



ACTIVATE BUSINESS WITH THE POWER OF I.T.™




Ins & Outs of Distributed Processing in DB2 on z/OS

Bill Arledge
BMC Software, Inc.




Overview



- › Distributed Processing Fundamentals
- › Distributed threads – What can I see?
- › WLM, enclaves and SRBs
- › DDF and DB2 system considerations
- › Impact on critical resources

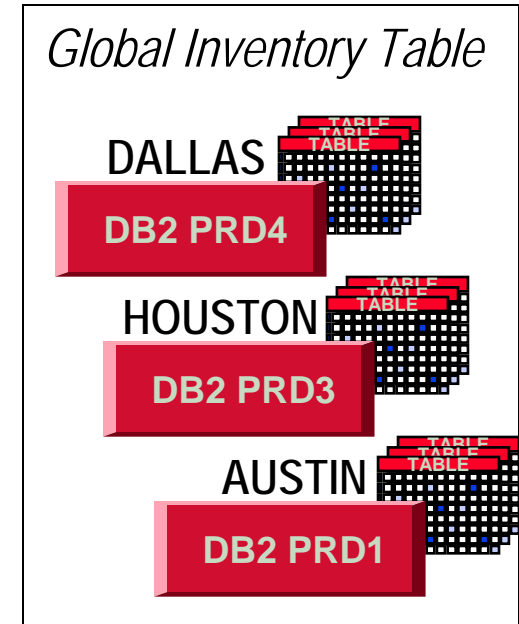
Basic Concepts – Intro

What is “Distributed”



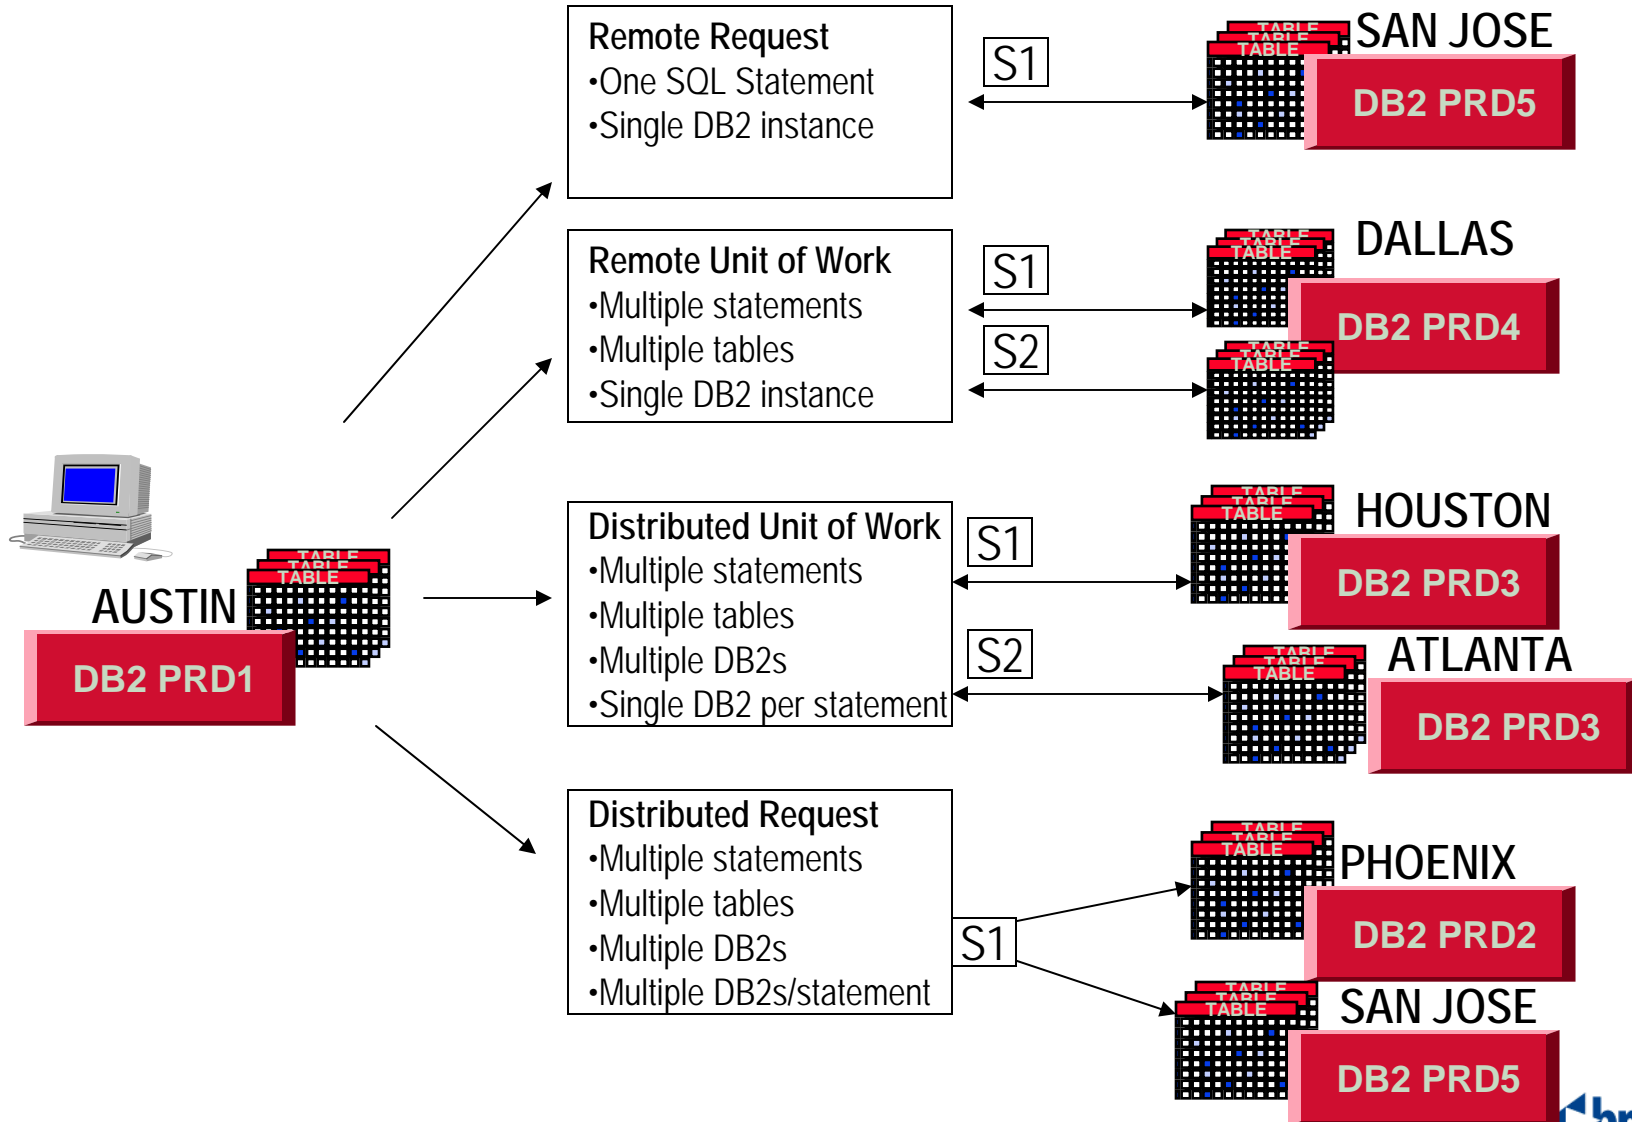

If not further qualified:

