

WLM and Enclaves for DB Servers DB2 and Oracle

**G. Tom Russell
IBM Canada Ltd.**

Tom_Russell@ca.ibm.com
CMG Canada
November 14, 2002
Toronto, Ontario

Thanks To:

- John Arwe
- Peter Yocom
- Gail Whistance

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

Copyrights and Trademarks

- (c) Copyright IBM Corporation 2002
- Database 2, MVS/ESA, OS/390, and RACF are trademarks of the International Business Machines Corporation.
- IBM and DB2 are registered trademarks of the International Business Machines Corporation.
- Oracle and OSDI are trademarks of Oracle Corporation.
- All other products may be trademarks or registered trademarks of their respective companies.

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

Agenda

- What is an enclave?
- What problems are being solved?
 - DB2 distributed requests (DDF)
 - Oracle distributed requests (Oracle NET)
- Some Examples of 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.

What is an enclave?

- A "business transaction" without address space boundaries
 - Two types
 - Exist in both goal and compatibility mode
- Independent enclave
 - True SRM transaction
 - Separately classified and managed in service class or performance group
- Dependent enclaves
 - Logical extension of an existing address space transaction
 - Inherits srvcass/pgn from its owner's address space

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

How do 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

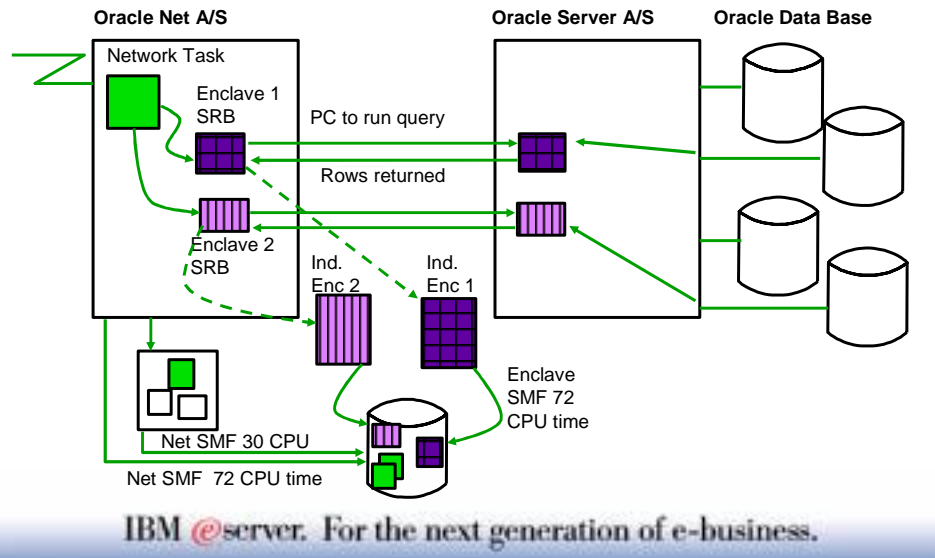
IBM @server. For the next generation of e-business.

Oracle Net and DDF

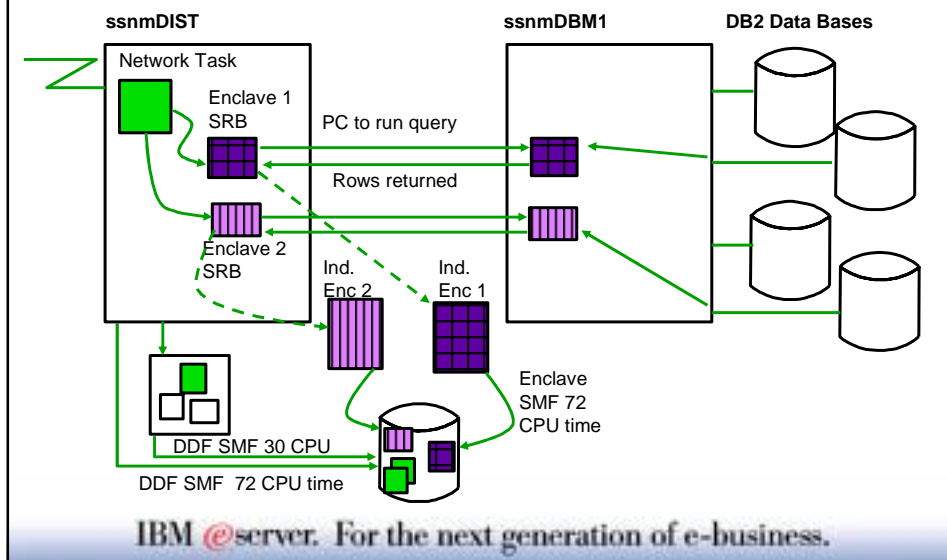
- | | |
|---|--|
| <ul style="list-style-type: none">• Problems<ul style="list-style-type: none">– No way to manage individual transactions. Everything runs with the address space controls of the ssnmDIST or NET address space controls– Heavy loads can prevent entry of new requests | <ul style="list-style-type: none">• Solution<ul style="list-style-type: none">– Make the transactions visible to SRM as independent enclaves<ul style="list-style-type: none">• Individually classified, managed, and reported• Subject to period switch– Run existing work at a dprty lower than the network atsk that creates new work |
|---|--|

