IBM Software Group

Understanding, Monitoring and Managing z/OS Enclaves

Ed Woods / IBM Corporation

Session 10667

Thursday, March 15, 2012: 4:30 PM-5:30 PM

Agenda

- Terminology
- What Is An Enclave?
- Enclave Examples
- Displaying Enclave Information
- DB2 Monitoring Considerations
- Enclaves And WLM Considerations

Some Important Terminology TCBs and SRBs

- z/OS uses a huge variety of control blocks, many with very specialized purposes
- The three most commonly used control blocks are:
  - Task control block (TCB) - represents a unit of work or task, such as an application program, that runs in an address space
  - Service request block (SRB) - represents a request for a system service
    - SRBs are typically created when one address space detects an event that affects a different address space
    - SRBs provide a mechanism for communication between address spaces.
  - Address space control block (ASCB), which represents an address space

More About TCBs And SRBs

- To identify and keep track of its work z/OS represents each unit of work on the system with a control block
- Standard dispatching units (TCBs and SRBs)
  - TCB - runs at dispatching priority of address space and is preemptible
  - SRB - runs at supervisor priority and is non-preemptible
- Advanced dispatching units
  - Enclave
    - Serves as an anchor for an address space independent transaction
    - Can consist of multiple tasks (TCBs or SRBs) executing across multiple address spaces
  - Client SRB
    - Similar to an ordinary SRB but runs with client dispatching priority and is preemptible
  - Enclave SRB
    - Similar to an ordinary SRB but runs with enclave dispatching priority and is preemptible

So What Are Enclaves?

- Enclaves represent a "business unit of work"
- Enclaves are managed separately from the z/OS address spaces
- Enclaves can include multiple SRBs/TCBs
  - Can span multiple address spaces
  - Can have many enclaves in a single address space
  - Assigned by WLM to a service class for prioritization by the system

What Is A Business "Unit Of Work"?

- A "unit of work" represents a WLM transaction
  - An item of work where WLM collects resource usage information
  - Represents a subsystem work request
    - WLM can measure resources used by the subsystem request
- Types of transactions
  - Address Space
    - WLM will measure all resource used by a subsystem request in a single address space
  - Enclave
    - Enclave created and used by a subsystem for each work request across multiple address spaces and systems
    - Used by a variety of workloads; DB2, DB2 DDF, WebSphere, MQ, LDAP, TCP/IP
  - CICS and IMS Transactions
    - Not address space or enclave oriented
    - Measures resource used by CICS/IMS transaction requests

Who Uses Enclaves?

- Enclaves have become a pervasive mechanism in the z/OS operating system
- DB2 was one of the early exploiters of the concept of enclaves
  - Enclaves provided a mechanism to manage and prioritize DB2 distributed (DDF) workload
  - More exploitation added with subsequent DB2 releases
    - DB2 stored procedure support
    - DB2 sysplex query parallelism
    - DB2 sequential prefetch and deferred write processing (DB2 10)
- Many core z/OS components use enclaves
  - MQseries, WebSphere, TCP/IP, LDAP

Categories Of Enclaves

- Independent Enclaves
  - Use an independent enclave to represent a new transaction
  - An independent enclave must be classified into a service class or performance group when it is created
- Dependent Enclaves
  - Use a dependent enclave when you have an existing address space defined with its own performance goal
  - Extends that goal to programs running under dispatchable units in other address spaces
- Work-dependent Enclaves
  - Use a work-dependent enclave to extend an existing independent enclave's transaction (for zIIP support – more on this later)

WLM Enclaves – An Example

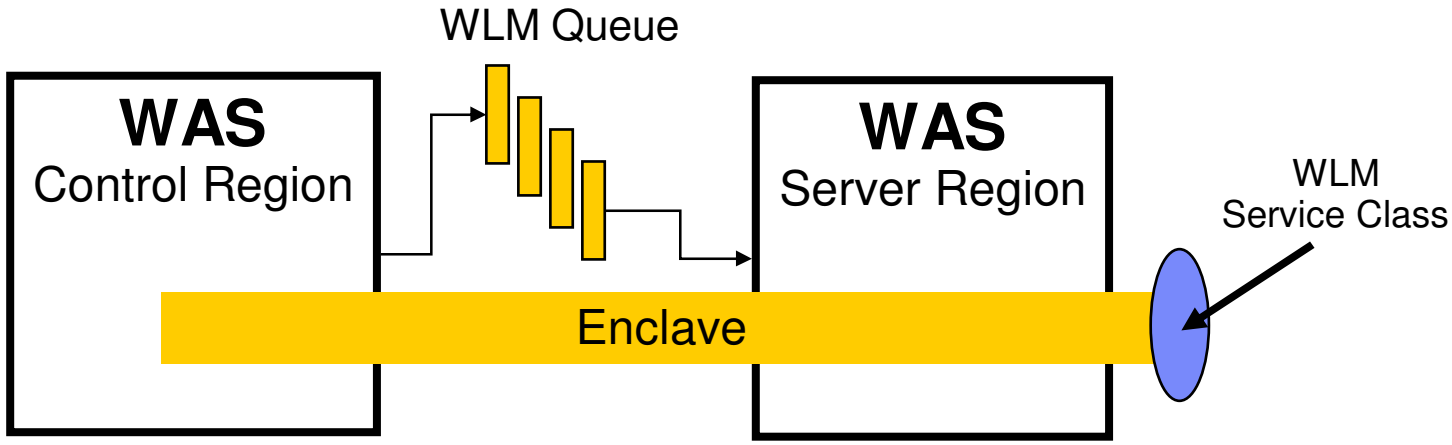

- The enclave is managed separately from the address spaces it runs in
  - CPU and I/O resources associated with processing the work request represented by the enclave may be managed by the transaction's performance goal
- • Storage resources may be managed as follows
  - To the goals of the enclaves it serves (if enclave server address space)
  - To the performance goal of the address space (if no server address space)

Example - Enclaves Provide DB2 DDF With Granularity And Control

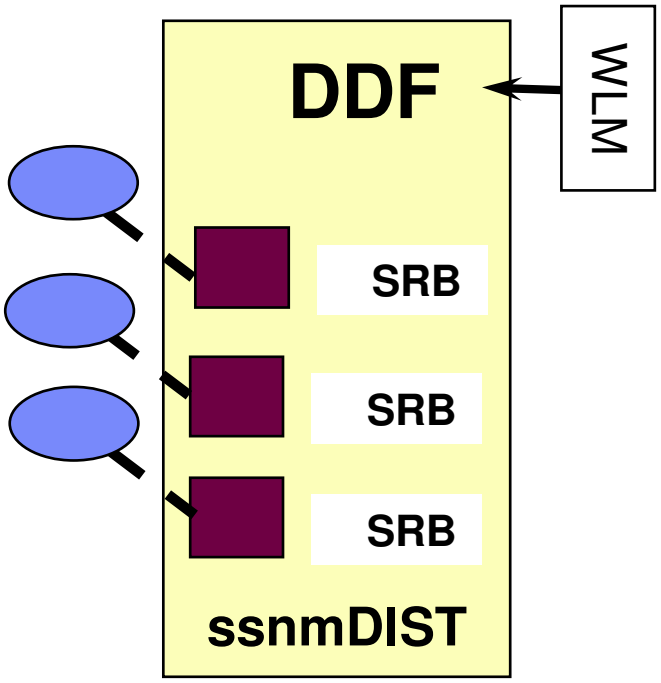

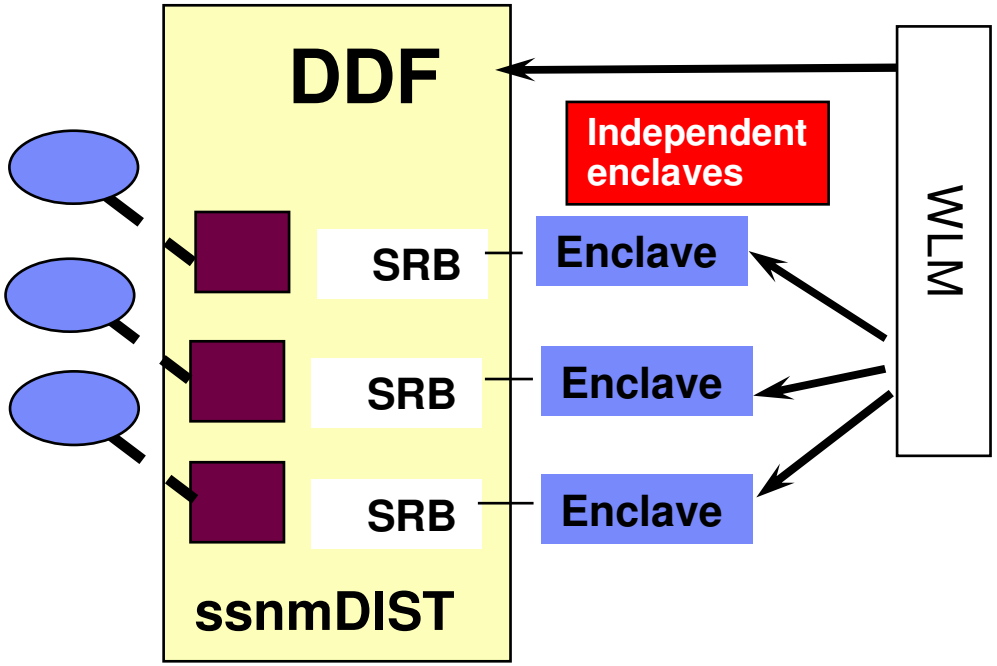

- DDF prior to enclaves
- Workload ran at the priority of the DDF task

- DDF with enclaves
- WLM has more control and granularity to prioritize work

Example - DDF Stored Procedure Priority

- When a Stored Procedure is called from DDF thread
  - DB2 references the enclave created for the DDF request for Stored Procedure
  - Stored Procedure priority is the priority of the DDF request