- › Could mean a distributed platform or database
 - UNIX, AIX, Windows
 - Oracle, DB2 LUW, and others
- › Could mean distributed DB2 tables
 - Portions of data for a single object distributed across multiple DB2s, probably geographically dispersed
 - Part of early relational model
 - Data distributed vertically or horizontally
 - Select from GLOBAL_INVENTORY_TABLE
 - Not directly supported by DB2 on z/OS
- › Probably means distributed processing
 - DB2 as a database server
 - Responding to requester on multiple platforms




Distributed Processing

Basic Concepts – Units of Work



Distributed Fundamentals

SQL Access



› System Directed Access

- Three Part Names
- Alias
- Supported by DRDA and PP

```
SELECT * FROM PRD1.RNDWDA.ORDER
        WHERE QUANTITY BETWEEN 1 and 100
or...
CREATE ALIAS RNDWDA.AUSTIN_ORDERS
        FOR PRD1.RNDWDA.ORDER
SELECT * FROM RNDWDA.AUSTIN_ORDERS
```

› Application Directed Access

- Explicit CONNECT by application
- Supported by DRDA only
- Remote BIND required

```
EXEC SQL CONNECT TO PRD1
SELECT * FROM RNDWDA.ORDER
```

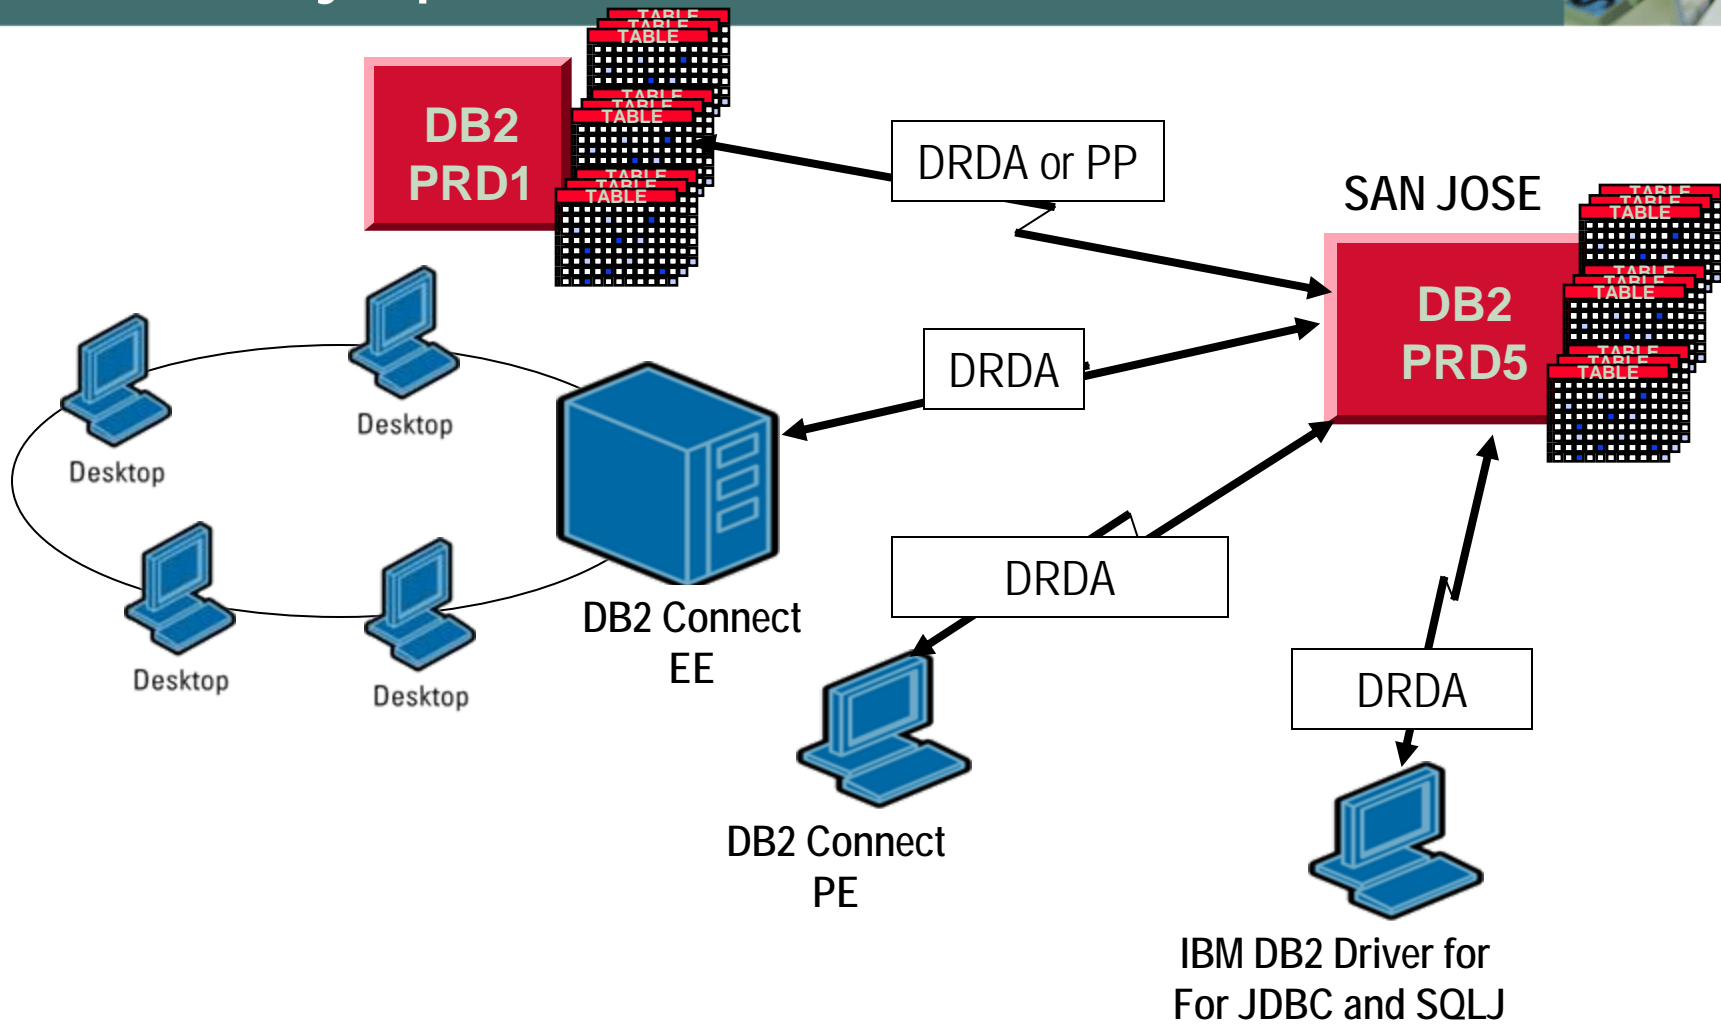
› Remote Stored Procedure Call

- Explicit CONNECT by application
- DRDA only


```
EXEC SQL CONNECT TO PRD1
EXEC SQL CALL PROCONE
```

```
PROCEDURE PROCONE
EXEC SQL SELECT.....
EXEC SQL UPDATE.....
```