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

Oracle NET Structure



DDF Structure



Independent Enclave Accounting

- Each enclave represents an individual SRM transaction
- Transaction counts and resource usage are recorded in the SMF 72 record for the enclave's service class and report class
- No SMF 30 records or equivalent for enclaves themselves. Transaction counts and resource usage are recorded in the SMF 30 of the owning address space.
- CPU time consumed by enclaves is recorded in the SMF 89 records of home address space for usage-based pricing.

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

Classifying Transactions

- All independent enclaves, are classified using the active MVS WLM policy in both compatibility and goal modes
- In compat mode the ICS can be used to map a service class to a performance group, and likewise for reporting
- Can classify using new attributes for each subsystem
- Defaults if you do not classify in WLM policy:
 - Goal mode: enclaves default to the SYSOTHER service class which has a discretionary goal!
 - Compat mode: any enclaves run in the pgn/rpgn of the owning address space as they did before

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

Managing Independent Enclaves

- Transactions are subject to period switch
- Goal mode:
 - All goal types allowed
- Compat mode:
 - Performance group and report performance group can be assigned using SRVCLASS=xxx in the IEAICSxx parmlib member
 - Limit of one report performance group
 - Time slicing, domain is ignored

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

WLM Support in Oracle OSDI 8.1.7 Net

- Define Net Service
 - DEFINE SERVICE ORANETW TYPE(NET)
PROC(ORANET) -
 - DESC('Oracle Network Supporting WLM') -
 - SID(NETW) -
 - PARM('HPNS PORT(1521) ENCLAVE(CALL)')
- ENCLAVE(CALL)
 - new behaviour, dynamic enclaves
- ENCLAVE(SESS)
 - old behaviour, static enclaves (default)

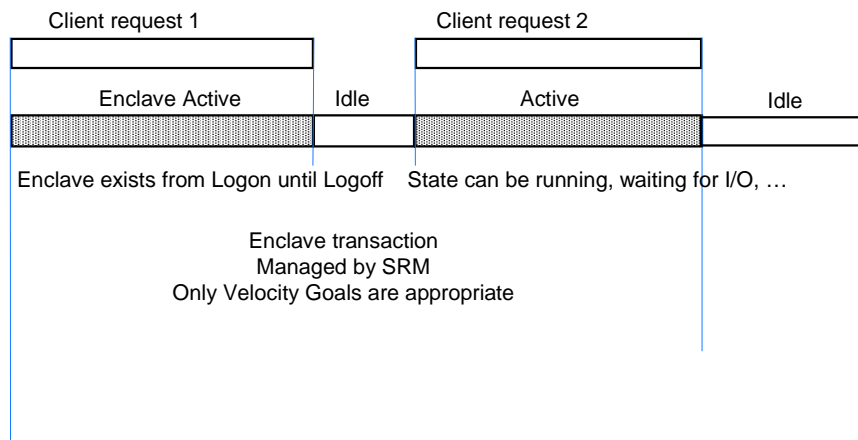
IBM @server. For the next generation of e-business.

WLM Support in Oracle OSDI 8.1.7 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

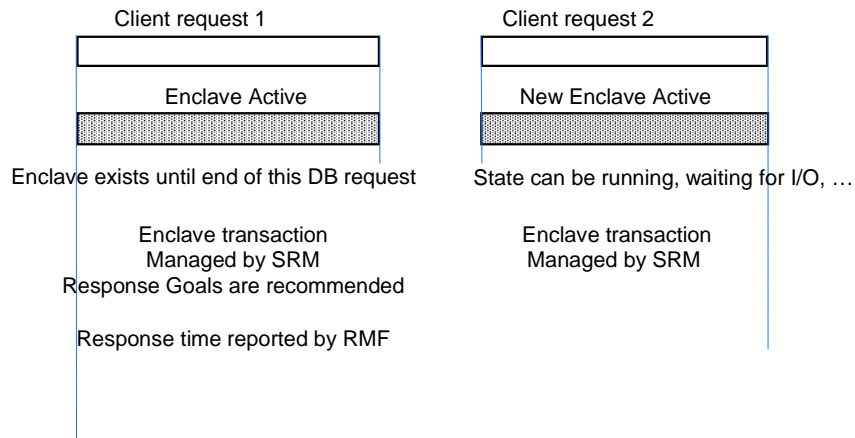
IBM @server. For the next generation of e-business.