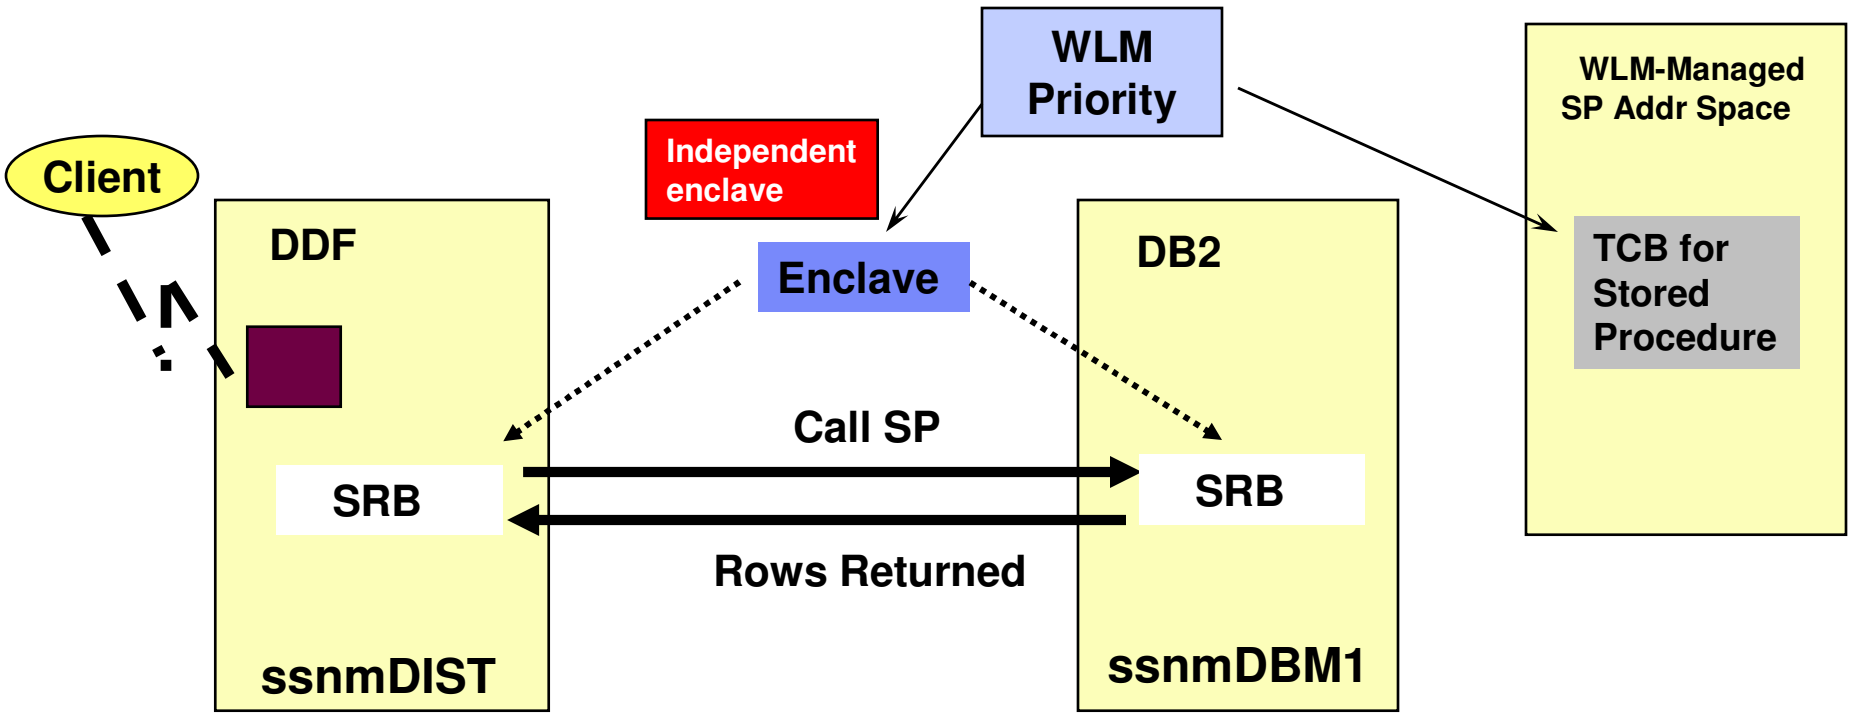

Stored Procedure Priority
Called From A Local Application

- When a Stored Procedure is called from an application on z/OS
  - DB2 creates an enclave for use by the Stored procedure
  - Stored Procedure priority is the priority of the calling application address space

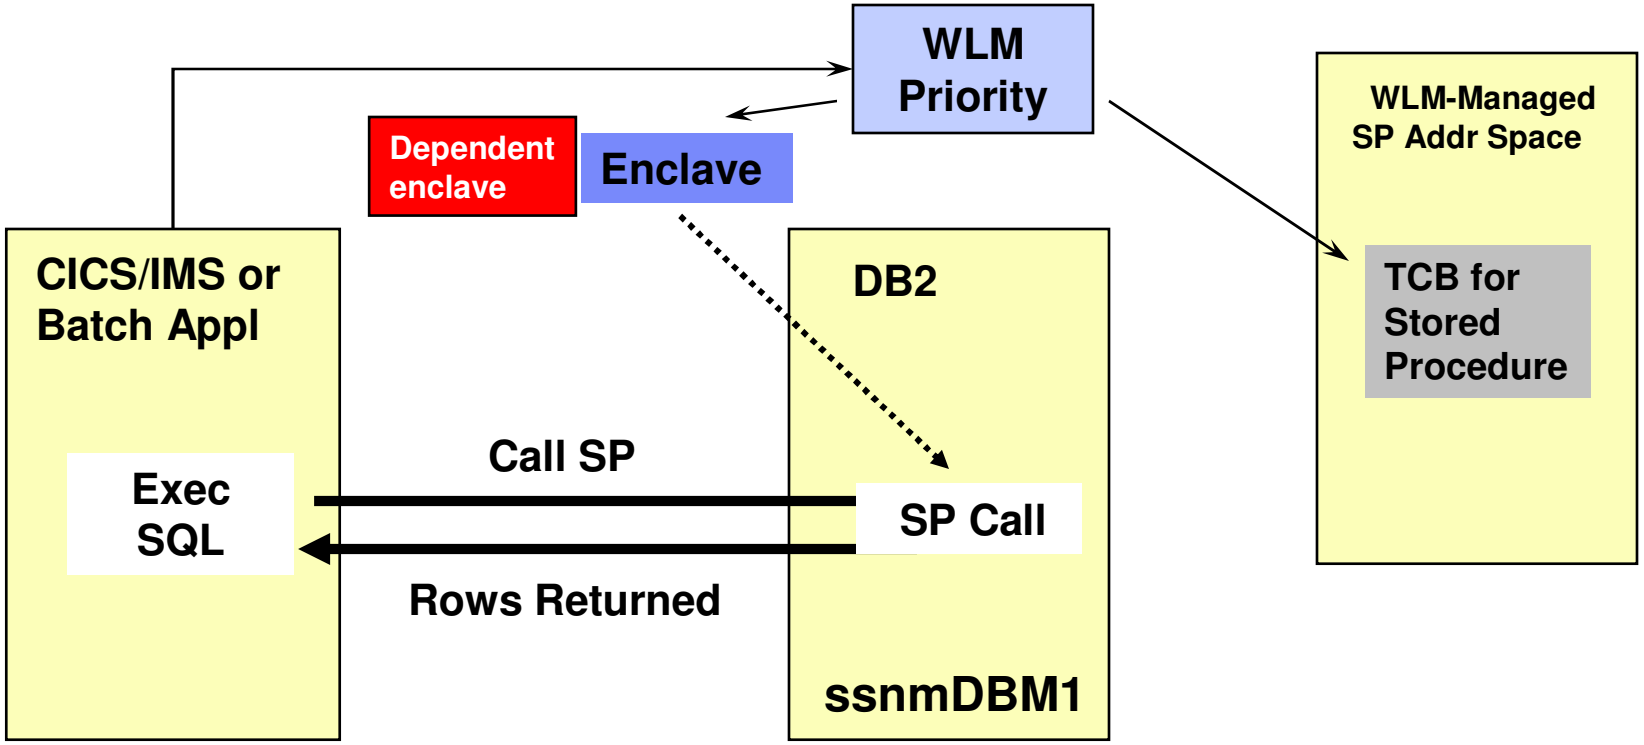

Summary - DB2 Workload Prioritization
How Does WLM Assign Priority To DB2 Workload?

