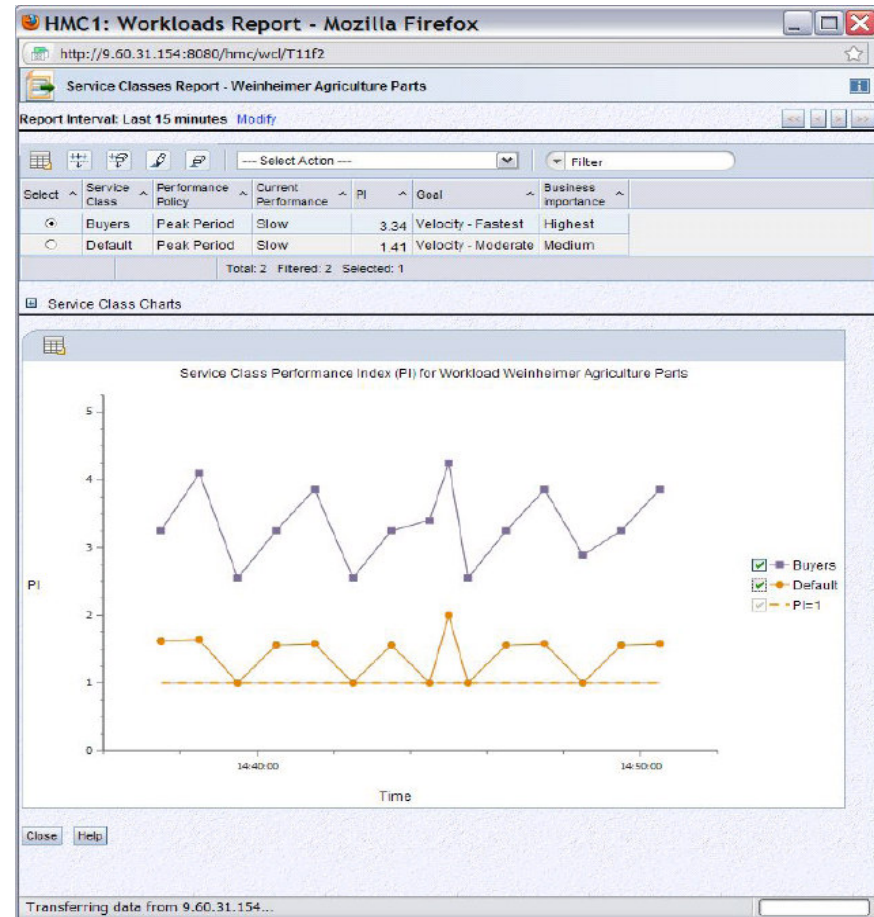
# zManager Workload Based Monitoring and Reporting



**SHARE**

Technology • Connections • Results


- Provide reporting capability that shows usage of platform resources in a Workload context with a zEnterprise Ensemble scope
  - Across virtual servers and optimizers supporting the Workload
- Workload goal vs. actual reporting
- Drill down from overall Workload “performance health” view to contributions of individual virtual server / optimizers
- Graphical views
  - Topology, trending graphs, etc.
- Links to system activity displays to show hardware utilization views
- Reporting limited to platform level resources, not trying to replace tools that report on intra-OS resources and performance



**SHARE**  
in Anaheim  
2011

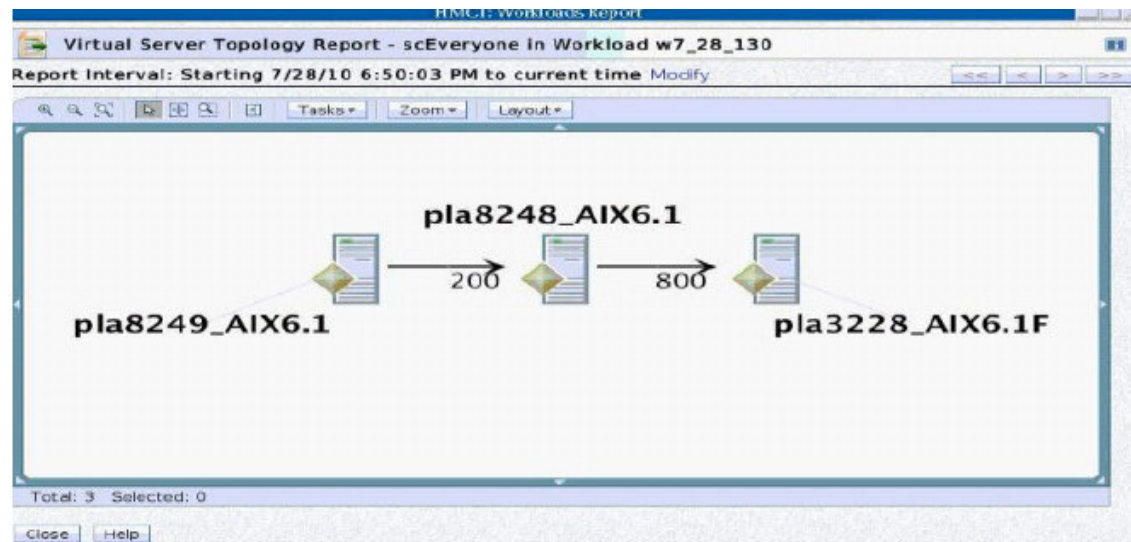


# zManager – Transaction Topology and Hops Report



**SHARE**  
Technology • Connections • Results

- Topology of virtual servers
- Transaction statistics




Hops Report - scEveryone in Workload w7\_28\_130  
Report Interval: Starting 7/28/10 6:50:03 PM to current time Modify

