



# IBM DB2 for z/OS Distributed Access Best Practices and Updates

**Adrian Burke**

DB2 SWAT Team SVL

[agburke@us.ibm.com](mailto:agburke@us.ibm.com)



SHARE is an independent volunteer-run information technology association  
that provides **education, professional networking and industry influence.**



# Acknowledgements and Disclaimers



**Availability.** References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

© Copyright IBM Corporation 2014. All rights reserved.

— **U.S. Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.**

— *Please update paragraph below for the particular product or family brand trademarks you mention such as WebSphere, DB2, Maximo, Clearcase, Lotus, etc*

IBM, the IBM logo, [ibm.com](http://ibm.com), [IBM Brand, if trademarked], and [IBM Product, if trademarked] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or TM), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at

“Copyright and trademark information” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

•If you have mentioned trademarks that are not from IBM, please update and add the following lines:[Insert any special 3rd party trademark names/attributions here]

•Other company, product, or service names may be trademarks or service marks of others.

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Agenda

- Driver Levels
- V11 and Newer Features
- T2 vs. T4
- Gateway
- Properties
- High Availability
- Tuning



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# DB2 11 and DB2 Connect



- Any in-service level of DB2 Connect drivers or DB2 Connect server should work with DB2 V11 CM and V11 NFM (V10.5 FP3 [JCC release 3.67/4.17] recommended for all new functions)
  - DB2 Connect in service versions are V9.5 and later (G.A. in 2007)
- Seamless DB2 z migration available with DB2 Connect V9.7 FP6, V10.1 FP2, V10.5 FP 4 or later (true for all current DB2 z releases)
  - DB2 Connect Server is not seamless, application will get connection error and must reconnect
  - Applications will continue to function as members are migrated one by one (without any exceptions seen)
  - While in V11 CM, applications will continue to see V10 function level
    - Driver returns V10R1
  - While fully at V11 NFM, new connections will see V11 function level using APPLCOMPAT
    - Driver returns V11R1

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# What levels have been tested?

- If you want to see what level IBM is testing look at the Consolidated Service Test
  - <http://www-03.ibm.com/systems/z/os/zos/support/servicetest/>
    - <http://www-03.ibm.com/systems/resources/RSU1403.pdf>
  - DB2 Connect Gateway is used between drivers and DB2 z/OS uses the following service levels:
    - DB2 Connect V9.7 Fixpack 9
    - DB2 Connect V10.1 Fixpack 3
- What level of the Driver equals what level of Connect?
  - <http://www-01.ibm.com/support/docview.wss?uid=swg21363866>
- What driver level came through DB2 for z/OS maintenance stream
  - [https://www-304.ibm.com/support/docview.wss?q1=pk77599&rs=0&uid=swg21428742&cs=utf-8&lang=en&loc=en\\_US&cc=us](https://www-304.ibm.com/support/docview.wss?q1=pk77599&rs=0&uid=swg21428742&cs=utf-8&lang=en&loc=en_US&cc=us)

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)





# What level of the driver is out there?

- **-DIS LOCATION DETAIL**
  - DSNL200I – JCC' is the driver, SQL' is DB2 Connect Server
    - *vv* A 2-digit number that identifies the product version number, such as 09 or 10.
    - *rr* A 2-digit number that identifies the product release level, such as 01 or 05.
    - *m* A 1-digit number that identifies the product modification level, such as 0 or 1
  - In the –DIS you can see the IP address it came from, product ID to let you know what level it is at.
    - The ATT column shows if the connection is using Sysplex WLB, XA two phase commit, or encryption (AES/ TLS)

```
JCC03640 <-- 9.7 FP6
SQL09076 <-- 9.7 FP6
JCC03660 <-- 10.5 GA or FP1
JCC04160 <-- 10.5 GA or FP1
```

*Example 2-2 DISPLAY LOCATION report*

DSNL200I	-D0Z2 DISPLAY LOCATION REPORT FOLLOWS-			
LOCATION		PRDID	T	ATT CONNS
:::9.12.4.142		JCC03640	S	50
			WLB	50
			XA	50