- The priority of the DB2 workload will vary depending upon the origin of the workload

- DB2 workload originating from a local application (examples - IMS, CICS, TSO, Batch, WebSphere)
  - Priority is inherited from the invoking application
  - This applies to Stored Procedures invoked locally

- DB2 Distributed requests (Subsystem type DDF)
  - Priority controlled by DDF Service Class definitions
  - DB2 Stored Procedure request via DDF - priority controlled by Service Class definitions

- DB2 Sysplex Query parallelism (Subsystem type DB2)
  - Classification done by DB2 Service Class definitions

DB2 And zIIP Processors

- Work on z/OS may have all or a portion of its resource usage on an enclave SRB
  - Enclave SRB work may be directed to the zIIP
- Certain types of DB2 work may take advantage of zIIP
  - DRDA - Queries that access DB2 for z/OS via DRDA over TCP/IP
  - Complex parallel queries
  - DB2 utilities for index maintenance
    - LOAD, REORG, and REBUILD
  - DB2 V10 – Sequential prefetch eligible for zIIP processor
- WLM and new enclave structures to manage zIIP related workload – work dependent enclave

About Work-dependent Enclaves

- A 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
  - Allows to have its own set of attributes (including zIIP offload percentage)

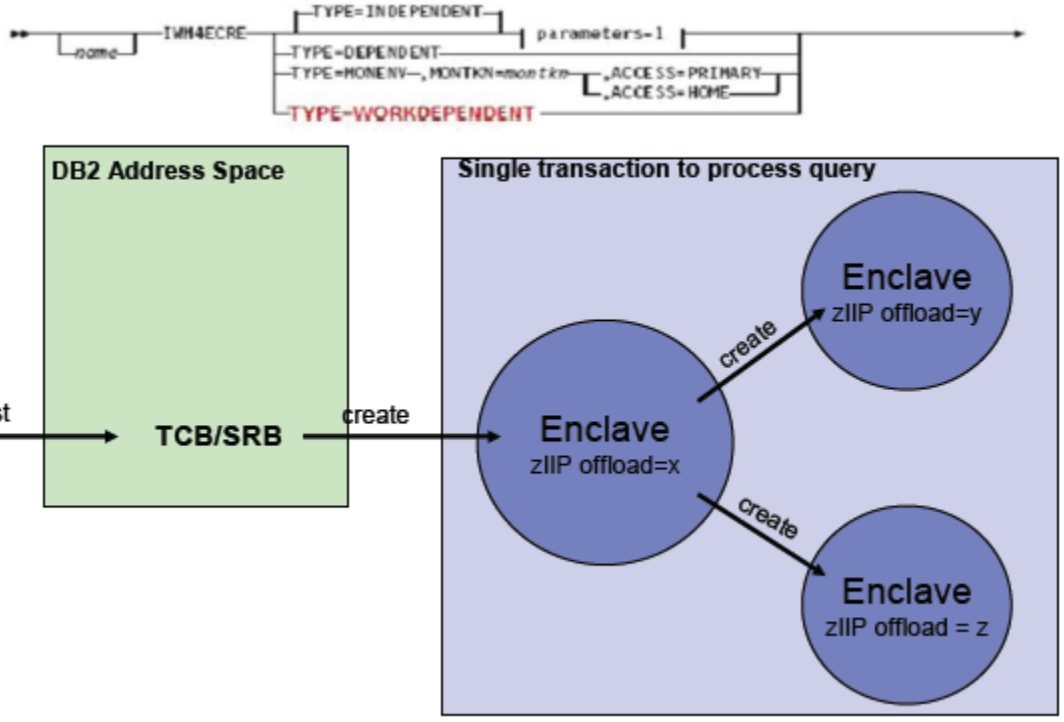

DB2 V10 Exploits Enclaves For Prefetch

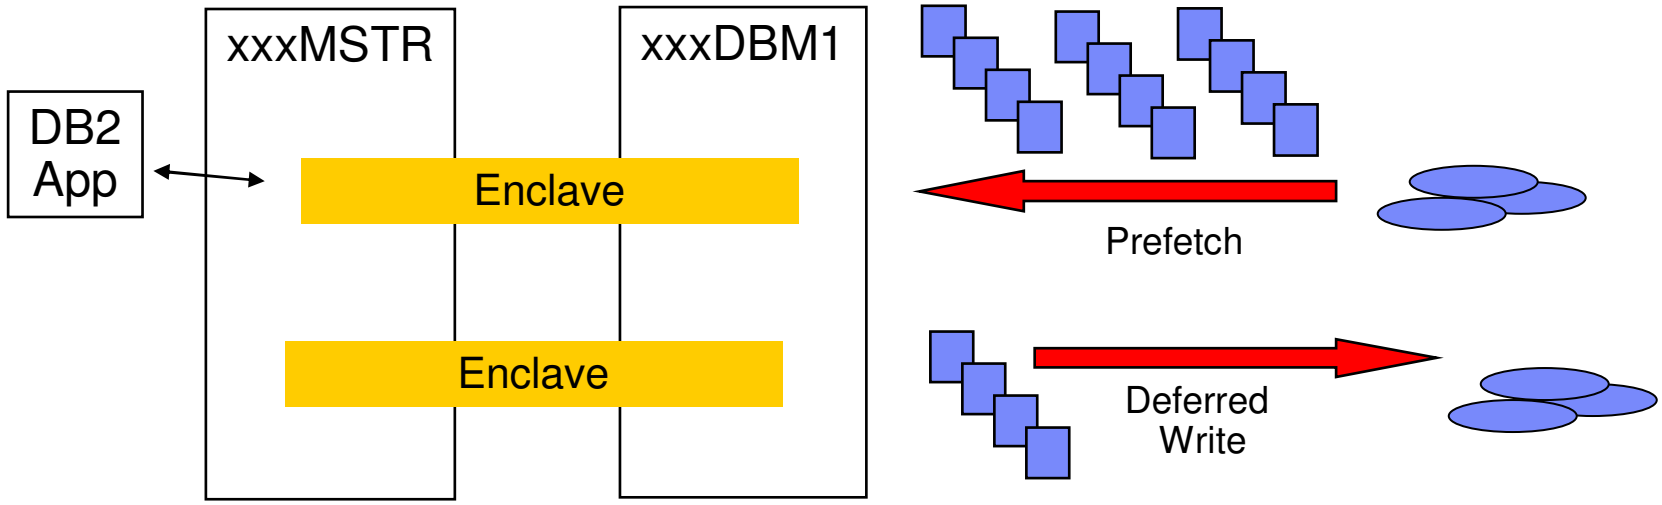

- Buffer pool prefetch activity (dynamic prefetch, list prefetch, sequential prefetch) is 100% zIIP eligible in DB2 10
- DB2 10 zIIP eligible buffer pool prefetch is asynchronously initiated by the DBM! address space
  - Executed with a dependent enclave owned by the MSTR address space
  - Deferred write also eligible for zIIP
- Asynchronous buffer pool prefetch activities are not accounted to the DB2 client
  - Shows up in the DB2 statistics report

DB2 V10 Exploits Enclaves For Prefetch - continued

CPU TIMES	TCB TIME	PREEMPT SRB	NONPREEMPT SRB	TOTAL TIME	PREEMPT IIP SRB
SYSTEM SERVICES ADDRESS SPACE	2.565794	0.825791	0.310654	3.702240	N/A
DATABASE SERVICES ADDRESS SPACE	0.187984	0.094044	0.003409	0.285437	0.486775
IRLM	0.000002	0.000000	0.065226	0.065228	N/A
DDF ADDRESS SPACE	0.005344	0.000000	0.000105	0.005448	0.000000
TOTAL	2.759124	0.919835	0.379394	4.058353	0.486775