**Details for scEveryone**

Workload: w7\_28\_130 Performance policy: p1  
Performance goal: Velocity - Moderate Business Importance: Medium  
PI: 0.40 Performance: Fastest

| Name                         | Hop Number | Group Name      | Successful Transactions | Failed Transactions | Stopped Transactions | Inflight Transactions | Queue Time (s) | Execution Time (s) | Successful Average Response Time (s) |
|------------------------------|------------|-----------------|-------------------------|---------------------|----------------------|-----------------------|----------------|--------------------|--------------------------------------|
| Hop 0                        | 0          |                 | 200                     | 0                   | 0                    | 0                     | 2              | 0.000              | 0.000                                |
| IBM DB2 Universal Database   | 0          | db2inst1        | 0                       | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| IBM Webserving Plugin        | 0          | IBM_HTTP_Server | 200                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| pla8249_AIX6.1               | 0          |                 | 200                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| WebSphere:APPLICATION_SERVER | 0          | server1         | 0                       | 0                   | 0                    | 0                     | 1              | 0.000              | 0.000                                |
| pla8248_AIX6.1               | 0          |                 | 0                       | 0                   | 0                    | 0                     | 1              | 0.000              | 0.000                                |
| HelloWorld                   | 0          | Examples        | 0                       | 0                   | 0                    | 0                     | 1              | 0.000              | 0.000                                |
| Hop 1                        | 1          |                 | 200                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| WebSphere:APPLICATION_SERVER | 1          | server1         | 200                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| pla8248_AIX6.1               | 1          |                 | 200                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| Hop 2                        | 2          |                 | 800                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |
| IBM DB2 Universal Database   | 2          | db2inst1        | 800                     | 0                   | 0                    | 0                     | 0              | 0.000              | 0.000                                |

# WLM support for zManager



**SHARE**

Technology • Connections • Results

- The *guest platform management provider* (GPMP) is the interface between the zManager and the z/OS Workload Manager
- GPMP
  - passes to WLM information about the platform wide performance goals of workloads in which the z/OS is participating
  - sends data provided by WLM to the HMC for platform performance monitoring
    - Server configuration and high level performance statistics collected on z/OS
    - Aggregated transaction response time and resource data for the ARM-instrumented applications
- WLM
  - supports GPMP configuration and management by new WLM service definition options, commands, and messages
  - manages the GPMP address space (start, stop, and restart)
  - displays GPMP status information
  - collects and aggregates performance measurements for GPMP

**SHARE**  
in Anaheim  
2011



# WLM support for zManager

## Service Definition Enhancements for GPMP



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

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

DEFAULTS:

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

# WLM support for zManager

## Service Definition Enhancements for GPMP



To configure GPMP

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

**After GPMP settings defined, functionality level changes to 25**

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

Definition data set . . . : none

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

Select one of the
following options. . . . . 11
1. Policies
2. Workloads
3. Resource Groups
4. Service Classes
5. Classification Groups
6. Classification Rules
7. Report Classes
8. Service Coefficients/Options
9. Application Environments
10. Scheduling Environments
11. Guest Platform Management Provider
```

**GPMP-Settings Notes Options Help**

Guest Platform Management Provider (GPMP) Settings

Command ==>

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

Names of systems to be excluded:

| SYS04 | SYS09 |
|-------|-------|
|       |       |
|       |       |
|       |       |
|       |       |
|       |       |
|       |       |
|       |       |
|       |       |
|       |       |

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

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

# WLM support for zManager GPMP Configuration and Management



- On **policy activation**
  - WLM checks whether the service definition has valid GPMP settings
  - If activate=yes and system name not specified on excluded-list, GPMP is started automatically
- Also, you can use the **MODIFY WLM** command
  - To start the GPMP on a system
  - To stop the GPMP on a system
  - Intended for recovery actions. Recommended is to manage GPMP through WLM
- Once you stopped the GPMP manually, the GPMP switches into “manual mode”. It is not automatically restarted even if a WLM policy with a valid GPMP configuration gets activated
  - Status maintained until next IPL

# WLM support for zManager

## GPMP related Commands



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

- **F WLM,GPMP,START**

- Indicates that you want to start the GPMP

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

- **F WLM,GPMP,STOP**

- Indicates that WLM stops the currently active GPMP instance

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

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

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

# WLM support for zManager GPMP related Commands



- DISPLAY WLM command extensions:

```
IWM025I 11.42.45 WLM DISPLAY 231
ACTIVE WORKLOAD MANAGEMENT SERVICE POLICY NAME: BASEPOL
ACTIVATED: 2010/02/18 AT: 12:57:55 BY: BMAI FROM: TRX2
DESCRIPTION: Base policy for system test
RELATED SERVICE DEFINITION NAME: Ralfpol
INSTALLED: 2010/02/18 AT: 12:57:48 BY: BMAI FROM: TRX2
WLM VERSION LEVEL: LEVEL025
WLM FUNCTIONALITY LEVEL: LEVEL025
WLM CDS FORMAT LEVEL: FORMAT 3
STRUCTURE SYSZWLM_WORKUNIT STATUS: CONNECTED
STRUCTURE SYSZWLM_EBAE2097 STATUS: CONNECTED
STATE OF GUEST PLATFORM MANAGEMENT PROVIDER (GPMP): ACTIVE
*SYSNAME* *MODE* *POLICY* *WORKLOAD MANAGEMENT STATUS*
TRX1 GOAL BASEPOL ACTIVE
TRX2 GOAL BASEPOL ACTIVE
*SYSNAME* *GPMP STATUS*
TRX1 INACTIVE
TRX2 ACTIVE
```

To display system and  
GPMP status information,  
enter:

**D WLM,SYSTEMS,GPMP**

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

To display whether ARM  
is enabled or disabled,  
enter:

**D WLM,AM**

# WLM support for zManager GPMP related Commands



- The existing **MODIFY WLM,AM=DISABLE|ENABLE** command is not changed, but the logic for DISABLE/ENABLE changed in the following way:
  - Disabling ARM (Application Response Measurement) will terminate a running GPMP
  - Manually starting the GPMP (using the MODIFY WLM,GPMP,START command) when ARM is disabled will result in message IWM078I
  - Activating a WLM policy that contains valid GPMP settings will not result in the start of a GPMP instance, if ARM is disabled
  - The state of the GPMP will be displayed as “DISABLED”, if ARM is disabled
  - If ARM is enabled again, the state of the GPMP will change to “STOPPED”. To start the GPMP again, it has to be started manually


# WLM support for zManager

## Prerequisites



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

# WLM Tools: A Summary



**SHARE**

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


<http://www-03.ibm.com/servers/eserver/zseries/zos/wlm/tools/>

**SHARE**  
in Anaheim  
2011



# WLM Tools

## Service Definition Editor



**SHARE**

Technology • Connections • Results

WLM Service Definition Editor - D:\SAMPLESD.xml

File Edit Options Help

local

| Classification Groups |                |        | Classifications        |       | Service Parameter |              | Application Environments |       |               | Scheduling Environments |                             | Report Classes |  |
|-----------------------|----------------|--------|------------------------|-------|-------------------|--------------|--------------------------|-------|---------------|-------------------------|-----------------------------|----------------|--|
| Service Definition    |                |        | Resource Groups        |       | Workloads         |              | Service Policies         |       |               | Report Classes          |                             |                |  |
| Name                  | ServiceClasses | Period | Goal                   | Im... | Duration          | ResponseTime | Perce...                 | Level | ResourceGr... | CPU...                  | Description                 |                |  |
| WKLDASC               |                |        |                        |       |                   |              |                          |       |               |                         | ALL APPC Transaction        |                |  |
| WKLDASC               | A3V30STD       |        |                        |       |                   |              |                          |       | GBATCH20      | No                      | ASCH default Service Class  |                |  |
| WKLDASC               | A3V30STD       | 1      | Velocity               | 2     | 500               |              |                          | 10    |               |                         |                             |                |  |
| WKLDASC               | A3V30STD       | 2      | Velocity               | 2     | -                 |              |                          | 8     |               |                         |                             |                |  |
| WKLTJK                |                |        |                        |       |                   |              |                          |       |               |                         | All Batch Jobs              |                |  |
| WKLTJK                | B4V10STD       |        |                        |       |                   |              |                          |       | -             | No                      | Batch Standard VEL 10 IMP 4 |                |  |
| WKLTJK                | B4V10STD       |        | Velocity               | 4     | -                 |              |                          | 10    |               |                         |                             |                |  |
| WKLTJK                | B4V20STD       |        |                        |       |                   |              |                          |       | -             | No                      | Batch Standard VEL 20 IMP 4 |                |  |
| WKLTJK                | B4V20STD       | 1      | Velocity               | 4     | 1000              |              |                          | 20    |               |                         |                             |                |  |
| WKLTJK                | B4V20STD       | 2      | Velocity               |       | -                 |              |                          | 10    |               |                         |                             |                |  |
| WKLDTSO               |                |        |                        |       |                   |              |                          |       |               |                         | ALL TSO USERIDS             |                |  |
| WKLDTSO               | T2335DEV       |        |                        |       |                   |              |                          |       | -             | No                      | 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             | 00:00:02.000 | 99                       |       |               |                         |                             |                |  |

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 "WKLTJK"/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

**SHARE**  
in Anaheim  
2011



# WLM Tools

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



**SHARE**

Technology • Connections • Results


```

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

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

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

## WLMOPT – WLM Application Environment Viewer



## SHARE

Technology • Connections • Results

```

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

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

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

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

```

# SHARE in Anaheim 2011