Distributed Fundamentals Connectivity Options




Basic Terminology Definitions



- › Location (DB2 for z/OS term)
 - Or: RDB-Name, VTAM nodes, TCP/IP partners
- › Connection – between a requester and a server
 - TCP/IP ports, or VTAM LUNAMES
- › Network protocol – TCP/IP or SNA (VTAM)
- › Conversation – handle traffic on a connection
 - Also referred to as a session
 - DRDA – one per requester to handle SQL & open cursors
 - Private protocol – may have more, one per open cursor

Distributed Fundamentals

A Tale of Two Protocols



› Private Protocol (PP)

- First delivered in DB2 2.2
- Uses 3 part names
 - Select C1,C2,C3 from DB2G.APPL1.TABLE1
- Dynamic SQL only
 - No remote bind
- DB2 on z/OS to DB2 on z/OS only
- No stored procedure support
- Functionally stabilized and on the way out

Specifying the protocol


- Default can be set at subsystem level using DBPROTCL DSNZPARAM value (DRDA or PRIVATE)
 - With DB2 9.1 PRIVATE can no longer be specified
- Can also be specified in the DBPROTOCOL parm of the BIND statement (D or P)
 - Only option in DB2 9

› DRDA (Distributed Relational Database Architecture)

- Introduced in DB2 2.3
- Supports 3-part names and explicit CONNECT statements
- Dynamic and Static SQL
 - Remote bind capabilities
- Supports stored procedures
- Supports all RDBMS implemented using DRDA protocol
- Supports SNA and TCP/IP
- Is the strategic architecture for distributed

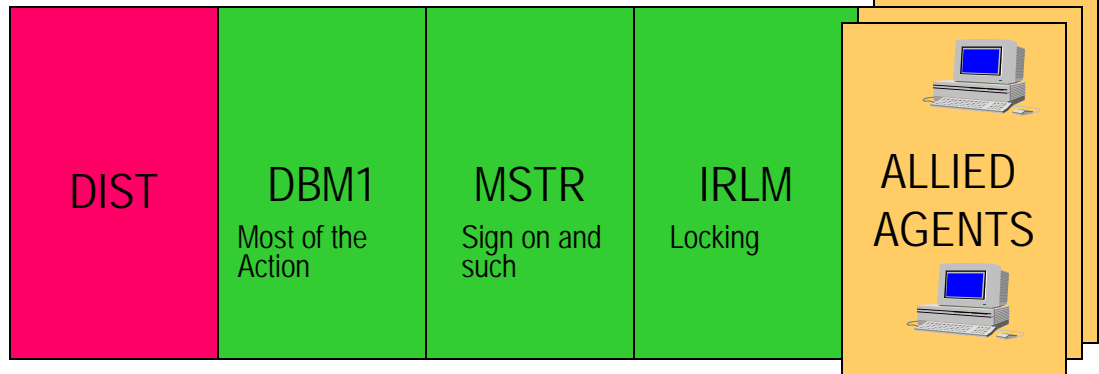
Distributed Fundamentals

DB2 on z/OS Distributed Implementation



Along came distributed

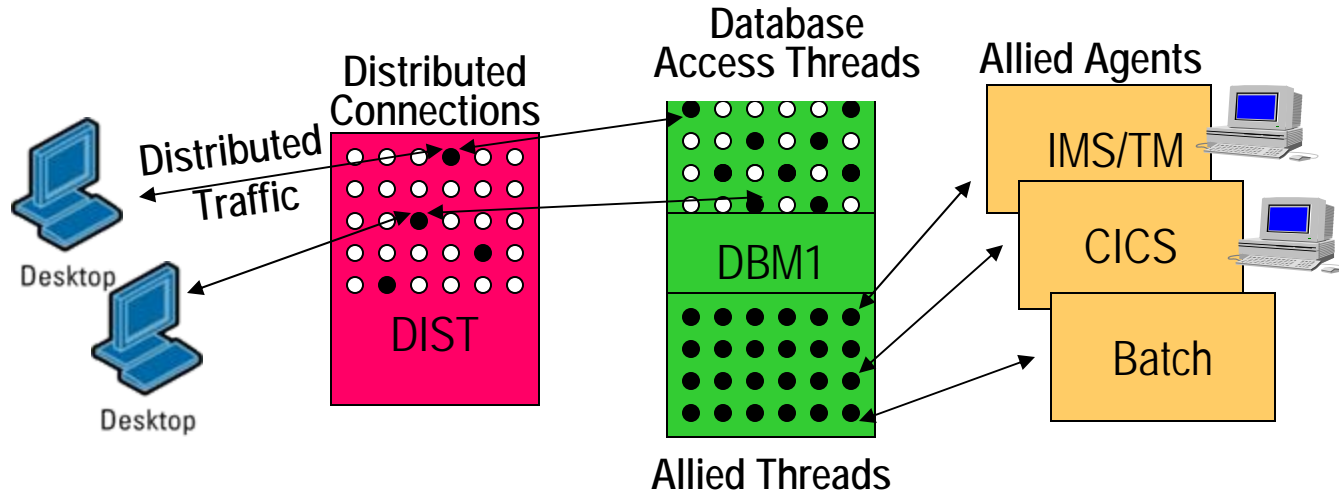

In the beginning



- › In the beginning - 3 DB2 operational address spaces
 - Plus all the allied agent address spaces
 - CICS, IMS, TSO Attach
 - TSO Batch, Call Attach Facility
- › Distributed Data Facility (DDF) in DB2 V2R2
 - Access using 3 part names or aliases
 - DIST address space introduced
- › DRDA (Distributed Relational Database Architecture) first implemented in DB2 V2R3
- › Major enhancements delivered in DB2 V4
 - DRDA support of stored procedures
 - DBAT user priority
- › More in DB2 V5
 - TCP/IP, ODBC, CLI, JDBC
 - Much more . . .
- › Web-based access comes of age
 - Java, JDBC Universal Driver, Websphere . . .

DB2 on z/OS Distributed Implementation

A Word About Threads




Database Access Threads (DBATs)

- › Service distributed workloads
- › Implemented as an MVS WLM enclave running in preemptive SRBs originating in the DIST address space (more coming)
- › DBAT Types
 - DBAT (Server)
 - DBAT (Dist)

Allied Threads


- › Service local attachment facilities
- › Run at the dispatching priority of the requesting application
- › Can become distributed requesters
- › Allied Agent Types
 - Allied
 - Allied Dist (requester)

DDF Setup Considerations Related ZPARM Review



- › CMTSTAT – DDF Threads
- › IDTHTOIN – Idle Thread Timeout
- › TCPKPALV – TCP/IP Keepalive
- › POOLINAC – Pool Thread Timeout
- › ACCUMACC and ACCUMUID
- › MAXTYPE1 (PP) – Max Inactive DBATs
- › KEEP DYNAMIC(YES) / MAXKEEPD
- › EXTRAREQ / SRV – Extra Blocks REQ / SRV
- › And of course:
 - MAXDBAT – Max Remote Active
 - CONDBAT – Max Remote Connected

DDF ZPARM View



W1 =ZPDDFD=====DECE=====*****=27FEB2007==19:23:12====MUDB2==

DDF - Dist Data Facility Definitions.....