- Asynchronous I/O processing is important to DB2 performance
- With DB2 10 buffer pool prefetch activities are asynchronously initiated by the DBM1 address space
  - This is executed in a dependent enclave
- Asynchronous buffer pool prefetch activities are not accounted to the DB2 client application
  - CPU time accounted to the zIIP appears in DB2 statistic report (PREEMPT IIP SRB)

Displaying Enclaves

SDSF Example – DA Display And The ENC Command

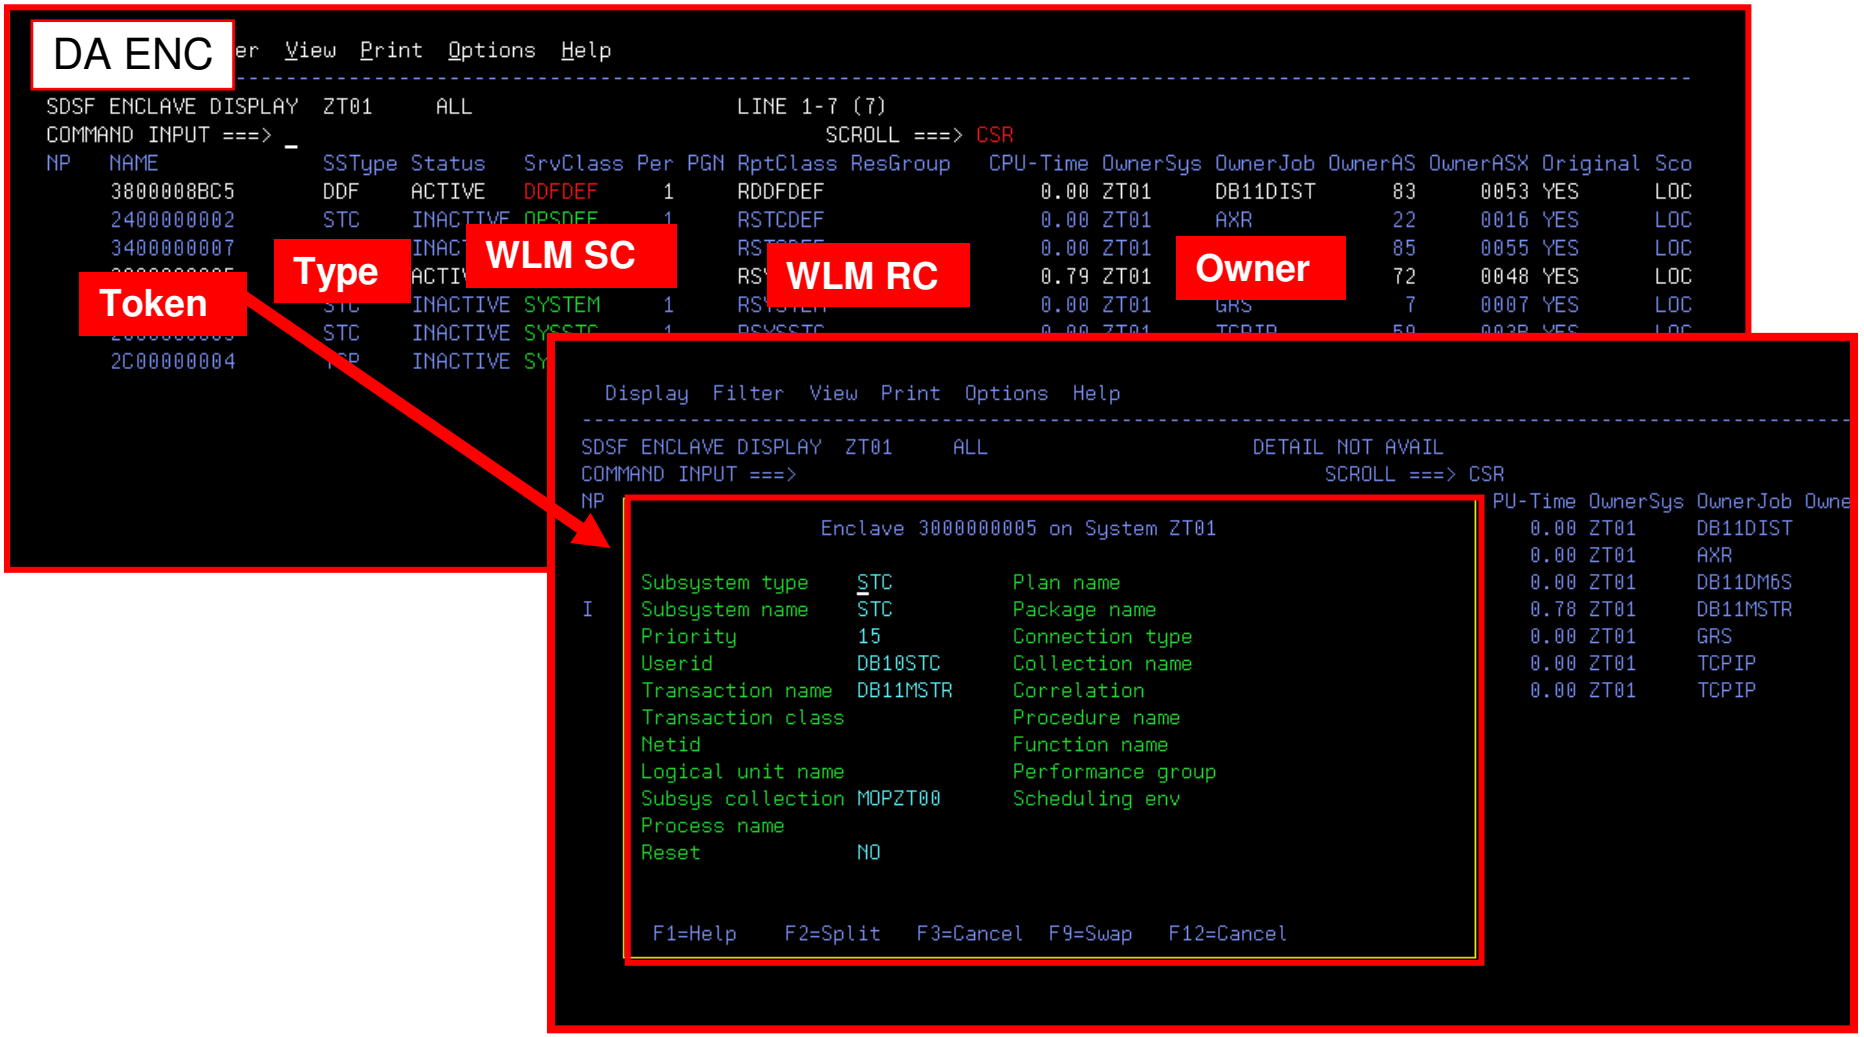

Information On Enclaves In SDSF - continued

```
  Display  Filter  View  Print  Options  Search  Help
------------------------------------------------------------------------------
SDSF ENCLAVE DISPLAY  DEMOMVS  ALL                        LINE 26-47 (47)
NP    NAME              P-Time zIIP-Time zICP-Time Promoted zAAP-NTime zIIP-NTime
      FC0000987A          0.00     58.58      0.90 NO             0.00      58.58
      100000075E7         0.00      0.00      0.00 NO             0.00       0.00
      108000098AD         0.00      0.08      0.03 NO             0.00       0.08
      11000012B0D         0.00      0.00      0.00 NO             0.00       0.00
      1140000ADB1         0.00    103.20      1.15 NO             0.00     103.20
      1180000B147         0.00      3.81      0.01 NO             0.00       3.81
      3800000007          0.00      0.00      0.00 NO             0.00       0.00
      2400000002          0.00      0.00      0.00 NO             0.00       0.00
      60000103AD          0.00      0.01      0.00 NO             0.00       0.01
      5C00010442          0.00      5.43      0.00 NO             0.00       5.43
      9800011514          0.00      0.00      0.00 NO             0.00       0.00
      A000010485          0.00      7.79      0.11 NO             0.00       7.79
      4400000010          0.00      0.00      0.00 NO             0.00       0.00
      3400000006          0.00      2.82      0.00 NO             0.00       2.82
      3000000005          0.00      0.00      0.00 NO             0.00       0.00
```

- SDSF ENC display will also show enclave CPU usage by various categories
  - zAAP-Time Accumulated zAAP time, in seconds
  - zACP-Time Accumulated zAAP on CP time, in seconds
  - zIIP-Time Accumulated zIIP time, in seconds
  - zICP-Time Accumulated zIIP on CP time, in seconds
  - zAAP-NTime Normalized zAAP time, in seconds
  - zIIP-NTime Normalized zIIP time, in seconds

