



# Oracle and the zIIP processor

G. Tom Russell  
IBM Canada Ltd.  
Tom\_Russell@ca.ibm.com

MVS Oracle SIG  
April 30, 2008  
Redwood Shores, CA

IBM Systems

© 2008 IBM Corporation

# Agenda

- **What is a zIIP processor?**
- **z/OS software support for zIIP**
- **Oracle support for zIIP**
  - ▶ **NET users**
  - ▶ **TG4DB2**
- **RMF WLM reports**


# Evolution of System z Specialty Engines

Building on a strong track record of technology innovation with specialty engines

- DB Compression, SORT, Encryption, Vector Facility



**IBM System z9  
Integrated  
Information  
Processor (IBM  
zIIP) 2006**



**System z Application  
Assist Processor  
(zAAP) 2004**


- Eligible for zAAP:**
- Java™ execution environment
  - z/OS XML

- Eligible for zIIP:**
- z/OS Oracle
  - DB2 remote access and BI/DW
  - New! IPSec encryption
  - z/OS XML
  - z/OS Global Mirror\*



**Integrated  
Facility for Linux  
(IFL) 2000**

\*SOD: IBM plans to enhance z/VM in a future release to support the new System z10 EC capability to allow any combination of CP, zIIP, zAAP, IFL, and ICF processor-types to reside in the same z/VM LPAR



**Internal Coupling  
Facility (ICF) 1997**

# New IBM System z10 Integrated Information Processor (IBM zIIP)

- **New specialty engine for the System z mainframes designed to help:**
  - ▶ Improve resource optimization and lower the cost of ownership for eligible data serving workloads
- **z/OS manages and directs work between the general purpose processor and the zIIP**
  - ▶ Number of zIIPs per System z not to exceed number of standard processors
  - ▶ No changes anticipated to Database applications
  - ▶ No IBM software charges on the zIIP – consistent with other specialty engines
- **DB2 UDB for z/OS V8 was the first IBM exploiter of the zIIP with**
  - ▶ System z9
  - ▶ z/OS 1.6 +
  - ▶ DB2 UDB for z/OS V8
- **Now z/OS Oracle 10g**



# How does the zIIP work ....

- **The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. DB2 and Oracle are similar in the kinds of work sent to the zIIP.**

- Example 1 = DB2 Distributed SQL requests (DRDA)

Queries that access DB2 UDB for z/OS V8 via DRDA<sup>®</sup> over a TCP/IP connection are dispatched within z/OS in enclave SRBs. z/OS directs a portion of this work to the zIIP.

- Example 2 = Oracle Distributed SQL requests (NET)

Queries that access Oracle for z/OS via NET are dispatched within z/OS in enclave SRBs. z/OS directs all of this work to the zIIP.

- Example 3 = DB2 other work

Index component of Utilities, and Parallel Query

## What are enclave SRBs?

- **z/OS dispatches work in either TCB (Task Control Block) mode or SRB (Service Request Block) mode.**
  - **Enclaves are used to do work on behalf of the originating TCB or SRB address space, even if the work is run in many different address spaces.**
  - **If the DB2 for z/OS V8 request is coming in over distributed (ie DRDA over TCP/IP) then that work is executed in enclave SRBs.**
  - **Oracle NET runs its work in enclave SRBs**
- **So..... regarding the zIIP, only the enclave SRB work, not enclave TCB work is eligible to be redirected to the zIIP.**

## z/OS zIIP software support

- **IBM System z Integrated Information Processor Support for z/OS and z/OS.e**
- **Support built-in to z/OS 1.8 and subsequent**
- **z/OS 1.7 FMID JBB772S Web deliverable**  
<http://www-3.ibm.com/servers/eserver/zseries/zos/downloads/#ziip>
- **PSP buckets have lists of PTFs that are also required so call IBM support**

# Oracle software support

- **10g NET**
- **Latest level of OSDI is available as a stand alone patch:**
  - ▶ **Patch 6138068**
  - ▶ **Currently OSDI 10.2.0.3**
  - ▶ **Gives you a new AUTHLOAD**
    - **Put in front of STEPLIB in NET and DB Service**
    - **Copy SUBSYSTEM code to LNKLST**
  - ▶ **Support in 10g NET, DB can be 9i**
  - ▶ **New OSDI 1.2 coming soon, same Patch Number**