Local Location.....	DECE
DDF Startup Facility Name.....	DDF
DDF Start Option.....(DDF).....	AUTO
Database Protocol for 3-Part Names.....(DBPROTCL).....	DRDA
DDF Max Number of Facility Entries.....	1
DBAT Status after Commit.....(CMTSTAT).....	INACTIVE
Idle Thread Timeout (Seconds).....(IDTHTOIN).....	1200
Minutes between Resync Periods.....(RESYNC).....	2
TCP/IP KEEPALIVE.....(TCPKPALU).....	ENABLE
DDF Interval Cycle Frequency.....(SPRMINT).....	120
DDF Queued Conversation Time.....(SPRMQCT).....	120
DDF Receive Buffer Size.....(SPRMDRB).....	30720
Max Extra DRDA Query Blocks for DB2 Req.....(EXTRAREQ).....	100
Max Extra DRDA Query Blocks for DB2 Svr.....(EXTRASRV).....	100
Check Connection State.....(PKGLDTOL).....	n/a

DBAT Thread Controls.....


Max Concurrent Database Access Threads.....(CONDBAT).....	10000
Maximum Remote Database Access Threads.....(MAXDBAT).....	3
Maximum Type 1 Inactive Threads.....(MAXTYPE1).....	10
DDF Pool Thread Timeout Value.....(POOLINAC).....	1200

DDF-Related Authorization.....

Extended Security.....(EXTSEC).....	N
ID Sent to Second Server.....(HOPAUTH).....	BOTH
Accept Already Verified TCP/IP Connects.....(TCPALVER).....	Y
DDF RLF Access Error Parameter.....(RLFERRD).....	NOLIMIT
DDF RLF Service Unit Limit.....(RLFERRD).....	0