Looking At Enclave Activity Using RMF And SMF

```
RMF V1R13  Enclave Report                    Line 1 of 7

Samples: 100     System: MVSA  Date: 02/24/12  Time: 10.13.20  Range: 100   Sec

Current options:   Subsystem Type: ALL                          -- CPU Util --
                   Enclave Owner:                               Appl%   EAppl%
                   Class/Group:                                   2.1      2.6

Enclave   Attribute  CLS/GRP  P Goal    % D X   EAppl%   TCPU    USG  DLY  IDL

*SUMMARY                                        0.031
ENC00006            DDFDEF   2        20 W     0.016  59.20    1.6  0.0  0.0
ENC00005            DDFDEF   2        20       0.013  39.54    0.0  0.0  0.0
ENC00004            OPSDEF   1        60 Y     0.001  23.31    0.0  0.0  0.0
ENC00003            OPSDEF   1        60 Y     0.001  14.21    0.0  0.0  0.0
ENC00002            OPSDEF   1        60 Y     0.000  3.269    0.0  0.0  0.0
ENC00001            OPSDEF   1        60 Y     0.000  2.254    0.0  0.0  0.0
```

- You can view total enclave usage interactively using RMF
  - Still, in the case of DB2 workload you may need DB2 accounting trace data to see detail on resource consumption for specific workload items
- About SMF type 30 and SMF type 72 records
  - SMF type 30 record contains resource consumption at the address space level
  - SMF type 72 contains information at the z/OS WLM service class (or report class) level
    - Note – use WLM classification rules to exploit the ability to classify and analyze workload by WLM report class

DB2 DDF Relevant SMF Information

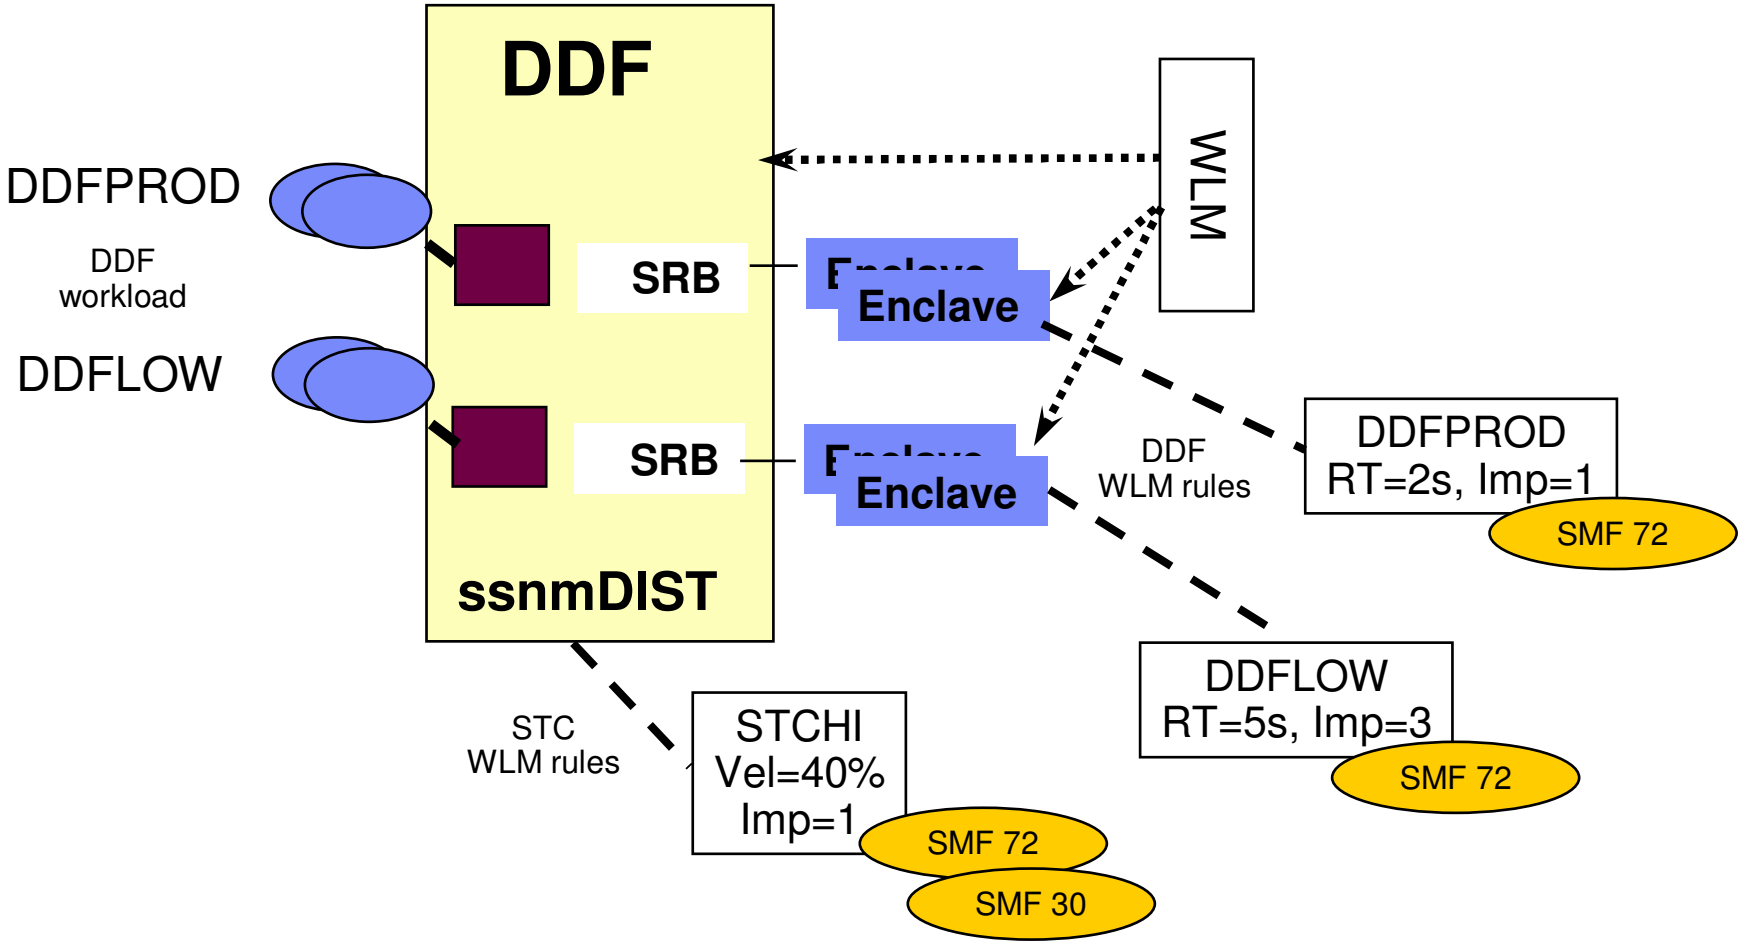

DB2 Example
Displaying A DB2 Thread

```
>                              THREAD DETAIL
 PLAN
+ Thread:  Plan=DISTSERV  Connid=SERVER   Corrid=db2bp.exe    Authid=DNET581
+ Dist   :  Type=DATABASE ACCESS, Luwid=G941491B.PA10.090604181240=168
+ Location  :  9.65.73.27
 act
+ Thread Activity                        User Defined Functions
+ ------------------------------------  ------------------------------------
+ DB2 Status            =   WAIT-REMREQ  TCB Time (SQL)          = 00:00:00.000
+ MVS Status            =                Wait for TCB Time       = 00:00:00.000
+ Total Elapsed Time    =  00:07:45.561  Elapsed Time            = 00:00:00.000
+ CP CPU Utilization    =         00.0%  Elapsed Time (SQL)      = 00:00:00.000
+ Total CP CPU Time     =  00:00:00.082  SQL Event               =             0
+ IIP CPU Time          =  00:00:00.000
+ Total Parallel Tasks  =             0
+ Current Parallel Tasks=             0
+
+ Stored Procedures                      Triggers
+ ------------------------------------  ------------------------------------
+ Total CPU             =  00:00:00.000  TCB not in Enclave      = 00:00:00.000
+ Elapsed time          =  00:00:00.000  Elapsed not in Enclave  = 00:00:00.000
+ Elapsed Time (SQL)    =  00:00:00.000  TCB prior to Enclave    = 00:00:00.000
+ Wait for TCB Time     =  00:00:00.000
+ Wait Event Count      =             0
+ Curr Wait TCB Time    =  00:00:00.000
+
+ SavePoints
+ ------------------------------------
+ Savepoint Requests    =             0
+ Release Savepoints    =             0
+ Rollback Savepoints   =             0
+
+ In-DB2 Times                                Total          Current
+ ----------------------------               ------------   ------------
+ Elapsed Time                               00:00:00.079   00:00:00.000
+ CP CPU Time                                00:00:00.033   00:00:00.000
+ IIP CPU Time                               00:00:00.041            N/A
+ Stored Procedure CPU Time                  00:00:00.000   00:00:00.000
+
+ Waits                                Count  Total          Current
+ ---------------------------------- ------- ------------   ------------
+ Synchronous I/O Wait                     0  00:00:00.000   00:00:00.000
+ Asynchronous Read I/O Wait               0  00:00:00.000   00:00:00.000
+ Asynchronous Write I/O Wait              0  00:00:00.000   00:00:00.000
+ Local Lock/Latch Wait                    0  00:00:00.000   00:00:00.000
+ Page Latch Wait                          0  00:00:00.000   00:00:00.000
+ Drain Lock Wait                          0  00:00:00.000   00:00:00.000
+ Drain of Claims Wait                     0  00:00:00.000   00:00:00.000
+ Archive Log Mode(Quiesce) Wait           0  00:00:00.000   00:00:00.000
+ Archive Read from Tape Wait              0  00:00:00.000   00:00:00.000
+ Switch to Open/Close Wait                0  00:00:00.000   00:00:00.000
+ Switch to SYSLGRNG Service Wait          0  00:00:00.000   00:00:00.000
+ Switch to DMS Waits                      0  00:00:00.000   00:00:00.000
+ Other Service Waits                      0  00:00:00.000   00:00:00.000
+ Force at Commit Waits                    0  00:00:00.000   00:00:00.000
+ Log Write I/O Wait                       0  00:00:00.000   00:00:00.000
+ Sync EX Unit Sw-com/abort/dealloc        0  00:00:00.000   00:00:00.000
+ ------------------------------------------------------------
+ Total Class 3 Wait Time                    00:00:00.000
```