# 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**

## Work redirected to zIIP

- **All work currently running in an OSDI service class**
  - ▶ **Velocity or Response Time Goal**
  - ▶ **ENCLAVE(CALL) or ENCLAVE(SESS)**
  - ▶ **You could measure by installing the z/OS and Oracle zIIP software, and adding projectcpu=yes to IEAOPTxx**
- **NOT work using batch XM connection**
- **NOT AM4CICS**
- **NOT AM4IMS**
- **NOT Parallel Query Slaves**
  - ▶ **Enhancement Request 3630671 rejected**

# Example for Enterprise Applications

Enterprise Applications that access Oracle z/OS via NET over a TCP/IP connection will have all of their SQL requests directed to the zIIP



# WLM Support in Oracle Net...

- **Define Net Service**

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

- **ENCLAVE(CALL)**

- ▶ dynamic enclaves, many WLM transactions

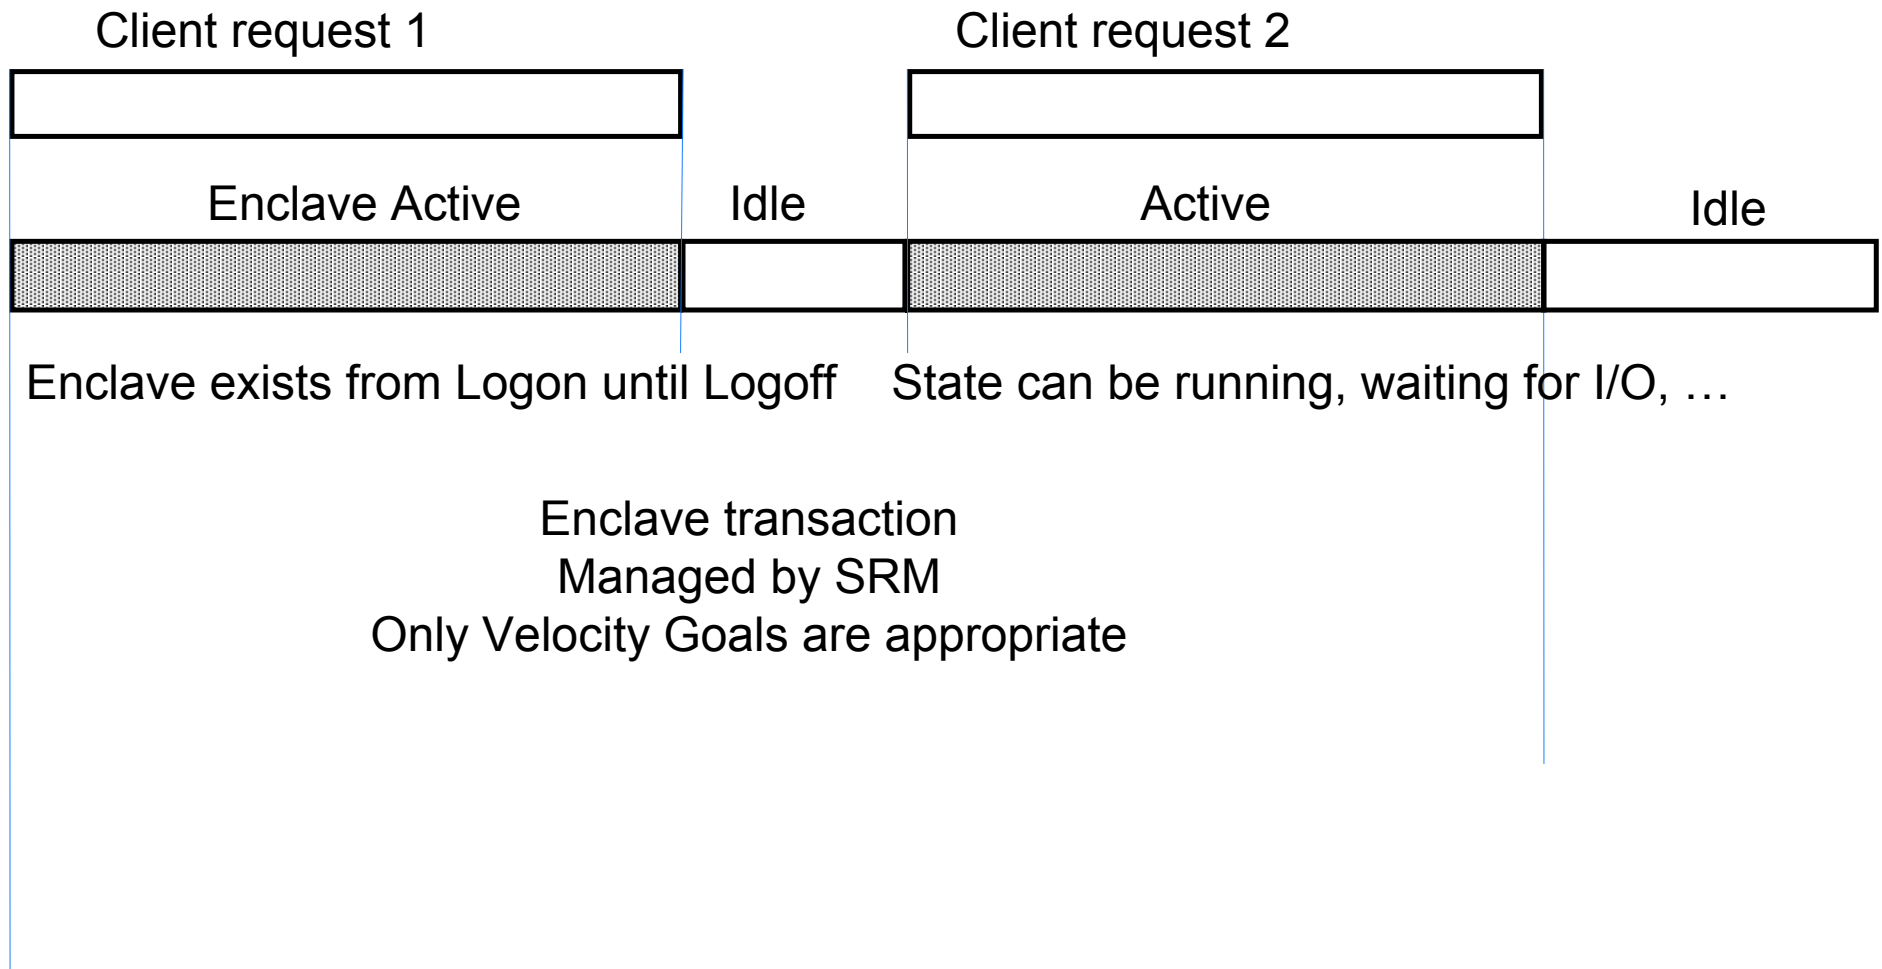