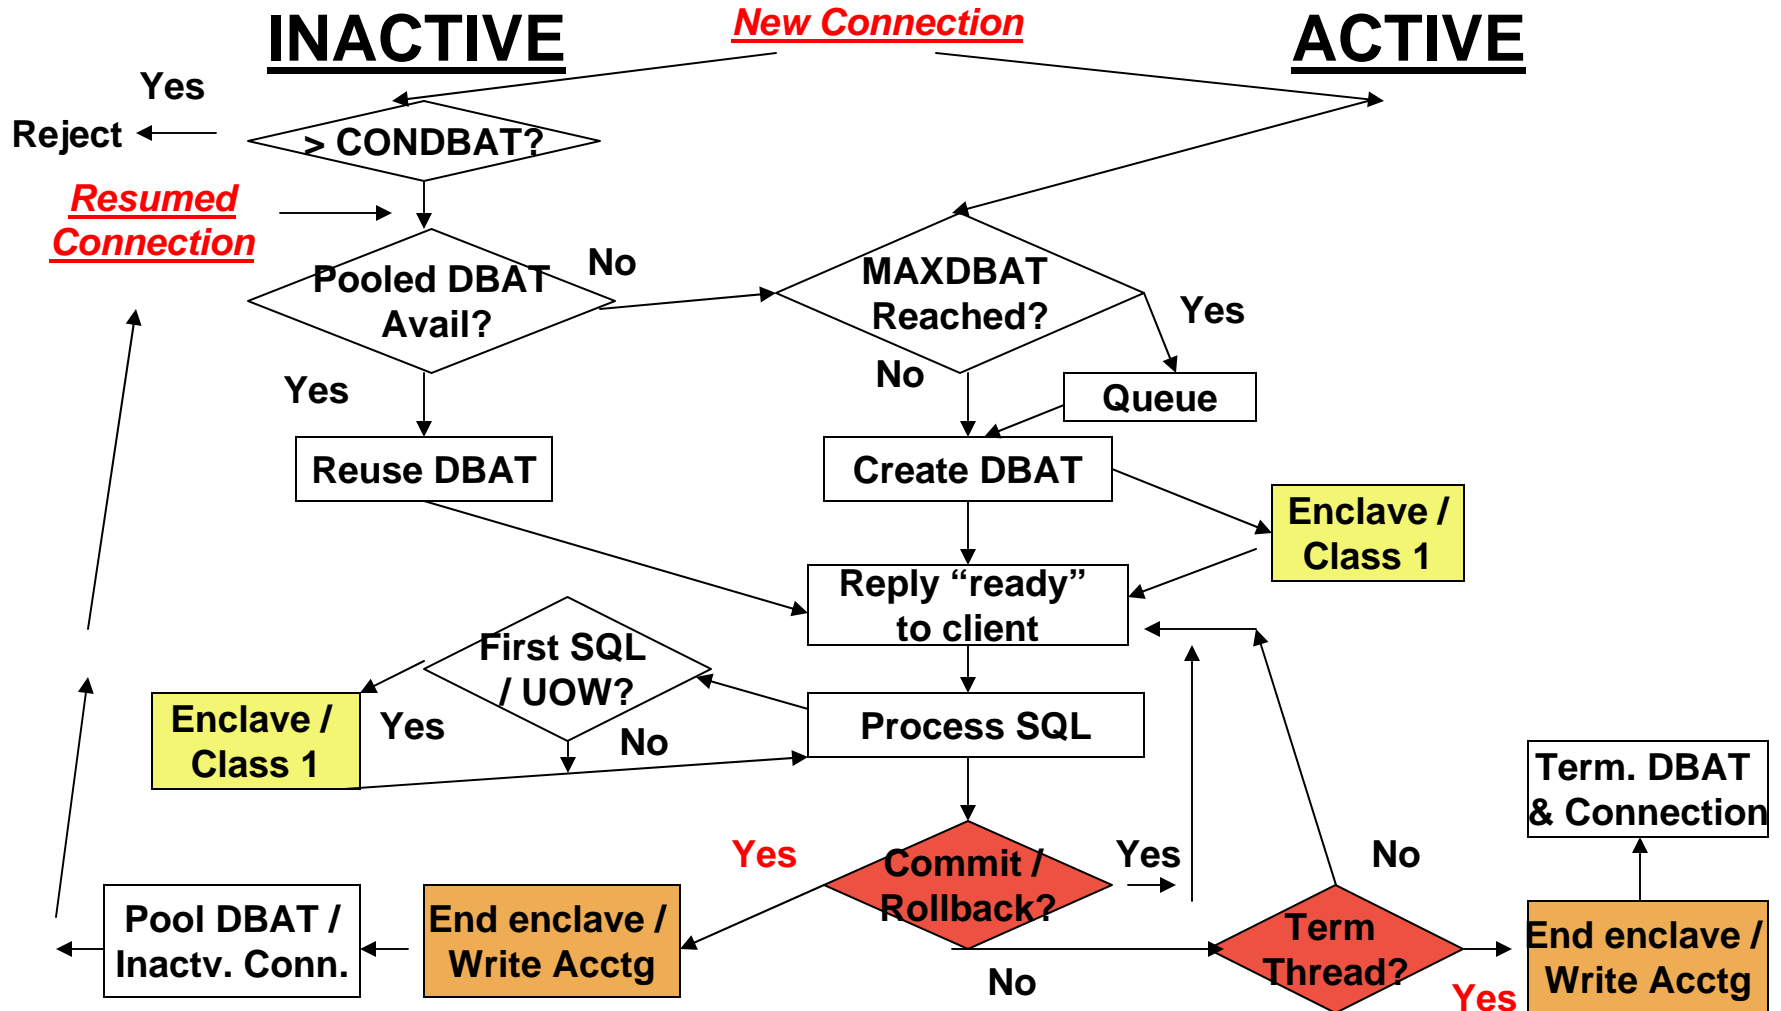

. NEXT

DBAT Processing Modes




- › Mode is defined with the ZPARM CMTSTAT
 - Very critical option for DDF that defines the flow of DBAT processing
 - “DDF Threads” on panel DSNTIPR
- › Two choices:
 - **INACTIVE** – highly recommended
 - Provides DBAT pooling for DRDA access
 - More effective WLM classification per UOW
 - Reduced Resource usage
 - **ACTIVE**
 - DBAT created for each new client application
 - DBAT held through commits
 - Use this only if the applications require it
- › ZPARM Considerations
 - CMTSTAT INACTIVE / ACTIVE
 - POOLINAC Used / not applicable
 - IDTHTOIN Used / not recommended
 - CONDBAT:MAXDBAT > / =

Processing Diagram



WLM Enclaves



- › WLM Enclaves are independent, dispatchable units-of-work that span multiple address spaces and can include combinations of SRBs and TCBs
 - Concept of “Business Unit-of-Work”
- › DB2 use enclaves for work coming into the system through DDF
 - Controlled by WLM
 - Can run on zIIP processors
- › Thread priority set by WLM workload classification
 - Providing good DDF classifications is vital
- › Enclave completes = accounting data is ready
 - Defines class 1 elapsed times of a thread
 - (Not affected by rollup option)
 - Enclave metrics externalized in SMF type 72 record


```
Subsystem Type . : DDF          Fold qualifier names?  Y  (Y or N)
Description . . . : DB2 DDF

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

Action  -----Qualifier-----  -----Class-----
Action  Type      Name      Start      Service      Report
-----  ---      ---      ---      ---          ---
_____  1      UI      DB*      _____  _____
_____  1      SSC     DIA      _____  _____
_____  1      PN      CWRPLAN  _____  _____
_____  1      CN      CWR*     _____  _____

          DEFAULTS: DDF
                   DDF
                   DDF
                   DDF
                   DDF
```

WLM Enclaves




› INACTIVE mode

- No end user “think time” included
- Enclave is created when the first SQL is received
- Enclave is deleted at commit / rollback (thread complete)
- New enclave for each UOW, reclassified by WLM
- Can use multi-period response time or velocity goals

› ACTIVE mode


- End user “think time” is included
- Enclave is created when the DBAT is created
- Enclave is only deleted at thread termination
- Only one enclave, no reclassification
- Can only use a single-period velocity goal

DBAT Thread identifiers - Basic



- › Connection Type **** WLM CT ****
 - DRDA or Private Protocol
- › Other IDs for DB2 to DB2 work (DRDA or PP)
 - All come from the remote requester thread
 - Even with a “hop”, they come from the requester
- › Other IDs for non-z/OS DRDA clients
 - Two unique identifiers
 - Connection Name = “SERVER”
 - Plan = “DISTSERV” **** WLM PN ****

More Identifiers from non-z/OS Clients



› Clients can flow other identifiers to DB2 for z/OS


– Values can be set by DDF threads via “Server Connect” and “Set Client” calls

- ODBC/CLI/VB (SQLSetConnectionAttr)
- Non-ODBC (sqleseti)
- JDBC (DB2Connection)
- DRDA (ACCRDB prddta / sqlstt in EXCSQLSET)

› Most important IDs supported in V8 with special registers

- Client Accounting (see QMDA below)
- Workstation Userid **** WLM SPM 1-16 ****
- Workstation Name **** WLM SPM 17-34 ****
- Workstation Application **** WLM PC 1-32 ****


Other Differences – DRDA Clients



- › **Package / Collection** ** WLM CN/PK **
 - First package accessed
- › **Stored procedure name** ** WLM PR **
 - If First SQL is a CALL
- › **AUTHID of client** ** WLM UI **
 - Often not unique for non-z/OS clients
- › **Original primary AUTHID**
 - Used to make initial connection to server
- › **Correlation ID** ** WLM CI **
 - DDM external name (EXTNAME) for client
- › **Accounting correlation token**
 - 22-byte token



More “Accounting” Information



- › **Special section for thread “accounting” data**
 - Used for additional client identification
 - Only in the accounting record IFCID 03
- › **Product ID - shows the client source product**
 - SQL – DB2 for LUW / DB2 Connect
 - JCC – Universal JDBC Driver
 - DSN – DB2 for z/OS requester
- › **DSN accounting string (z/OS)**
 - A repeat of the QWHC identifiers, except:
 - MVS accounting string (QMDAACCT)

Non-z/OS Accounting IDs




› SQL or JCC Accounting

- Client platform
- Client application name
- Client AUTHID of an application process
- Accounting String **** WLM AI ****

› Also, IDs for the DB2 for z/OS server processing the work


- Subsystem instance **** WLM SI ****
- Subsystem collection name
(Data sharing group) **** WLM SSC ****
- Sysplex name **** WLM PX ****

DBATs and Accounting



- › **ACTIVE mode**
 - Only cut at thread termination, not at commit
- › **INACTIVE mode**
 - DRDA – at “clean” COMMIT or ROLLBACK
 - “Type 2 inactive”
 - DRDA with KEEP DYNAMIC(YES)
 - At “clean” commit (DB2 V8 and above)
 - PP DBAT – at commit or termination
 - At commit, if “Type 1 Inactive” (MAXTYPE1) allowed
 - Else only at termination
- › **Active thread is idle too long and is canceled**
 - At “Idle Thread Timeout” (IDTHTOIN), if allowed
 - Checked every 2 minutes


Accounting and DDF Rollup



- › Option in DB2 V8 to reduce accounting volume
 - Turned on if ZPARM ACCUMACC > 1
 - Just going to V8 the ACCUMACC parm defaults to 10 which means it's on
 - You could see a big drop in SMF 101 records
- › Data accumulated for specified # of threads
 - For matching IDs, based on ACCUMUID
 - Combination of the 3 workstation IDs
 - Total of 10 combinations in V8
 - 7 new options in DB2 9
- › Accounting written when
 - “Too old” (staleness threshold)
 - “Too much” (internal storage threshold reached)
 - “Just enough” (limit threshold reached)
- › One accounting record reflects one or more threads
 - Currently no DDF statistics (QLAC) or QMDA accounting
 - Only one “ROLLUP” package
- › Active thread data only shows the current thread counts


