

Managing Oracle Workload with z/OS Workload Manager

G. Tom Russell
IBM Canada Ltd.
Tom_Russell@ca.ibm.com

MVS Oracle SIG
April 13, 2005
Redwood Shores, CA

IBM @server. For the next generation of e-business.

Agenda


- Brief Introduction to MVS Workload Management (WLM) Externals
- Classifying Different Oracle Workloads
 - TSO
 - Batch
 - CICS or IMS
 - Transparent Gateway for DB2
- What is an Enclave?
- Handling of Enclaves in Oracle Net
- Some Examples of Oracle Workload Classification
 - Oracle Apps Benchmark
 - Classify Oracle Net client work
- CPU Accounting considerations
 - Enclave CPU reporting

IBM @server. For the next generation of e-business.

Controls Managed by WLM

In goal mode, SRM uses the following controls to manage work to help meet customer specified goals.

- Dispatching Priority
- MPL Targets
- Swap Protect Time
- Storage Targets
 - Protective Storage Targets
 - Restrictive Storage Targets
- Storage Policies
 - Estor (protected, LRU, space available)
 - Swap, VIO, Hiperspace, Stolen pages
- Etc.



IBM @server. For the next generation of e-business.

Service Class and Classification

- WLM part of z/OS
- A “Work Manager” calls WLM to classify a work request
 - Ex: JES, CICS, Websphere, Oracle
- A Service Class is assigned based on information passed to Classify service
- Classification rules are in the workload definition in the WLM Policy
- Policy built with ISPF Dialog
 - EXEC ‘SYS1.SBLSCLI0(IWMARIN0)’

IBM @server. For the next generation of e-business.

Goals

- The Service Class has multiple Periods
 - Based on the amount of CPU used
- The SC Period has two attributes
 - Importance
 - 1-5 and discretionary
 - Performance Goal
 - Response, or percentile response
 - Velocity

IBM @server. For the next generation of e-business.

Velocity Goals

- A goal for long running, non response oriented work
- Velocity = $\text{Using samples} \div (\text{Using} + \text{Delay samples}) \times 100$
 - Includes Storage Delay
 - MPL
 - Paging
 - Includes CPU Delay
 - Higher priority work
 - Optionally includes I/O delay

IBM @server. For the next generation of e-business.

Response Goals

- A goal for interactive, response oriented work
- Average response time
 - Ex: Average response of transactions = .2 sec.
- Percentile
 - Ex: 95% of transactions with response less than .5 sec.

IBM @server. For the next generation of e-business.

Service Class Period

- After consuming a number of Service Units, the service class migrates to another period
 - Different goal and importance
 - Allows longer work to be treated at lower importance than new work

IBM @server. For the next generation of e-business.

CPU Service Unit

- Service Units per CPU second is a constant associated with the CPU model.
- Examples:
 - 9672-R56 delivers 5158 SU per CPU second
 - 2066-002 delivers 8588.3 SU per CPU second
 - 2084-301 delivers 21858 SU per CPU second
- WLM uses Service Units to make Policy independent of processor speed

IBM @server. For the next generation of e-business.

Resource Group

- Define a maximum or minimum service rate
 - CPU service units
- A set of service classes
- Usually used as a “cap” on resource
- Can be very effective as a minimum rate
 - Service class with discretionary goal, but guaranteed a minimum number of service units corresponding to 5% of system capacity

IBM @server. For the next generation of e-business.

Classifying TSO

- Usually classified by USER ID
- Recommend a Goal with Multiple Periods
- Oracle processing performed at priority of the TSO transaction
 - Response Time Goal for first and second period
 - Maintain high importance for short transactions
 - Velocity Goal for last period, or large response
 - Lower importance for long running transactions
- Example:
 - Period 1: Importance 1, Average response .1 second
 - Period 2: Importance 3, Average response .5 second
 - Period 3: Importance 5, Velocity 10 or
 - Period 3: Importance 5, Average response 10 seconds

IBM @server. For the next generation of e-business.

Classifying Batch Jobs

- Usually classified by Job Name or Job Class
- Recommend a Goal with a Single Period
- Oracle processing performed at priority of the Batch Job
 - Velocity Goal
 - Lower importance to prevent interference with interactive work
- Example:
 - Importance 5, Velocity 10
- Note different classes of work may require different service classes, with differing goals

IBM @server. For the next generation of e-business.