- **ENCLAVE(SESS)**

- ▶ static enclaves (default), one WLM transaction per session

## 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**

# NET Behaviour with Enclave(sess)



# NET Behaviour with Enclave(call)

Client request 1




Enclave exists until NET has no more data

Enclave transaction  
Managed by SRM

Response time reported by RMF

Client request 2



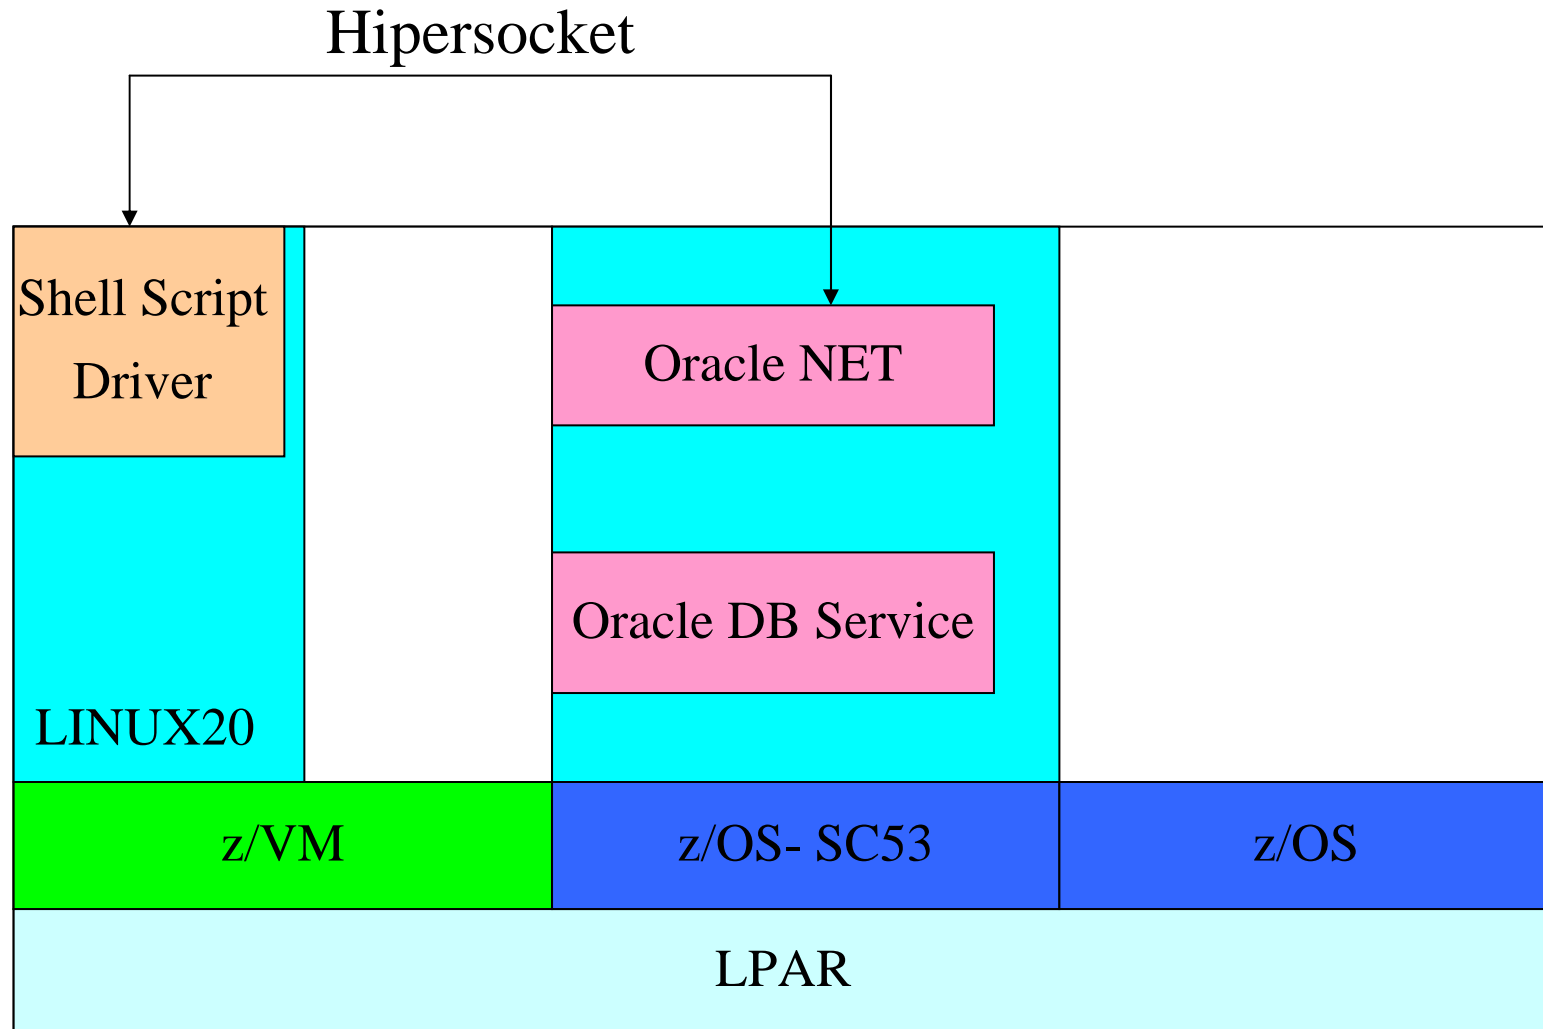
State can be running, waiting for I/O, ...