Connection and Thread Processing

- Review of INACTIVE Mode




- › (1) A new connection (in DIST) is established
- › (2) DB2 attempts to allocate a DBAT
 - Use a pooled DBAT if possible
 - Allocate a new DBAT if possible (expensive)
 - Queue if MAXDBAT reached (RQ)
 - DBAT shows as pooled until SQL is received (DA)
- › (3) UOW processes SQL (RA)
 - Idle thread timeout can cause it to be canceled
- › (4) “Clean” commit or rollback completes the UOW
 - Frees the DBAT to be pooled, connection goes inactive (R2)
 - KEEP DYNAMIC(YES) keeps the DBAT until termination
- › (5) New SQL “resumes” the connection and a new UOW
- › (6) Disconnect frees the connection

"Real" DBAT Thread Status (#1)




- › Assigned to a remote client (RA or RX)
 - Actively processing executing SQL
 - Active but idle waiting for more SQL
 - Waiting for more work after “clean” commit, if:
 - INACTIVE mode – only:
 - KEEPDYNAMIC(YES) – all resources & DBAT kept
 - Type 1 inactive – PP only / some resources freed
 - ACTIVE mode – even after commit
 - All resources & DBAT kept until thread termination
 - Suspended to connect (PP only, temporary) (RN)

“Real” DBAT Thread Status (#2)



- › Pooled (DA)
- › DRDA clients only, with INACTIVE mode
 - Freed or newly created DBATs are pooled
 - Also referred to as “DBAT slots”
- › Available for reuse by any new / resumed request
 - (Still somewhat in “standby” for previous client)
- › Still uses resources (esp. DBM1 storage)
 - 200K in DBM1
 - Occasionally terminated to free storage
- › Still shown and counted as “active threads”
 - But connection name is “DISCONN”
 - Can be terminated if not used (POOLINAC)


Where are the Inactive Type 2 DBATs?



- › They are referenced often in various manuals
 - Pooled DBATs? **Not DBATs at all !**
- › **Actually, they are the inactive connections**
 - Associated with a remote requester
 - Waiting for more work
 - This speeds up response to additional SQL
 - Tracked in DIST, and use less storage (7.5K)
- › **Shown only with DIS THREAD TYPE(INACTIVE)**
 - Connection name is now "SERVER"
 - "Thread" status **(R2)**




And Inactive Type 1 DBATs?



- › These are real DBATs
- › Idle between UOWs
- › Only Private Protocol
 - Old style of inactive processing
- › The DBAT is still assigned
 - But resources are reduced
- › This can only occur if $\text{MAXTYPE1} > 0$
 - And limit is not reached

Understanding Thread Status



› Active thread displays


- Show both assigned and pooled DBATs
- Even though pooled DBATs aren't really "active"

› Inactive thread displays

- Show the inactive connections in DIST
 - While still "associated" with a pooled DBAT
 - Looks like the same requester is both active and inactive
- When pooled DBAT is terminated or reassigned
 - The requester "disappears" from active
 - Still shows as inactive until connection terminated



Conversation Processing



- › Conversations are used for actual traffic on a connection between two remote partners
- › When processing, the conversation is
 - Shown under the active thread
- › Otherwise, the conversation is
 - Shown under the inactive connection
 - After the initial connection until the first SQL
 - After a successful commit

Conversation Processing

Display DBATs and associated conversations



› Here is an example of a DBAT that has just committed its UOW.

› DISPLAY ACTIVE

DSNV402I *DECE ACTIVE THREADS -


› NAME	ST	A	REQ	ID	AUTHID	PLAN	ASID	TOKEN
› DISCONN	DA	*	3	NONE	NONE	DISTSERV	00CA	13491
› V471-USBMCN01.LUDECE.C018343A2C5F=13491								

› DISPLAY INACTIVE

› DSNV424I *DECE INACTIVE THREADS -


› NAME	ST	A	REQ	ID	AUTHID	PLAN	ASID	TOKEN
› SERVER	R2		0	db2bp.exe	BOLJX01	DISTSERV	00CA	13502
› V437-WORKSTATION=joott-SJC-70 USERID=boljxo1								
› APPLICATION NAME=db2bp.exe								
› V445-AC173AFE.F70D.01F581235657=13502 ACCESSING DATA FOR								
› 172.23.58.254								
› V447--LOCATION				SESSID	A	ST	TIME	
› V448--172.23.58.254				7623:3575	W	R2	0703216095818	

Viewing Active Threads



- › Assigned DBATs are identified with SERVER
- › Pooled DBATs with DISCONN
 - Only the number is interesting (see statistics)
- › Extra DDF activity counts
- › Data sharing considerations
 - Various routing mechanisms across members
 - Need a group view of DBATs
 - To see complete distributed workload
 - Evaluate DBAT's across data sharing groups if possible

Active DBATs (Data Sharing Members)




```
>W1 =THDDBAT===== (DBGK=====*) 01MAR2007==13:24:15====MUDB2====U====6
CMD Correlation DB2 Package Elapsed DDF Msgs DDF Msgs Total
--- Id ID Name Time CPU Time Sent Received SQL
db2bp.exe DB2K 00:00:00.00 00:00:00.00 0 0 0
javaw.exe DB2K SYSSH200 00:08:32.48 00:00:00.01 1 1 214
db2bp.exe DB2K 00:00:00.00 00:00:00.00 0 0 0
db2bp.exe DB1K 00:01:08.62 00:00:00.00 2 2 0
javaw.exe DB1K SYSSH200 00:01:07.22 00:00:00.00 3 3 123
db2bp.exe DB1K SQLC2F0A 00:04:02.51 00:00:06.59 5 5 3
```

Customize with workstation IDs, other values as needed
Scroll right to see more fields

```
+W1 =THDDBAT===== (DBGK=====*) 01MAR2007==13:24:15====MUDB2====U====6
CMD Correlation Workstation Remote Enclave
--- Id Name Location Token Auth ID
db2bp.exe JBARTHEL-HOU-98 172.18.60.221 0000000000000000 DMRQA01
javaw.exe JBARTHEL-HOU-98 172.18.60.221 0000002400000007 DMRQA01
db2bp.exe JBARTHEL-HOU-98 172.18.60.221 0000000000000000 DMRQA01
db2bp.exe JBARTHEL-HOU-00 172.18.61.200 0000000000000000 DMRQA02
javaw.exe JBARTHEL-HOU-00 172.18.61.200 0000002C0000000B DMRQA02
db2bp.exe dwitkows-SJC-04 172.23.59.211 0000002800000009 BOLDJW1
```