Elapsed time includes ‘think’ time between calls to DB2

IN-DB2 time represents time executing the calls

CPU times shows general CP time and time on zIIP

No delays in this example

Thread Enclave Information

Enclave token

What service class is this thread executing in?

WLM qualifiers used to select service class

WLM samples

See enclave token, WLM service class, and service class performance index (PI)

```
>                          ENCLAVE DETAIL INFORMATION
 PLAN
+ Thread:  Plan=DISTSERV  Connid=SERVER   Corrid=db2bp.exe    Authid=DNET581
+ Dist   :  Type=DATABASE ACCESS, Luwid=G941491B.FA05.090605140121=294
+ Location  :  9.65.73.27
 enc
+ ENCLAVE TOKEN:      5C0002D6C9          Enclave Type:       Original Indep
+ Owning System:      DEMOMVS             Owning Job:         DSNCDIST
+ WLM Mode:           Goal                Enclave CPU Time    00:00:00.007
+
+ SERVICE PERIOD INFORMATION
+ Period(s) for Service Class DDFDEF:    2
+ Current Period for This Thread:        1
+ Performance Index This Period:            .50
+
+ +------------------------------ Period Number
+ | +---------------------------- Importance
+ | | +-------------------------- Duration in Servic  Units
+ | | |     +-------------------- Percentile
+ | | |     |     +-------------- Response Time Goal or Velocity
+ | | |     |     |     +-------- Response Time Unit
+ | | |     |     |     |         Goal Description
+ - - +---- +--- +---- +-------- --------------------
+ 1 3   500  80%  2500 MilliSecs Percentile Response Time Goa
+ 2 4              .20           Velocity Goal
+
+ SERVICE CLASS INFORMATION
+ CPU Critical:       No                  Storage Protection: No
+
+                     Name                Description
+                     --------            ------------------------------
+ Service Class:      DDFDEF              DDF Default
+ Workload:           DB_WKL              database workloads
+ Resource Group:
+ Report Class:       RDDFDEF             default for ddf
+
+ WLM APPLICATION ENVIRONMENT
+ Application Environment Name:           NO WLM ENVIRONMENT
+ Description:
+ Subsystem Type:
+ WLM Started Task Procedure Name:
+ Start Parameters:
+
+
+ CLASSIFICATION WORK QUALIFIERS
+ Subsystem Type:     DDF                 Correlation:        DB2BP.EXE
+ Proc Name:                              Trans Program Name:
+ UserId:             DNET581             Transaction Class:
+ Network ID:                             Logical Unit Name:
+ Plan Name:          DISTSERV            Package Name:       SQLC2F0A
+ Connection:         SERVER              Collection:         NULLID
+ Function Name:      DB2_DRDA            Subsystem Name:     DSNC
+ Accounting Info:    SQL09013NT          Subsystem Parm:     DNET581
+ Perform:                               Subsystem Priority: N/A
+ Scheduling Env:                        Subsys Coll Name:
+ Process Name:       DB2BP.EXE
+
+ Performance Index Input Data for Percentile Response Time Goal
+ Observations:             14            Percent             Count of Transacti
+                                         ---------           ------------------
+ Goal Percentile This Observation --->          50%              165707
+                                               60%                  62
```

Thread Reuse Complicates The Analysis

```
 _________________ ZALLT     VTM      O2        V510./C DSNA S 02/23/12 13:27:15
>.LOGPOP    >> LOG status restored to Inactive. <<
>.LOC                                                          >> Done <<
> Help PF1       Back PF3        Up PF7      Down PF8      Sort PF10     Zoom PF11
> T.A
>          Thread Activity:  Enter a selection letter on the top line.

> *-All-Idle     B-TSO       C-CICS
> G-Dist DBAC    H-Util      I-Inact
> M-Triggers     N-Sysplex   O-Enclaves
=====================================
>                     Threads Summary Excluding Idle Threads
 THDA
+ *
+ Elapsed       Planname   CPU     Status        GetPg  Update Commit CORRID/JOBN
+ ----------    --------   -----   ------------  ------ ------ ------ ------------
+ 02-21:14      DISTSERV   00.0%   WAIT-REMREQ     3410   1048    104 db2jcc_appli
+ 02-04:07      DISTSERV   00.0%   WAIT-REMREQ    33766K 136231  13600 db2jcc_appli
+ 02-04:07      DISTSERV   00.0%   WAIT-REMREQ    93873  16054   1032 db2jcc_appli
+ 02-04:06      DISTSERV   00.0%   WAIT-REMREQ    59668  11031    637 db2jcc_appli
+ 02-04:06      DISTSERV   00.0%   WAIT-REMREQ    26845K 31927   2843 db2jcc_appli
+ 02-04:06      DISTSERV   00.0%   WAIT-REMREQ      145M 475626  52175 db2jcc_appli
                           00.0%   WAIT-REMREQ    38639K 158876  15815 db2jcc_appli
                           00.0%   WAIT-REMREQ    21644K 21825   2696 db2jcc_appli
                           00.0%   WAIT-REMREQ    69752    369     99 db2jcc_appli
                           00.0%   WAIT-REMREQ    23463    178     43 db2jcc_appli
                           00.0%   WAIT-REMREQ    13805    215    116 db2jcc_appli
                           00.0%   WAIT-REMREQ      736      8      5 db2jcc_appli
```

With thread resue getpage, In-DB2 times, CPU times may all reflect multiple transaction executions

Many distributed/enclave based applications will employ thread reuse for efficiency

```
>                                THREAD DETAIL
 PLAN
+ Thread:  Plan=DISTSERV  Connid=SERVER   Corrid=db2jcc_appli Authid=JAZZ
+ Dist   :  Type=DATABASE ACCESS, Luwid=G9274493.A2B4.C928B3BEE858=993
+ Location  :  9.39.68.147,Host Name=DEMOMVS
 act
+ Thread Activity                          User Defined Functions
+ ----------------------------------------  ------------------------------------
+ DB2 Status             =  WAIT-REMREQ     TCB Time (SQL)          = 00:00:00.000
+ MVS Status             =                  Wait for TCB Time       = 00:00:00.000
+ Total Elapsed Time     =  03-00:28        Elapsed Time            = 00:00:00.000
+ CP CPU Utilization     =         00.0%    Elapsed Time (SQL)      = 00:00:00.000
+ Total CP CPU Time      =          N/A     SQL Events              =             0
+ IIP CPU Time           =  00:00:00.000
+ Total Parallel Tasks   =             0
+ Current Parallel Tasks=              0
+
+ Stored Procedures                         Triggers
+ ----------------------------------------  ------------------------------------
+ Total CPU              =  00:00:00.000    TCB not in Enclave      = 00:00:00.000
+ Elapsed time           =  00:00:00.000    Elapsed not in Enclave  = 00:00:00.000
+ IIP CPU Time           =  00:00:00.000
+ Total Parallel Tasks   =             0
+ Current Parallel Tasks=              0
+
+ Stored Procedures                         Triggers
+ ----------------------------------------  ------------------------------------
+ Total CPU              =  00:00:00.000    TCB not in Enclave      = 00:00:00.000
+ Elapsed time           =  00:00:00.000    Elapsed not in Enclave  = 00:00:00.000
+ Elapsed Time (SQL)     =  00:00:00.000    TCB prior to Enclave    = 00:00:00.000
+ Wait for TCB Time      =  00:00:00.000
+ Wait Event Count       =             0
+ Curr Wait TCB Time     =  00:00:00.000
+
+ SavePoints
+ ----------------------------------------
+ Savepoint Requests     =             0
+ Release Savepoints     =             0
+ Rollback Savepoints    =             0
+
+ In-DB2 Times                                 Total          Current
+ ------------------------------               -------------  -------------
+ Elapsed Time                                 00:02:47.604   00:00:00.000
+ CP CPU Time                                  00:01:00.160   00:00:00.000
+ IIP CPU Time                                 00:01:26.226            N/A
+ Stored Procedure CPU Time                    00:00:00.000   00:00:00.000
+ UDF CP CPU Time                              00:00:00.000
+ UDF IIP CPU Time                             00:00:00.000
+ UDF Elapsed Time Main                        00:00
+
```

Elapsed time

In- DB2 time

In-DB2 CPU time

Thread Reuse Complicates The Analysis - continued

A large difference between Total Elapsed time and In-DB2 times indicate thread reuse. Consider these numbers when analyzing DB2 accounting data.