Classifying CICS Transactions

- Usually classified by CICS Transaction ID
- Oracle processing performed at priority of the CICS address space
 - If MRO, Classification done in TOR
 - Response Time Goal is recommended
 - Maintain high importance for short transactions
 - All transactions in same region should have similar response requirements
 - Must be single period
 - Velocity Goal for CICS address space
 - Transactions can be managed by this goal if desired
 - Similar to (obsolete) WLM Compatibility Mode
- Example:
 - Importance 2, Average response .2 second

IBM @server. For the next generation of e-business.

What Is an Enclave?

- A "business transaction" without address space boundaries
- Independent enclave used by Oracle
 - True SRM transaction
 - Separately classified and managed in service class or performance group
- LE uses the word Enclave for something completely different
 - the set of resources and processing associated with a single logical LE application
 - Defines the scope of the application. Owns storage, files, etc.
 - Not discussed here


IBM @server. For the next generation of e-business.

How Do MVS Enclaves Behave?

- 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

IBM @server. For the next generation of e-business.

NET Structure and CPU Recording



IBM @server. For the next generation of e-business.

Classifying NET Transactions

- All independent enclaves are classified using the active MVS WLM policy
- Classified using attributes associated with each subsystem
- Defaults if you do not classify in WLM policy:
 - enclaves default to the SYSOTHER service class which has a discretionary goal.
- Note that goal mode is required on all currently supported MVS releases

IBM @server. For the next generation of e-business.

WLM Support in Oracle Net

- Support in Oracle OSDI 8.1.7 or subsequent
 - Multi-Process Monitor (MPM) support ended December 2003
- Inbound client work gets WLM classification
 - Client IP/LU, target service, userID
- Client requests run under a preemptable enclave SRB
- Each enclave classified into a service class
 - Resource consumption and transaction rates in RMF workload report
 - Separately managed by WLM
 - Dispatching priority, storage, I/O priority

IBM @server. For the next generation of e-business.

Transparent Gateway for DB2 (TG4DB2)

- TG4DB2 in 8i was based on MPM
 - All work in TG4DB2 ran at priority of TG4DB2
- TG4DB2 in 9i is based on OSDI
 - Work from network comes through Oracle NET
 - Runs in an enclave until it gets to TG4DB2 A/S
 - TCB per session, uses Call Attach Facility (CAF)
 - You *must* classify this enclave
 - Velocity goal with high importance is appropriate

IBM @server. For the next generation of e-business.

TG4DB2 ...

- With latest OSDI patches the work in TG4DB2 TCB is joined to the enclave
 - Some of DB2 code runs in the enclave
 - Background DB2 processes charged to DB2
 - Response time goals possible
- In 10g, CAF removed, RRSAF used to connect to DB2
 - **All** DB2 work runs in the enclave
 - Response time goals are recommended

IBM @server. For the next generation of e-business.

Code for Database Available in MetaLink

- Ensure you are at current Oracle OSDI maintenance level
 - 8.1.7.4.50, 9.2.0.5.21, or 10.1.0.3.0 plus the latest OSDI cumulative patches
 - Check with Oracle support
 - Notes 237007.1, and 290738.1 in Metalink
 - Critical patch update scheduled for April 12, 2005
- Must migrate from MPM

IBM @server. For the next generation of e-business.

WLM Support in Oracle Net...

- Define Net Service

```
DEFINE SERVICE ORANETW TYPE(NET)
PROC(ORANET) -
DESC('Oracle Network Supporting WLM
Transactions') SID(NETW) -
PARM('HPNS PORT(1521) ENCLAVE(CALL)')
```

- ENCLAVE(CALL)
 - dynamic enclaves, many WLM transactions
- ENCLAVE(SESS)
 - static enclaves (default), one WLM transaction per session