NET Behaviour with Enclave(sess) DDF Behaviour with Threads=Active



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

NET Behaviour with Enclave(call) DDF Behaviour with Threads=Inactive



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

What Goals Should I Use?

- THREADS=Inactive and RELEASE(Commit) for DDF
- ENCLAVE(Call) for Oracle Net
 - One enclave per active interval
 - Response times do not include think time
 - Response time goals and multiple periods can be used
- THREADS=Active or RELEASE(Deallocate) for DDF
- ENCLAVE(Sess) for Oracle Net
 - One one enclave for the life of the thread/session
 - Enclave response time includes think time
 - Response time goals should not be used
 - Multiple periods should not be used

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

New Management Capabilities

- Recommend ENCLAVE(Call) or Threads=Inactive
- 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

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

CPU Service has Changed

- Philosophical refinement, compatible with the existing definition
- CPU time/service are no longer synonymous with "task time/service"
- CPU time/service now includes contributions from all preemptible dispatchable units: tasks, client SRBs, enclave SRBs
- SRB time/service is unchanged – local and global SRBs

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

SMF Type 30 Record

- No changes are required to existing accounting packages
- Existing CPU time and service fields include enclave contributions
- For all enclaves created by DDF or Oracle NET, new fields contain:
 - enclave active time
 - transaction count
 - CPU time and service
 - I/O times and I/O count

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

SMF Type 72 Record

- No changes are required to existing accounting packages
- Existing CPU service fields include enclave contributions
- Other existing fields include data from enclave transactions:
 - Active time
 - Response time
 - Transaction count
 - Number of address spaces (now: plus enclaves) sampled
 - I/O times and I/O count
- Resources consumed by the DDF or NET transactions are no longer reported with DDF or NET

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

Enclave System Effects

- (Lots) more transactions. Each DDF or NET transaction used to be part of the DDF or NET address space SRM transaction; now SRM sees each transaction.
- Increased active time since more transactions exist
- Large decrease in SRB time/service in SMF 30 record
- Corresponding increase in CPU time/service in the service classes/performance groups where enclaves are running
- MSO and I/O service is unchanged

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.

Classification of DDF Enclaves

Attribute	Value
SI	Subsystem name of the DB2 Subsystem
UI	The DDF thread's AUTHID
AI	Accounting Information for DB2 thread
LU/NET	The client LU name and NETid
CT	"DIST "
CI	DB2 Correlation ID of the DDF thread
CN	DB2 Collection Name
PK	DB2 Package Name
PN	DB2 Plan name
PR	Called procedure name if first statement is a CALL, else blank

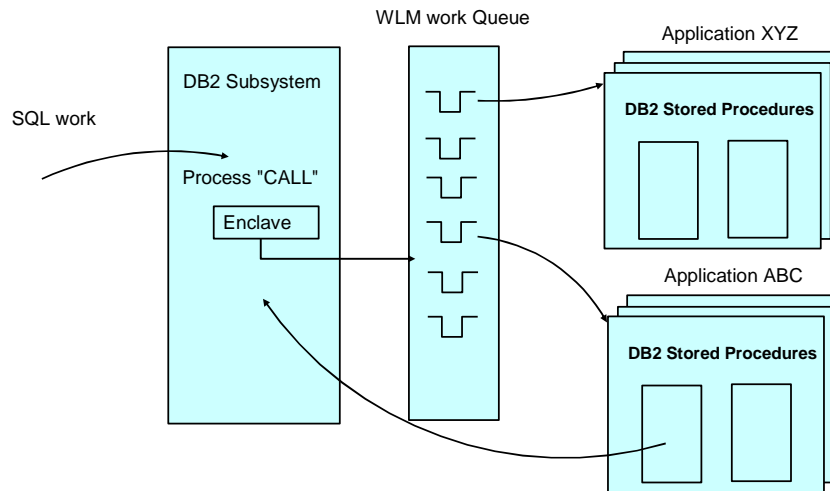
IBM @server. For the next generation of e-business.