DB2 DDF Threads
WLM And Enclave Considerations

- DB2 thread options may influence enclave creation and how DB2 interacts with WLM
  - Impacted by such things as KEEPDYNAMIC options, cursor with hold
  - Enclave creation may drive using velocity versus response time goals

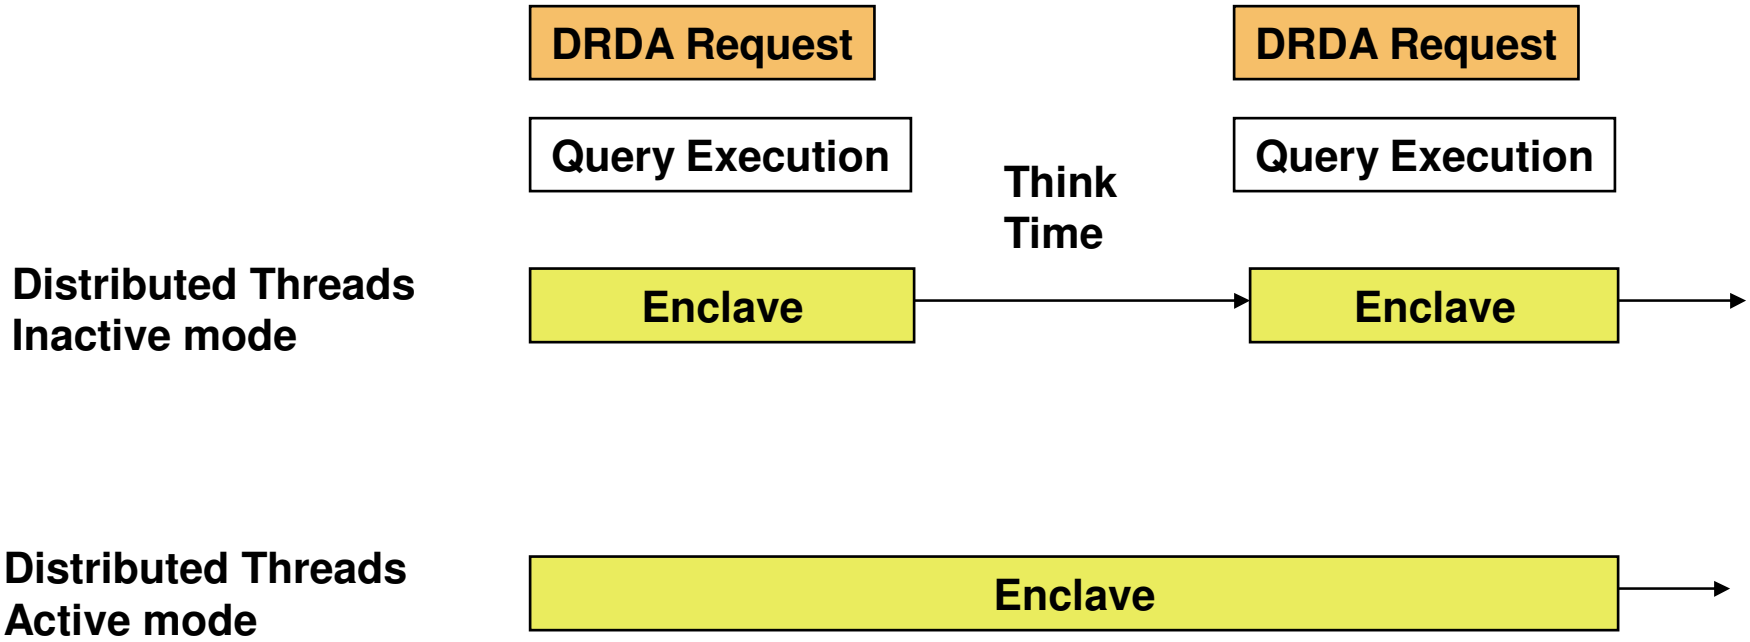

DB2 Accounting Data Considerations

- DB2 provides options to control frequency and granularity of accounting record creation
  - ACCUMACC controls whether and when DB2 accounting data is accumulated by the user for DDF and RRSAF threads
    - Data accumulated for specified # of threads
    - Turned on if ZPARM ACCUMACC > 1
    - How it is summarized is based upon ACCUMUID setting
      - ACCUMID may be set as a combination of user id, workstation id, transaction id, etc…
- Rollup of accounting information can be useful for reducing the amount of SMF data created
  - Summarized information may be limited for problem investigation
  - Summarized information may hide the effects of problem thread in the rollup
  - Note – ACCUMAC and ACCUMID may be changed online

WLM
Service Classes Categorize Workload

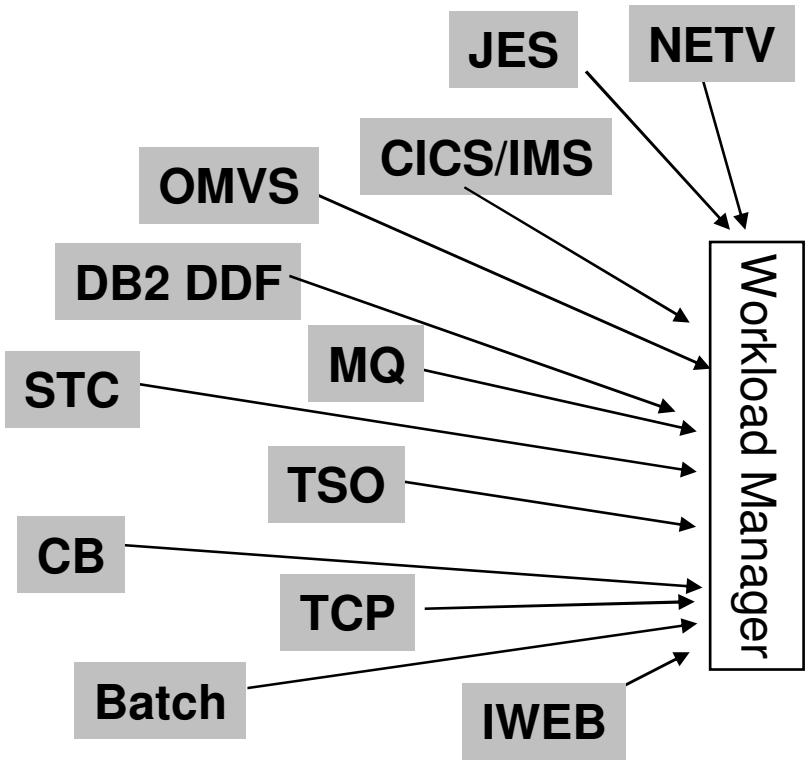

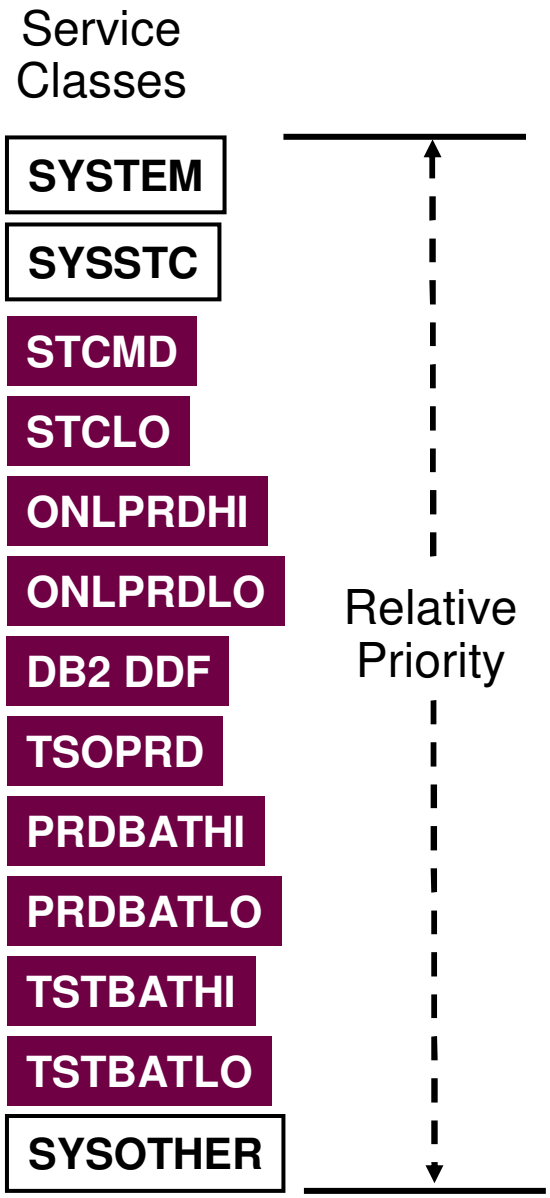

- Classification rules assign incoming work to the appropriate WLM Service Class
- Classification rules group together logically related work