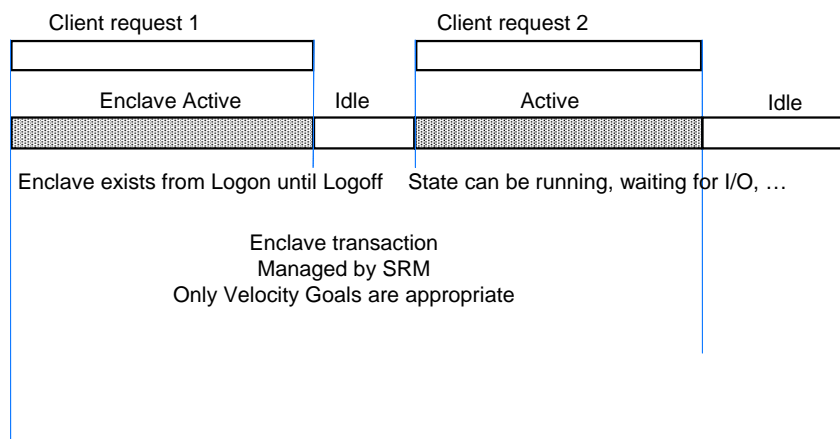
IBM @server. For the next generation of e-business.

WLM Support in Oracle Net...

- Enclave(Sess)
 - Classification done once at Logon
 - Enclave deleted at Logoff
 - Entire session is a single WLM transaction
 - Only Velocity Goals are appropriate
- Enclave(Call)
 - Classification done every time a request arrives from client
 - Enclave deleted when NET has to wait for next request
 - Each client request is a separate WLM transaction
 - Response Time or Percentile Goals should be used

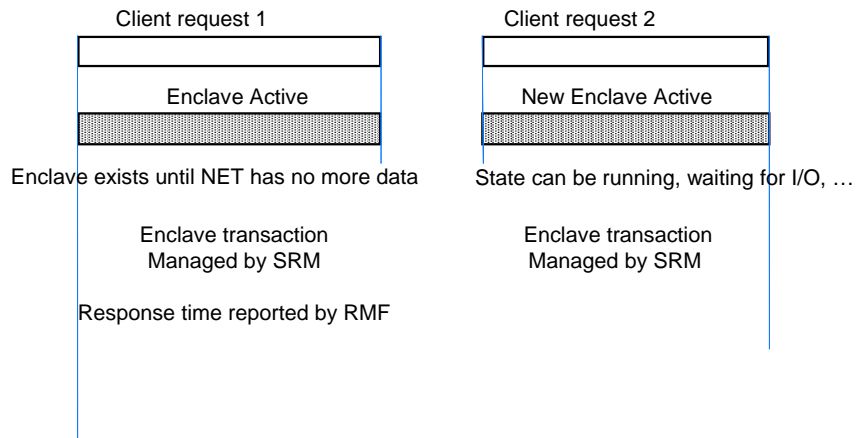
IBM @server. For the next generation of e-business.

NET Behaviour with Enclave(sess)



IBM @server. For the next generation of e-business.

NET Behaviour with Enclave(call)



IBM @server. For the next generation of e-business.


New Management Capabilities

- Recommend ENCLAVE(CALL)
- Establish Response Goal at high importance for first period
- Migrate to less importance for second period
- Third Period for very low importance
 - This allows you to maintain high response for trivial work
 - Treat heavier, less sociable work at an appropriate priority
 - Especially if running in capped Logical Partition or capped resource group

IBM @server. For the next generation of e-business.

Both Modes Can Co-exist


- Separate Net can be used for greater flexibility
- Segregate users based upon workload and performance needs
- Distinguish by IP address or Port number



IBM @server. For the next generation of e-business.

Server Consolidation

- Work can be assigned different Service Classes
 - Multiple Oracle instances handled based on goals for service classes assigned by NET, TSO, JES, ...
 - Single copy of Oracle Subsystem
 - Single copy of NET possible
 - Can safely consolidate multiple independent Oracle instances onto a single MVS image




IBM @server. For the next generation of e-business.

Classification of Oracle NET Enclaves

Attribute	Value
SI	OSDI subsystem name
UI	User ID from the client. For Oracle Applications this is the UID of the user running the application server on the middle-tier processor
NET	if SNA: client Network Name from VTAM if TCP: First eight characters of dotted IP address. (ex.100.024.)
LU	if SNA: The client LU name. if TCP: Last eight characters of dotted IP address. Note that the IP address requires leading zeros to be specified.
CT	Protocol from connect, TCP or LU6.2
SPM	Position 1 to 8. Oracle Service Name for this connection. The service name is defined in the parameters used to initialize the Oracle OSDI subsystem.
SPM	Position 9 to 89. TCP/IP hostname (left justified)

IBM @server. For the next generation of e-business.