## DB2 11 Features (Distributed workload vs. V10)

- Sync Receive CM mode
  - z/OS 1.13 via APAR (PM80004) or z/OS 2.1
- Non-Stored procedures
  - CL1 elapsed time down 3-6%
  - Total CPU time/commit reduced 4-6%
- Stored procedure workloads
  - CL2 elapsed time down 4%
  - Total CPU time/commit reduced 3-5%
  - With implicit commit 50-70% CL1 elapsed improvement
- Implicit Commit for stored procedures
  - Relies on AUTOCOMMIT being on, but bundles commits
  - CM but need Connect V10.5 FP 2

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Cancel thread and Interrupt statement ...



- Pre DB2 11 DDF CANCEL THREAD may not work in some cases
  - Historically "soft", or reactive in nature
  - Cancel detection point may not be encountered in a timely manner
- DB2 11 New FORCE keyword on CANCEL DDF THREAD command
  - Purges thread of a remote connection in the DB2
  - FORCE option accepted only after CANCEL THREAD issued without the FORCE option
  - Available in CM
- Purge running DDF SRB using new z/OS purge SRB function
  - Requires new z/OS CALLRTM TYPE=SRBTERM service
    - Available in z/OS 1.13 or later via APAR OA39392

## Syntax

```
>>-CANCEL ---THREAD(token)-----+-----+-----+-----+-----+-----+-----+-----X  
      +--DDF THREAD(--luwid--)-+ +-DUMP-+ +-LOCAL-+ +-NOBACKOUT-+ +-FORCE-+  
      +-token-+
```

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)





# DBATs in DB2 10



- DB2 10 high performance database access threads
  - Provides additional CPU benefits in DB2 10 – reintroduce RELEASE(DEALLOCATE) to DDF
  - Requires rebinding of DB2 Connect client packages (*SYSLN\*\*\*, SYSLH\*\*\*, etc*)
- Should use the set **currentPackageSet** in driver for 'good' applications to point towards DEALLOCATE packages
- DB2Binder utility defaults to **releasePackageResourcesAtCommit(false)** for DB2 10 packages (9.7 FP 3a) – *so you get DEALLOCATE by default*
- Can be altered on and off with -MODIFY DDF PKGREL command
  - » COMMIT / BINDOPT
    - POOLINAC will determine how long a high performance DBAT remains idle before being destroyed
    - Threads are recycled after 200 uses, but can be pooled if the application continues to issue connection requests (PI20352)

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# High Perf. DBAT Monitoring



- Starting in DB2 10
  - Stats Class 7 record trace (IFCID 365 Remote Location Statistics)
  - Can be used to track rows, bytes, messages and initiated conversations or logical connections between a requestor and a server (IP address)
- Initiated conversations [from remote site] are logical connections
  - High number of conversations deallocated would indicate a poor candidate for HPD

```

DDF DATA BY LOCATION

LOCATION NAME (SHORT).....: ::FFFF:192.0.1.6
LOCATION NAME (LONG).....: ::FFFF:192.0.1.68

INITIATED CONVERSATIONS....:          0
INITIATED FROM REMOTE SITE.:          5
MESSAGES SENT TO REMOTE....:      10815873
SQL STMTS SENT TO REMOTE...:          0
BYTES SENT TO REMOTE.....:      2005208369
ROWS SENT TO REMOTE.....:
BLOCKS TRANSMITTED.....:
    
```

- In the accounting short report compare #DDFS (occurrences) to CONVI (conversations)
  - If > 200 it is an easy fit
  - PI20352 should eliminate most of the overhead if the ratio is < 200:1
  - Using client acct info instead of IP address could make this easier