Hyperlink ↑


Enclave views (MVzOS)



```
>W1 =WMENCLUZ=====SYSBDEMO=*=====05APR2007==14:18:45====MUMUS====D====1
Enclave Token      Service  Owner      Cumulative Exectn  Total  Total  %Idle  %
----- Class      Jobname    CPU Time  Velcty  Dly%   Use%  ----- Unkn
0000002400000087 DDF       DHN1DIST  00:00:22.1  30.00  70.00  30.00
```

```
>W1 =WMENCLAS=WMJINFO==SYSBDEMO=*=====05APR2007==14:19:31====MUMUS====D====1
Timeframe... Interval 0....50..100
Jobname..... ENCLAVE Sysplex Name INTLPLEX Total Use%.. 25.00
Type..... ENCLAVE System Name. SYSB %Use CPU.... 25.00
Serv. Class. DDF SMF I.D..... SYSB %Use DASD... 0.00
Rept. Class. Enclave Cnt. 0 Total Dly... 75.00
ASID..... 0 Velocity.... 25.00 %Dly CPU.... 75.00
Dmn..... 0 Velocity 2.. 25.00 %Dly DASD... 0.00
Period No... 2 Using Sampls 3 %Dly Stor... 0.00
Workload.... DDF Delay Sampls 9 %Dly Srvr... 0.00
Resource.... MPL Delays.. 0 %Dly MPL.... 0.00
Trxn RPGN... 0 Swpin Delays 0 %Dly Swpin. 0.00
Userid RPGN. 0 Idle Samples 0 %Dly Quiesce 0.00
TrxC RPGN... 0 Unk. Delays. 0 %Idle... 0.00
Acct RPGN... 0 Sample Count 12 %Dly Unknown 0.00
Status..... Active Us/Del Count 12
```


Checking Client Connections



- › Many will be inactive connections
 - Shown as inactive threads (with client IDs)
- › Some have active DBATs
 - In-flight accounting data is available
- › Conversation is with DBAT or inactive connection and shows:
 - Whether the conversation is active in the network or suspended in DB2 waiting for a response
 - Last send/receive time stamp
 - Whether it is receiving or sending
 - The remote location (IP address) and “Sessid” - local and partner ports (for TCP/IP)




Inactive Thread (Connection) View



```
>W1 =THDINACT===== (DBGK=====*)=====) 01MAR2007==13:22:53====MUDB2====D====2
Connect  Current      Correlation      Plan          LUW          Workstatio
Name     Activity    Id              Auth ID      Name         ASID   Token   Name
SERVER  Inactive  DBAT db2bp.exe      DMRQA01     DISTSERV   273    3467   JBARTHEL-H
SERVER  Inactive  DBAT db2bp.exe      DMRQA01     DISTSERV   273    3493   JBARTHEL-H
```

Connection / Conversation Views




```
>W1 =DDFLOC=====DEDM=====*****02MAR2007==12:09:43====MUDB2====U====4
DB2                               Product  Tot  Req  Serv  Inact  Tot
Target  Remote Location  Link Name  ID      Conn Conn Conn  Conn  Cnv
DEDM    ::172.17.8.86                1      0    1     0     1
DEDM    ::172.21.22.183             4      0    4     1     4
DEDM    DECE                    ::172.17.8.86  DSN08015  2      2    0     0     2
DEDM    DHH                      LUDHH2        DSN08015  2      2    0     0     2
```

Hyperlink on “Tot Conn” to see details

```
>W1 =DDFLOC===DDFTHD===DEDM=====*****02MAR2007==12:14:17====MUDB2====U====4
DB2      Age      Cnv Workst      Correlation  Latest      Wor
Target  Typ  Status  Cnt  User ID      ID          Send/Receive Time  Nam
DEDM    R/S  Active  2  boljxo1  db2bp.exe  2007.061  11:14:30.89  joo
DEDM    R/S  Active  2  boljxo1  db2bp.exe  2007.061  11:13:48.06  joo
DEDM    Srv  Active  1  boljxo   db2bp.exe  2007.061  11:12:53.16  joo
DEDM    Srv  Idle   1  boljxo   db2bp.exe  2007.061  11:11:46.91  joo
```



Analyzing DDF Thread Data




- › The accounting data is the first source
- › Still analyze other application considerations
 - Elapsed and CPU times, I/O, SQL counts . . .
- › But in addition:
 - Elapsed time inside / outside the DB2 server
 - Number of messages and blocks sent / received
- › Batch reports summarized by
 - The important DDF identifiers for your workloads



Thread Accounting – DBAT Specifics

Summary Information



```

STOP.....21MAR 04.01.14.60 PLAN.....DISTSERV TYPE.....DBAT
START....21MAR 04.01.10.99 AUTHID.....APPDEMO CONNECT.....SERVER/DRDA
ELAPSED.....3,605 ms ORIG PRIM AUTH.....APPDEMO CORR ID....db2jcc_appli
TERM.....NORMAL/DEALLOC COMMITS......0 ROLLBACKS.....1
-----
RUNTIME ANALYSIS      IN DB2      IN APPL.      TOTAL      %IN DB2(=)      TOTAL(*)
-----
ELAPSED TIME          807 ms     2,798 ms     3,605 ms     0 ..25...50...75..100%
CPU TIME              83 ms     1,099 us     84 ms
DB2 WAIT TIME        620 ms
ZIIP CPU TIME         0 us           0 us
ZIIP-ELIGIBLE CP          46 ms
  
```


Distributed Indicators/
Identifiers

- > DBAT Accounting Metrics
 - "Application" processing is occurring in the workstation client
 - Class 2 – Class 1 CPU measures time not processing SQL in DB2
 - Includes parts of thread creation and termination and moving data from and to communication buffers
 - CMTSTAT ACTIVE is one explanation for the elongated IN APPL elapsed time
- > zIIP Metrics
 - Actual CPU time on the zIIP
 - Not included in "standard" class 1 & 2 CPU times
 - CPU eligible for the zIIP



Thread Accounting – DBAT Specifics

DDF Summary Information



```

----- DDF SUMMARY -----
CURRENT LOCATION: DIA (DBAT SERVER)
WORKSTATION USER ID...AFDQA User WORKSTATION NAME..AFDQA Workstation
WORKSTATION TRANSACTION ID...AFDQA TEST appdemo
MVS WLM SERVICE CLASS NAME..... DDF
Client = DB2 Universal JDBC Driver. / Vers. 02 Rel. 10 Mod. 0
Client Platform = AFDQA Workstation
Job accounting data = AFDQA CLIENT ACCOUNTING012345678910111213141
51617181920212223242526272829303132333435363738394041424344454647484950515253545
REMOTE FROM
LOCATION: ::192.168.7.5 SQL STATEMENTS 8 n/a
WAIT ON MAX DBAT ....0 us
REMOTE PROD-ID: JCC02100 ROWS n/a 1,526
DB AGENT CPU.....n/a BYTES 1,771 83,337
ELAPSED LOCAL.....0 us TRANSACTIONS 1 n/a
ELAPSED REMOTE.....n/a MESSAGES 12 12
CONVERSATIONS QUEUED.....n/a CONVERSATIONS 1 n/a
BLOCK FETCH %.....100% COMMITS 0 n/a
MAX CONVERSATIONS.....n/a ABORTS 0 n/a
UNSUCCESSFUL CONV'S.....n/a BLOCK FETCH n/a 5


```

Client Identifiers Available in DB2 V8

WLM Service Class

- > DBAT Details
 - Wait on MAX DBATs
 - PROD-ID = IBM Universal Driver for JDBC and SQLJ
 - Traffic Information


Tracing Distributed Workloads



- › Additional focus on one workload
 - Summary exception trace (accounting)
 - Detail trace with important event IFCIDs
- › All the usual qualifiers are available
- › For DDF, important to reduce the data:
 - Filter by requesting location
 - Filter by Workstation ID(s)
 - In V9, DB2 also allows qualification by these IDs
- › Exception Filters can be used to keep only threads that may need analysis (high In-DB2 elapsed, etc.)