An Example
Providing Attributes To DB2

```
!DIS THD(*)
DSNV401I  ! DISPLAY THREAD REPORT FOLLOWS -
DSNV402I  ! ACTIVE THREADS -  241
NAME     ST A    REQ ID           AUTHID   PLAN     ASID TOKEN
SERVER   RA *     33 javaw.exe    DNET581  DISTSERV 008D   282
 V437-WORKSTATION=Ed@IBM, USERID=dnet581,
      APPLICATION NAME=javaw.exe
 V445-G941491B.E205.090605130859=282 ACCESSING DATA FOR 9.65.73.27
CICSACB7 N        3              SYSSTC            009E     0
CICSACB6 N        3              SYSSTC            009D     0
```

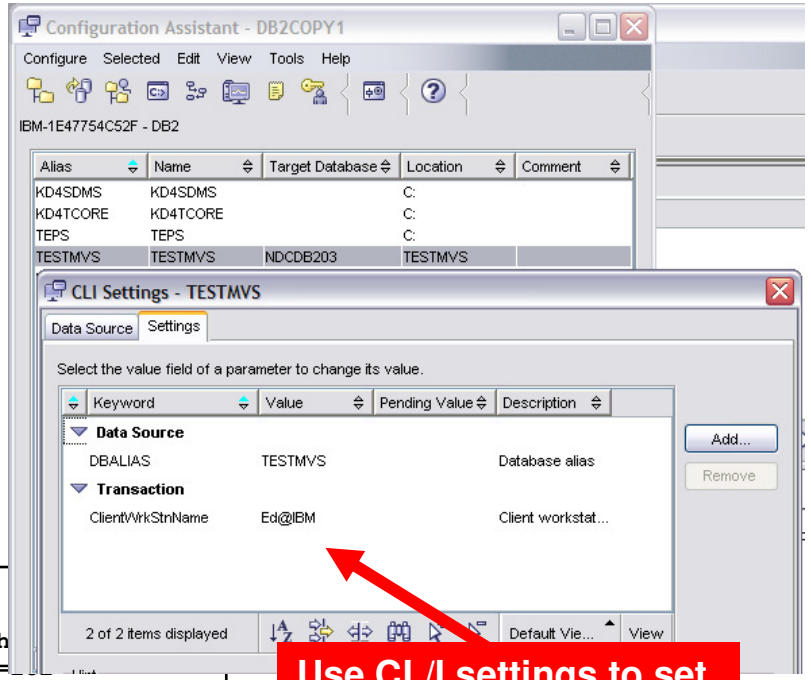


```
>                        DISTRIBUTED THREAD DETAIL
 PLAN
+ Thread:  Plan=DISTSERV  Connid=SERVER   Corrid=javaw.exe     Auth
+ Dist   :  Type=DATABASE ACCESS, Luwid=G941491B.E205.090605130859=
+ Location   :  9.65.73.27
 rsum
+                        Distributed TCP/IP Data
+Location          IP Addr  Port Ctbuser  Srvclsnam  Prod ID  Workstation Nam
+----------------- -------- ---- -------- ---------- -------- ----------------
+9.65.73.27        0941491B  448 dnet581  NT         SQL09013 Ed@IBM
+
+ Transaction name: javaw.exe
+ TCP/IP Userid:    dnet581
+
+                         Distributed SQL Statistics
+
+ Remote Location Name = 9.65.73.27       Remote Location Luname =
```

Use CL/I settings to set the workstation name

Attributes may be passed that may in turn be used by WLM classification

Classifying The Workload
Specifying The WLM Objectives

```
File  Utilities  Notes  Options  Help
-------------------------------------------------------------------------
Functionality LEVEL011          Definition Menu          WLM Appl LEVEL013
Command ===> ____________________________________________________________

Definition data set  . . : none

Definition name  . . . . . DEMO       (Required)
Description  . . . . . . . No overrides

Select one of the
following options. . . . . ___  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
```

All the various WLM constructs on z/OS are defined using the WLM ISPF dialogs

- WLM provides an ISPF interface to define and manage the WLM Service Definition
- Note – z/OS Management Facility V1.12 provided a new management interface

Workload Manager As A DB2 Priority Mechanism

```
  Subsystem-Type  Xref  Notes  Options  Help
------------------------------------------------------------------------------
                Create Rules for the Subsystem Type          Row 1 to 5 of 5

Subsystem Type . . . . . . . . DDF     (Required)
Description  . . . . . . . . . Example DB2
Fold qualifier names?  . . . . Y   (Y or N)

Enter one or more action codes: A=After  B=Before  C=Copy  D=Delete
M=Move I=Insert rule IS=Insert Sub-rule  R=Repeat

          -------Qualifier-------------              -------Class--------
Action    Type        Name      Start                Service      Report
                                                DEFAULTS: PRDBATCH     _________
 ____  1  SI          DB2A      ___                  PRDBATCH     _________
 ____  2    CN          ONLINE    ___                PRDONLIN     _________
 ____  2    UI          SYSADM    ___                PRDONLIN     _________
 ____  2    PK          QMFOS2    ___                TSTQUERY     _________
 ____  1  SI          DB2B      ___                  TESTUSER     _________
```

Examples of Thread Attributes

- **AI (Accounting Information)**
- **CI (Correlation Information)**
- **CN (Collection Name)**
- **CT (Collection Type)**
- **LU (LU Name)**
- **NET (Net ID)**
- **PK (Package Name)**
- **PN (Plan Name)**
- **SI (Subsystem Instance)**
- **UI (Userid)**

Thread attributes in WLM allow for considerable granularity in the classification of DB2 workloads into the appropriate Service Class

Exploit the granularity to prioritize higher versus lower importance workload

DB2 Enclave Workloads
Setting Optimal Goals

- Use Response Time goals when possible
  - Less need for ongoing maintenance and review
  - WLM will manage resources dynamically to achieve goals
- Response Time goals work well for certain types of DB2 workloads
  - DB2 Distributed workloads in e-business and WebSphere transactional type workloads
  - Transactional type workloads in general including distributed workloads that invoke Stored Procedures
  - Repetitive workloads that have multiple events for WLM to measure and manage
- Use a Velocity goal for the DB2 DDF address space
  - DDF address space has internal tasks that govern thread creation that should have high performance goal

Setting WLM Goals Things To Note

- Considerations for DDF threads
  - For DDF inactive threads
    - Consider a two-period service class with a response time goal where 80-90% of the transactions complete in first period
  - For DDF always active threads
    - Consider velocity goals and use a single-period service class
- Look for overly simplistic Service Class definitions
  - Example – type DDF and nothing more than DB2 subsystem name
    - Does little to exploit the ability of WLM to prioritize DB2 workloads
    - Some workloads will inherently be more important than others
- Look for workloads that run longer than expected but use less resource than anticipated
  - Indicative of workload that may not being optimally classified
- Avoid too many service classes/periods
  - WLM analyzes service classes/periods in a round-robin manner
  - Too many and WLM is unable to manage them all effectively
  - Consider WLM reporting classes for report/analysis granularity and detail

Summary

- Enclaves are a pervasive mechanism for z/OS workload priority management
- DB2 is one of the primary exploiters of enclaves
- Many functions of DB2 exploit enclaves
  - DDF workload, Stored procedures, sequential prefetch and deferred writes
- Effective analysis of DB2 enclave based workload requires an understanding on the interaction of DB2 and z/OS enclaves
  - Understand WLM service classes, reporting classes and how they are defined in your environment
- Setting optimal WLM goals and priorities for DB2 requires an understanding DB2 and enclaves
  - Understand application flow and logic
  - Response time versus velocity goals

Thank You!

Check Out My Blog

http://tivoliwithaz.blogspot.com