Enclave transaction  
Managed by SRM

## 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**

# Test Environment in ITSO Poughkeepsie



## Test Workload

- **Shell script driver written by a colleague**
- **Very light weight transactions**
- **Oracle NET set for Enclave(CALL)**
- **WLM service class ORATRAN with response time goals, 3 periods**

# RMF III System Summary Report

----- Goals versus Actuals -----													
		Exec Vel		--- Response Time ---		Perf		Trans	--Avg. Resp. Time--				
Name	T	I	Goal	Act	---Goal---	--Actual--	Indx	Ended	WAIT	EXECUT	ACTUAL		
			Rate	Time	Rate	Time		Rate	Time	Time	Time		
ORACLE	W			53				413.4	0.000	0.003	0.003		
ORATRAN	S			53				413.4	0.000	0.003	0.003		
	1	1		55	0.100	AVG	0.003	AVG	0.03	409.4	0.000	0.003	0.003
	2	2		46	1.000	AVG	0.015	AVG	0.01	4.040	0.000	0.014	0.015

# RMF WLM Report fragment

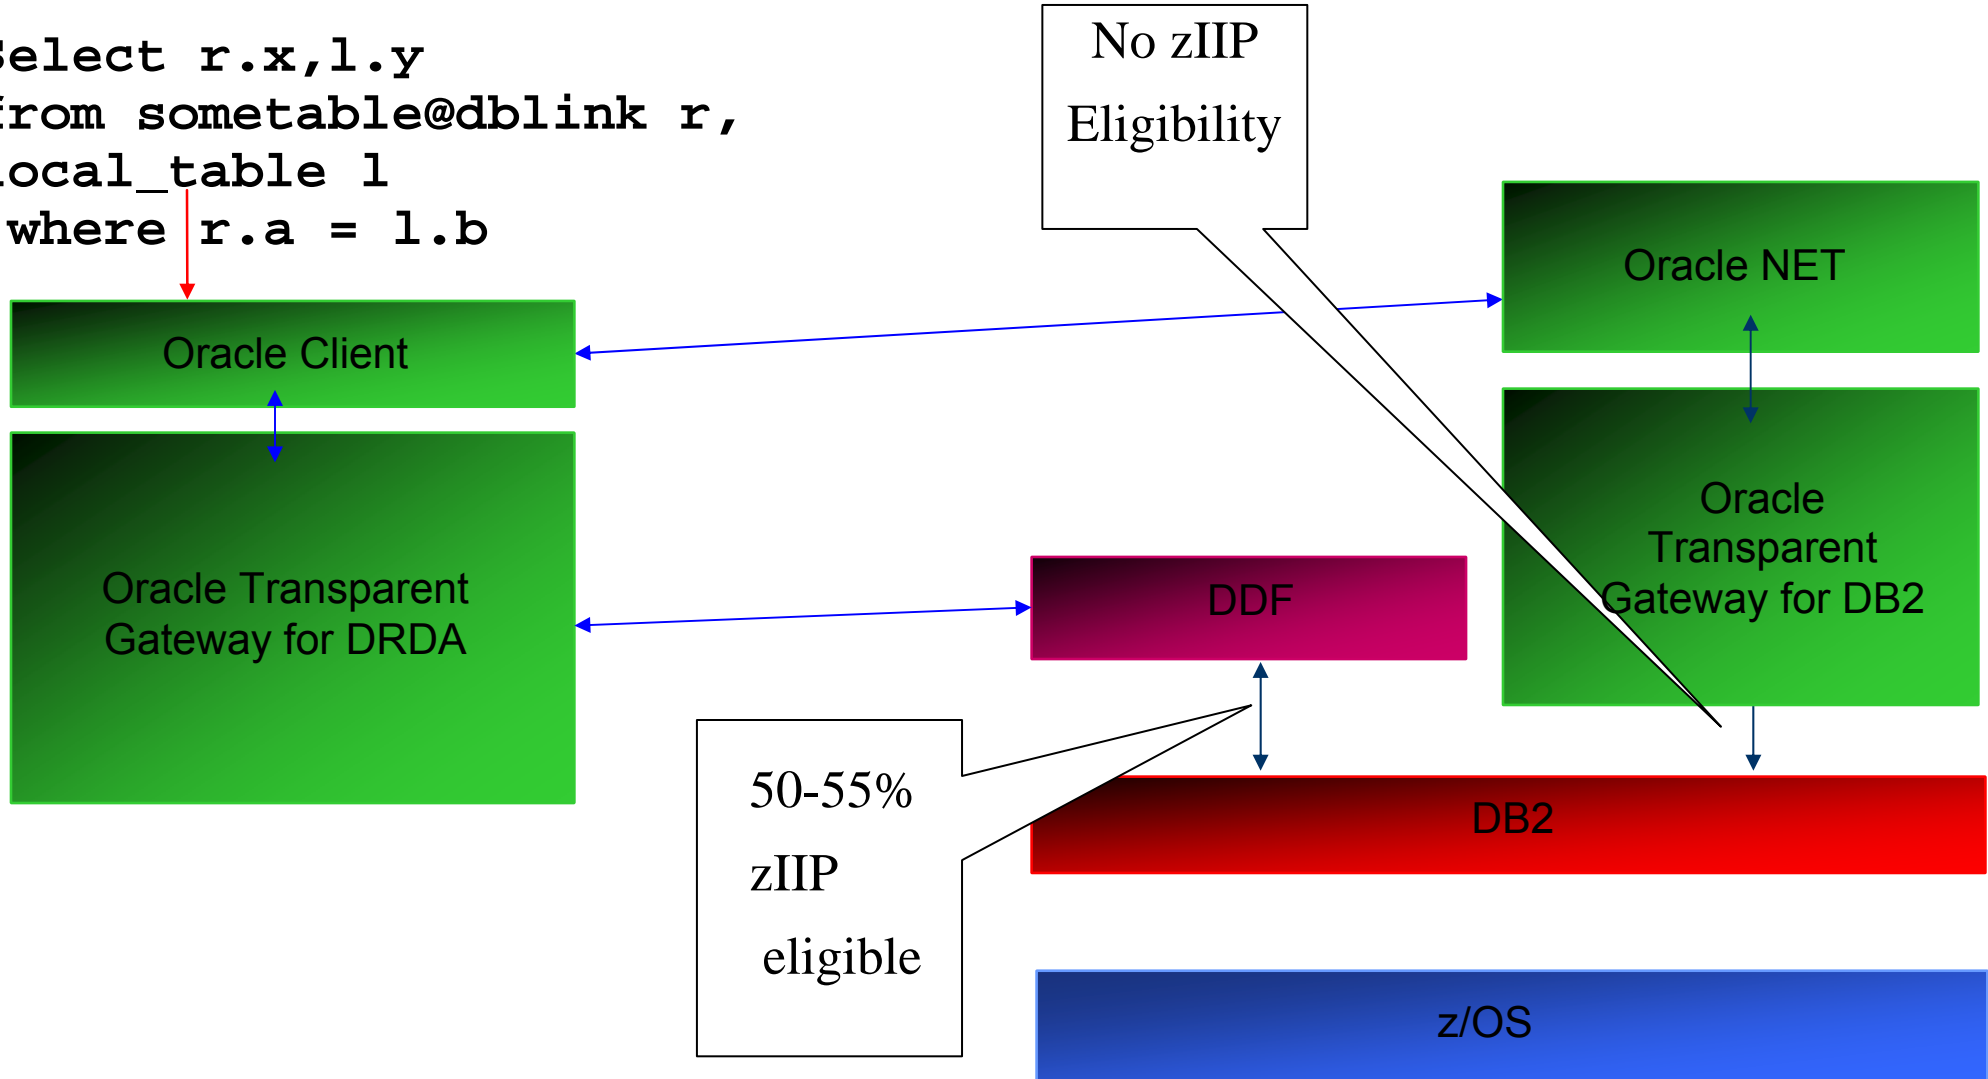
REPORT BY: WORKLOAD=ORACLE SERVICE CLASS=ORATRAN

APPL%  
is % of one  
engine

TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	---	APPL	%---
AVG	1.33	ACTUAL	3	<b>CP</b>	<b>0.35</b>
MPL	1.33	EXECUTION	3	AAPCP	0.00
ENDED	244535	QUEUED	0	<b>IIPCP</b>	<b>0.31</b>
<b>END/S</b>	<b>407.42</b>	R/S AFFIN	0		
#SWAPS	0	INELIGIBLE	0	AAP	0.00
EXCTD	0	CONVERSION	0	<b>IIP</b>	<b>16.19</b>
AVG ENC	1.33	STD DEV	12		