Detail Traces



- › Detail traces can include selected event groups
 - Basic thread flow and SQL
 - Also can choose to add scans, I/O, locks
- › Another group to include specific DDF events
 - The volume can be high
 - Use it only when needed
 - To understand the conversation flow
- › Each event has a pop-up view with the IFCID details




DDF Statistics



- › The next place to look are the statistics
- › Global statistics
 - Critical DB2 subsystem tuning information
- › Location statistics
 - Application impact on DB2 and network
 - DRDA_Remote_Locs (combined)
 - Private Protocol locations (separate)
- › DDF Address Space CPU usage
 - TCB and SRB

Global DDF Statistics - STDISTD



W1 =STDISTD=====DECE=====*****01MAR2007==13:33:48===MUDB2

DBAT Statistics Detail.....

	Interval	Session
Maximums Reached.....		
New DBATs Queued (MAXDBAT).....	0	5
Conversations Deallocated (CONDBAT).....	0	0
New/Resumed (Type 2) DBATs Queued (MAXDBAT)	1	40
Connections Terminated (MAXTYPE1).....	0	0
Status Values.....		
Remote Connections - Maximum.....		10
Active DBATs - Current.....		3
- Maximum.....		3
DBAT Slots Not Used - Current.....		0
- Maximum.....		1
Type 1 Inactive DBATs - Current.....		0
- Maximum.....		1
Type 2 Inactive DBATs - Current.....		6
- Maximum.....		6
Type 2 Queued (New/Resumed) - Current.....		2
- Maximum.....		4
Two-Phase Commit Activity.....		
Cold Start Connections.....	0	0
Warm Start Connections.....	0	0
Resync Attempts.....	0	0
Resync Succeeds.....	0	0
Statistics.....		
Requests that Required a DBAT.....		4
Requests that Used a Pool Thread.....		72

STDISTD View - Revised



```
W1 =STDISTD=====DECE=====*****=01MAR2007==13:36:33====MUDB2==
DDF Global Statistics Detail.....
. DDF ZPARAMS.....
Status - Current and High Water Mark.....      Current      HWM
Total DBATs - Active & Pooled.....              3             3
DBATs Pooled for Reuse (Type 2).....             0             1
Inactive DBATs (Type 1).....                    0             1

Total Remote Connections.....                    10
Type 2 Inactive Connections.....                 6             6
Type 2 Connections Queued for DBAT.....          3             4


Maximums Reached.....      Interval      Session
Queued for DBAT (MAXDBAT Reached).....           1             6
Connections Deallocated (CONDBAT Reached).....   0             0
Type 1 Connections Terminated (MAXTYPE1 Reached) 0             0

DBAT Usage Statistics.....      Interval      Session
New DBATs Created.....                          4             4
Pooled DBATs Reused.....                         72            72
New/Resumed (Type 2) Requests.....               2            41

Two-Phase Commit Activity.....      Interval      Session
Cold Start Connections.....                     0             0
Warm Start Connections.....                     0             0
Resync Attempts.....                             0             0
Resync Succeeds.....                             0             0
Resync Failures.....                             0             0
```




Exception Monitoring



- › Review your current exceptions
 - Are DDF conditions being monitored?
- › Statistics
 - DBAT high water mark
 - Queuing for a DBAT?
 - DDF still active?
 - DBM1 storage usage
- › Accounting
 - Focus on DDF service levels
 - Filter for DBATs / most important work
 - Elapsed time / CPU usage



DDF Resource Usage



- › CPU – TCB and SRB
 - In the DIST address space
 - Management of the DBATs and connections
 - For the threads themselves (enclave SRBs)
- › DBM1 storage (MAXDBAT, and CTHREAD)
 - Management of thread storage is critical
- › DIST address space
 - Storage likely not an issue (CONDBAT)
- › Dynamic SQL cache
 - Most distributed SQL is still dynamic
 - The cache is critical for good performance
 - Aim for an 80% or better hit ratio for SQL reuse

DBM1 Storage – DB2STORD View



```
W1 =DB2STORD=====DECE=====*****=27FEB2007==19:25:03
DBM1 And MUS Storage Usage..... Quantity in MB

DBM1 Storage Summary Below 2 GB.....
Storage Available to DBM1..... 1515.32
Total DBM1 Storage In Use (1)..... 58.27
Total MUS Storage In Use Below 2GB (5).. 106.67
Storage Cushion (4)..... 118.61
Average Thread Footprint..... 0.84
Maximum Number of Possible Threads..... 1611

1) Total DBM1 Storage In Use..... 58.27
Total Getmaind Storage (2)..... 32.68
Total Variable Storage (3)..... 17.70
Total Fixed Storage..... 0.11
Total Getmaind Stack Storage..... 7.78

3) Total Variable Storage..... 17.70
Total Agent Local Storage..... 15.19
Total Agent System Storage..... 14.12
  Number of PreFetch Engines..... 135
  Number of Deferred Write Engines... 15
  Number of Castout Engines..... 0
  Number of GBP Write Engines..... 0
  Number of P-Lock/Notify Exit engines 0
Total Agent Non-System Storage..... 1.07
Total Number Of Active User Threads.. 1
Total Number of Active DBATs..... 1
Thread High Water Mark..... 9
DBAT High Water Mark..... 3
```



Dynamic Cache – STCACHED View



```
W1 =STCACHED=====DECE=====*=27FEB2007==19:31:43
```

Dynamic SQL Cache Details.....

. SQL Cache Statement Analysis.....

SQL Cache in Statement Pool.....

Total Pages.....		25600
Pages Used.....	0.07	17
Free Pages.....	99.93	25583

Global Cache Usage..... Interval Session

Requests.....	0	8
Inserts.....	0	6
Found in Cache(Short Prepare).....	0	2.0
Not Found in Cache(Long Prepare)....	0	6
Global Cache Hit Ratio.....	0.0	25.0
Failures - Data Space Full.....	n/a	n/a
Failures - Statement Pool Full.....	0	0




Local Cache Effectiveness..... Interval Session

Avoided PREPARE (Match).....	0	0
Implicit PREPARE (No Match).....	0	0
Local Cache Hit Ratio.....	0.0	0.0
Statement Discarded (>MAXKEEPD)....	0	0
Statement Purged (Drop/Alter/Revoke)	0	0






Extended Reporting



- › Distributed workloads are often volatile
 - Less insight and control
- › Can be useful to track activity over time
 - Store and query summary data in DB2 tables
- › When needed, distributed traces and monitoring
- › z/OS reporting on WLM can be helpful
 - Enclaves – SMF 30
 - Workloads by service class – SMF 72
- › MVzOS provides online views as well as reports



Questions?