An Example: Oracle Applications Benchmark



IBM @server. For the next generation of e-business.

Classification Rules for ORANET

```

Subsystem-Type Xref Notes Options Help
-----
          Modify Rules for the Subsystem Type          Row 1 to 5 of 5
Command ==>> _____ SCROLL ==>> PAGE

Subsystem Type . : OSDI          Fold qualifier names?  Y (Y or N)
Description . . . ORACLE Subsystem

Action codes:  A=After      C=Copy          M=Move      I=Insert rule
                B=Before    D=Delete row  R=Repeat    IS=Insert Sub-rule
                More ==>>

-----Qualifier-----
Action  Type      Name      Start      Service      Report
-----
_____ 1  SI        ORAC      _____  _____
_____ 2  NET      010.100.  _____  _____
_____ 3  LU        001.080  _____  ORAMT1
_____ 3  LU        002.082  _____  ORAMT2
_____ 3  LU        003.081  _____  ORAMT3
***** BOTTOM OF DATA *****

```

IBM @server. For the next generation of e-business.

Service Class Goal for Important Work

```

. Service-Class Xref Notes Options Help
-----
          Modify a Service Class          Row 1 to 2 of 2
Command ==>> _____

Service Class Name . . . . . : ORAMT1
Description . . . . . Oracle Mid Tier #1
Workload Name . . . . . ORACLE (name or ?)
Base Resource Group . . . . . _____ (name or ?)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- -----Goal-----
Action # Duration Imp. Description
-----
_____ 1 50 1 Average response time of 00:00:00.015
_____ 2 1000 3 Average response time of 00:00:00.500
_____ 3 5 5 Execution velocity of 10
***** Bottom of data *****

```

IBM @server. For the next generation of e-business.

Goal for Concurrent Manager Work

```
. Service-Class Xref Notes Options Help
-----
          Modify a Service Class                Row 1 to 2 of 2
Command ==>> _____

Service Class Name . . . . . : ORAMT2
Description . . . . . Oracle Mid Tier #2
Workload Name . . . . . ORACLE (name or ?)
Base Resource Group . . . . . _____ (name or ?)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

      ---Period--- -----Goal-----
Action # Duration Imp. Description
-----
   1          Discretionary
***** Bottom of data *****
```

IBM @server. For the next generation of e-business.

Effect of Implementing Enclave(call)

- RMF III Enclave report only shows “active” enclaves
 - Many fewer in display
 - Use the RMFPP reports for the service class data
- RMF now has transaction rate, and related stats
 - “Transaction” is a network interaction with new code
 - Better definitions will come in later releases of Oracle
 - This code available today
- Slight increase in NET TCB time

IBM @server. For the next generation of e-business.

RMF Monitor III Enclave Report

```

RMF 2.10.0 Enclave Report                               Line 1 of 2
Command ==>>>                                         Scroll ==>>> CSR

Samples: 100      System: MVS4  Date: 02/20/02  Time: 17.18.20  Range: 100  Sec

Current options:  Subsystem Type: ALL                  -- CPU Util --
Enclave Owner:                                     Appl%   EAppl%
Class/Group:                                       13.4    69.0

Enclave  Attribute  CLS/GRP  P Goal  % D X  EAppl%  TCPU   USG  DLY  IDL
*SUMMARY
ENCO0001          ORAMT1  3                28.78
                28.78  71.79    49  41  0.0
    
```

IBM @server. For the next generation of e-business.

RMF Monitor III Sysplex Summary

```

RMF 2.10.0 Sysplex Summary - R26PLEX                   Line 5 of 33
Command ==>>>                                         Scroll ==>>> PAGE

WLM Samples: 400      Systems: 1  Date: 06/10/02  Time: 13.25.00  Range: 100  Sec

>>>>>>>XXXXXXXXXXXXXXXXXXXX<<<<<<<<

Service Definition: PBCORAC                            Installed at: 06/07/02, 10.16.52
Active Policy: ORAPOL                                  Activated at: 06/07/02, 10.17.00

----- Goals versus Actuals -----   Trans --Avg. Resp. Time-
Exec Vel  --- Response Time ---   Perf  Ended  WAIT EXECUT ACTUAL
Name      T  I  Goal Act  ---Goal--- --Actual--  Indx  Rate  Time  Time  Time
ORACLE    W           46                               237.1 0.000 0.031 0.031
ORACLE    S           88                               235.1 0.000 0.013 0.013
          1  1       98 0.015 AVG 0.005  AVG 0.30 232.6 0.000 0.004 0.005
          2  2       95 0.500 AVG 0.190  AVG 0.38 2.020 0.000 0.190 0.190
          3  3       90 84                               1.07 0.490 0.000 3.403 3.403
ORACTC    S  5       20 26                               0.76 1.940 0.000 2.142 2.142
    
```