Stored Procedures

- DB2
 - Procedures written in Cobol, PL/I, C, ...
 - Executed in Application Environment (APPLENV) defined in WLM
 - WLM manages queue
- Oracle
 - Procedures written in PL/SQL
 - Handled like any other SQL
 - No separate address space, same enclave as ordinary SQL requests

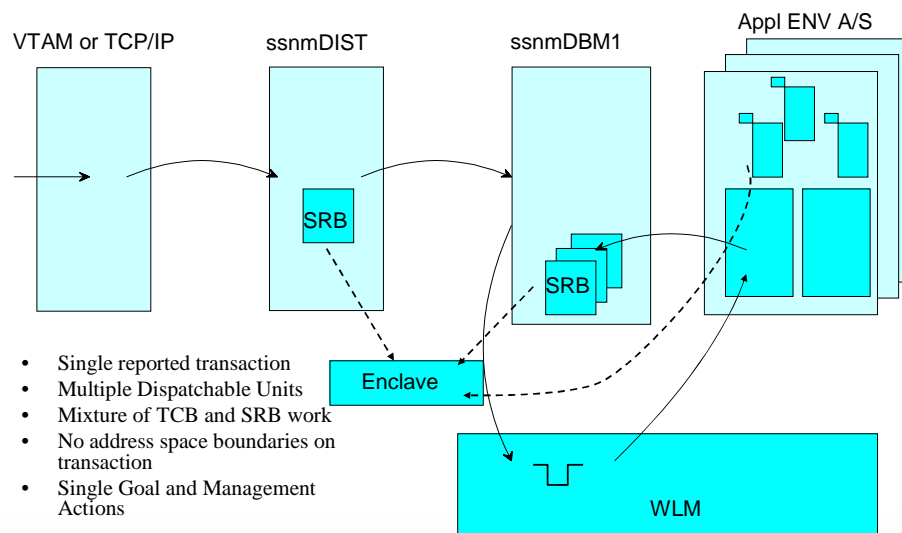
IBM @server. For the next generation of e-business.

DB2 Stored Procedures



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

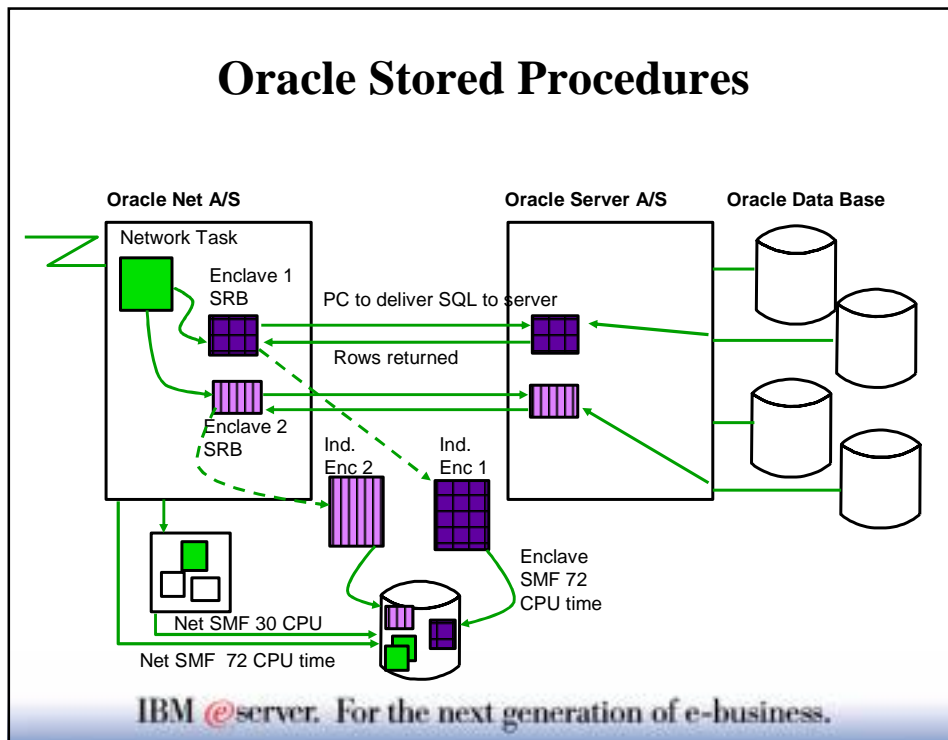
DB2 Stored Procedures ...