# Oracle10g Transparent Gateways

```
Select r.x,l.y
from sometable@dblink r,
local_table l
where r.a = l.b
```



# Batch XM Connections

- **Remove //ORA@xxxx DD DUMMY**
- **Change connect string to:**
  - ▶ **User/password@xxxx\_TCP**
  - ▶ **Name in tnsnames file using TCP/IP**
- **Job shows very little CPU time**
- **All CPU time shows up in IIP time for the OSDI service class assigned to the enclave by NET**
- **Limited information to use for classification**
  - ▶ **User ID and IP Address**

# Added entry for ORA1

```
ORA1_TCP_LB=  
    (DESCRIPTION=  
        (ADDRESS=  
            (PROTOCOL = TCP)  
            (HOST = 127.0.0.1)  
            (PORT = 1501)  
        )  
        (CONNECT_DATA=(SID=ORA1))  
    )
```

## Two batch jobs running

```
SDSF ENCLAVE DISPLAY SC53      ALL
```

```
COMMAND INPUT ===>
```

NP	TOKEN	SSType	Status	SrvClass	Per
	20000000FA	OSDI	ACTIVE	ORAVEL	1
	2400000111	OSDI	ACTIVE	ORAVEL	1

- **Allows classifying batch into a different Service Class based on IP address used to connect**

## Classification data for the Batch Enclave

Subsystem type	OSDI	Plan name	
Subsystem name	ORSS	Package name	
Priority		Connection type	TCP
Userid	RUSSELL	Collection name	
Transaction name		Correlation	
Transaction class		Procedure name	
Netid	127.000.	Function name	OSDINET8
Logical unit name	000.001	Performance group	
Subsys collection		Scheduling env	
Process name			
Netid	009.012.	Function name	OSDINET8
Logical unit name	004.145	Performance group	

# RMF WLM for Batch enclaves

APPL% is % of  
one engine

REPORT BY: WORKLOAD=ORACLE SERVICE CLASS=ORAVEL

TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	APPL %
AVG	2.00	ACTUAL	0
MPL	2.00	EXECUTION	0 <b>CP</b> 0.00
ENDED	0	QUEUED	0 AAPCP 0.00
END/S	0.00	R/S AFFIN	0 <b>IIPCP</b> 0.00
#SWAPS	0	INELIGIBLE	0
EXCTD	0	CONVERSION	0 AAP 0.00
AVG ENC	2.00	STD DEV	0 <b>IIP</b> 198.70

## Considerations

- **Examples used earlier all used Oracle NET**
  - ▶ Supports 8i, 9i, 10g Databases
- **Generic Listener is also supported**
  - ▶ DB must be 10gR2 or later
  - ▶ Code to build enclave, support zIIP, do TCP/IP is all in new code in 10gR2 DB server.
    - **NEW\_SOCKETS(YES)** parameter
- **Long term, I recommend Generic Listener**
  - ▶ New hidden OSDI options.. Call support
  - ▶ Enhancement Request 3675452 for improving the **KILL SESSION** command depends on Generic Listener

# Summary

- **z9 or z10 processor with a zIIP can significantly reduce the cost of ownership of a z/OS Oracle**
- **Cheaper hardware**
  - ▶ zIIP much less \$ than CP
- **zIIP processor faster than CP on many z9 BC machines**
- **Resource group capping can no longer be used.**
  - ▶ CPU time on zIIP not counted by WLM for capping
- **OSDI SMF record will contain zIIP and zIIP on CP time**
- **No software charges for capacity used on zIIP**
  - ▶ Look at WLM report to determine your potential
- **Move to latest level of Oracle**

Thank you – Any Questions?

**Tom\_Russell@ca.ibm.com**