IBM @server. For the next generation of e-business.

CPU Accounting for Oracle Work

- CPU time for user work through Oracle Net now accumulated separately
 - May require changes to customer accounting and capacity planning methodologies
 - CPU now charged to NET not DB Server or Gateway
 - Operations staff may need education
 - SDSF shows almost no CPU in Server or Gateway A/S
 - Lots in ORANET
- Overhead of Enclave(Call) in NET TCB time

IBM @server. For the next generation of e-business.

Example SDSF display

Display Filter View Print Options Help										
SDSF	DA	OU02	OU02	PAG	0	SIO	2	CPU	1	LINE 1-26 (26)
NP	JOBNAME	STEPNAME	SRVCLASS	DP	CPU%	ECP%	CPU-TIME	ECP-TIME	REAL	
	ORACNET	ORACNET	SYSSTC	FE	0.00	0.00	2001.54	45261.24	7218	
	ORAC015	ORAC015	SYSSTC	FE	0.00	0.00	3.06	3.06	4381	
	ORAC025	ORAC025	SYSSTC	FE	0.00	0.00	3.33	3.33	4402	
	ORAC024	ORAC024	SYSSTC	FE	0.00	0.00	3.83	3.83	4246	
	ORAC019	ORAC019	SYSSTC	FE	0.00	0.00	5.26	5.26	4022	
									
	ORAC007	ORAC007	SYSSTC	FE	0.00	0.00	3.34	3.34	3978	
	ORAC006	ORAC006	SYSSTC	FE	0.00	0.00	3.09	3.09	3984	
	ORAC005	ORAC005	SYSSTC	FE	0.00	0.00	3.76	3.76	3888	
	ORAC004	ORAC004	SYSSTC	FE	0.00	0.00	11.07	11.07	4058	
	ORAC003	ORAC003	SYSSTC	FE	0.00	0.00	3.09	3.09	4072	
	ORAC002	ORAC002	SYSSTC	FE	0.00	0.00	3.46	3.46	4048	
	ORAC001	ORAC001	SYSSTC	FE	0.00	0.00	547.68	547.68	3802	

IBM @server. For the next generation of e-business.

Comments and Recommendations

- Goal Mode required since z/OS 1.3
 - Enclave(call) with velocity goals will not hurt
 - Velocity still appropriate, but not optimal
 - Enclave(sess) with response time goals not appropriate
 - Enclave goes to last period shortly after Logon
- Must have subsystem OSDI defined, with a default service class specified
 - If no service class assigned by policy then SYSOTHER is used
 - SYSOTHER has a discretionary goal. *Very Bad*
- Recommend ENCLAVE(CALL) in Oracle Net Service
 - Establish a short 1st period importance 1 to maintain response for trivial requests
 - Lower importance for 2nd period
 - Importance 5 or discretionary 3rd period
- Consider setting PARALLEL_MAX_SERVERS = 1
 - Parallel query processing still in DB service address space

IBM @server. For the next generation of e-business.

Summary

- OSDI WLM support exploits z/OS function not available to other Oracle versions
- Consolidation of multiple smaller Oracle instances on single S/390 now possible
 - Either multiple or single instance of Oracle on MVS
- Transparent Gateway product allows coexistence with DB2, with new support for response time goals
- Each client's transactions can be separately managed
 - The most important work gets the resources
 - Unsociable work can be segregated
- Resource group can be used to guarantee minimum (or maximum) service
- New response time and transaction rate recording in RMF

IBM @server. For the next generation of e-business.

Thank you – Any Questions?

Tom_Russell@ca.ibm.com

IBM server. For the next generation of e-business.