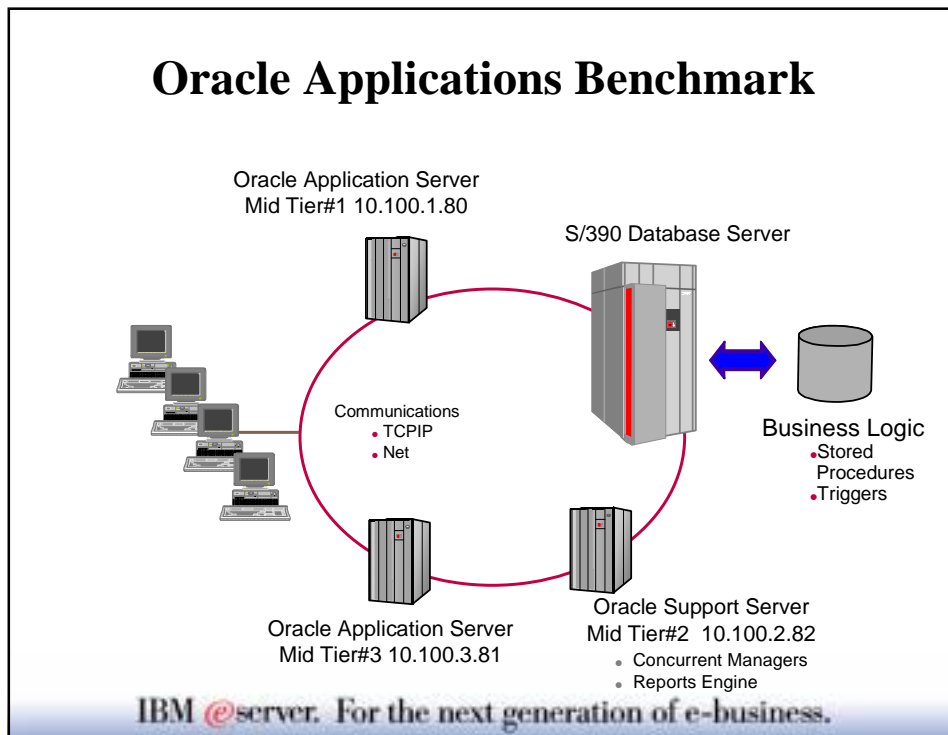
- Single reported transaction
- Multiple Dispatchable Units
- Mixture of TCB and SRB work
- No address space boundaries on transaction
- Single Goal and Management Actions

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

Oracle Stored Procedures



Oracle Applications Benchmark



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
_____ 2 500 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 Enclaves

- 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
 - With DDF transaction is DRDA request

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
ENC00001          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

```

2.10.0 Sysplex Summary - ORACLE                       Line 1 of 18
Command ==>>>                                         Scroll ==>>> CSR

WLM Samples: 400      Systems: 2  Date: 02/20/02  Time: 16.41.40  Range: 100  Sec

>>>>>>>-----<<<<<<<

Service Definition: ORACLE1                            Installed at: 02/20/02, 15.55.26
Active Policy: POLORAL                                Activated at: 02/20/02, 15.56.09

----- Goals versus Actuals -----
Exec Vel  --- Response Time --- Perf  Ended  WAIT EXECUT ACTUAL
Name      T  I  Goal Act  ---Goal--- --Actual--  Indx  Rate  Time  Time  Time
ORACLE   W          23                1897 0.000  0.004  0.004
ORACLES  S  D          29                0.000 0.000  0.000  0.000
ORAMT1   S          15                1897 0.000  0.004  0.004
          1  1          14  0.015 AVG  0.003  AVG  0.23  1895 0.000  0.003  0.003
          2  3          64  0.050 AVG  0.025  AVG  0.50  1.390 0.000  0.025  0.025
          3  D          33                .490 0.000  0.071  0.071
SYSTEM   W          82                0.010 0.000  46.8M  46.8M
SYSSTC   S      N/A  82  N/A                0.010 0.000  46.8M  46.8M
SYSTEM   S      N/A  91  N/A                0.000 0.000  0.000  0.000
    
```

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

Comments and Recommendations

- New code with old WLM policy will not hurt
 - Velocity still appropriate
- Old code with response time goals not appropriate
 - Enclave goes to last period shortly after Logon
- Must have subsystem OSDI or DDF defined, with a default service class specified
 - Mistake in classification rules will result in SYSOTHER being used, which has a discretionary goal. *Very Bad*
- Use ENCLAVE(CALL) in Oracle Net Service
- THREADS=INACTIVE in DDF
 - Establish a short 1st period importance 1 to maintain response for trivial requests
 - Lower importance for 2nd period
 - Importance 5 or discretionary 3rd period

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

Summary

- Enclave support in DDF and Oracle OSDI unique to zSeries
 - Consolidation of multiple smaller DB instances on single S/390 now possible
 - Either multiple or single instance of DB on MVS
 - 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.