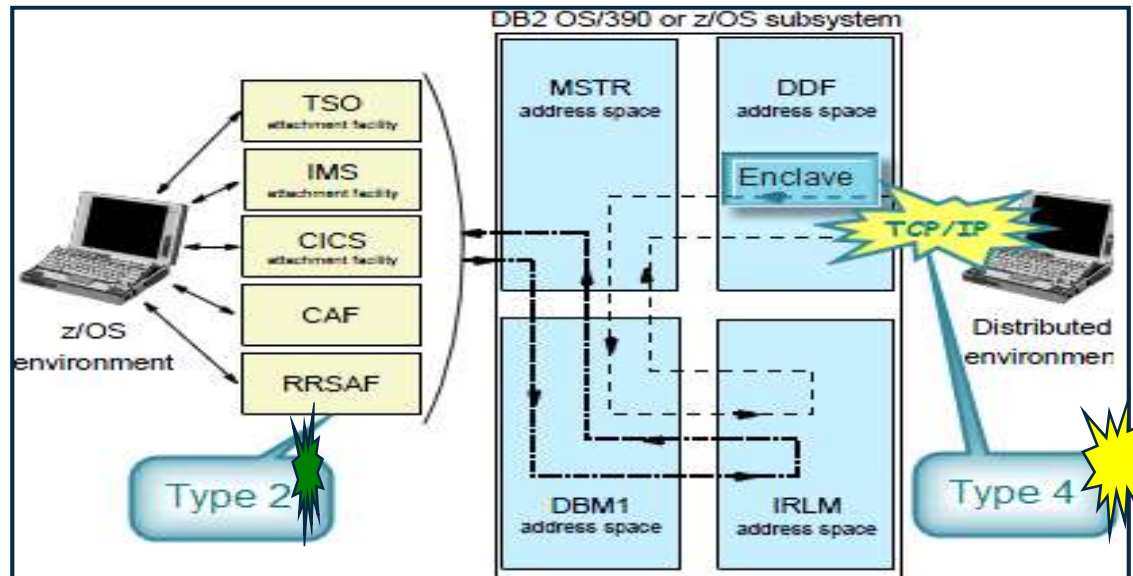
REQUESTER	METH	#DDFS	TRANS	#ROLLBK	#COMMIT	SQLRECV	ROWSENT	CONVI
::FFFF:11.190#04	DRDA	4	N/A	0	6	1.50	88.75	0.25
::FFFF:11.190#05	DRDA	10	N/A	0	12	1.20	36.70	0.10
::FFFF:11.190#07	DRDA	4	N/A	0	6	1.50	27.00	0.25
::FFFF:11.190#08	DRDA	1	N/A	0	1	1.00	1.00	0.00
::FFFF:11.190#09	DRDA	244	N/A	0	246	1.01	3.81	0.00

Complete your session evaluations online at [www.share.com](http://www.share.com)



# Thread footwork T2 vs. T4 Connections same LPAR

- Going through DIST adds network translations, another address space, and context switch to an SRB – but you get zIIP offload
  - Even **local** T4 connection hits DIST address space – need DDF WLM service class
  - Moving T2 to T4 is not 1 for 1 MIP exchange due to overhead
    - But, the worse behaving the application the more there is to offload so in the end it must be tested and compared



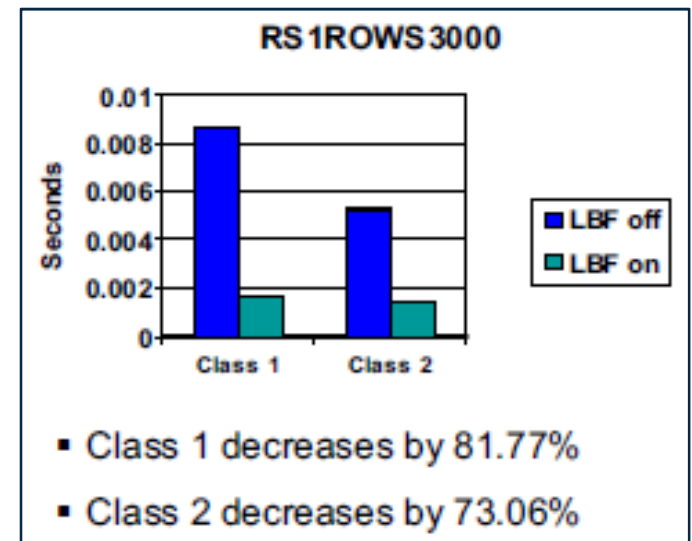
Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

## Cost of execution for type 2 vs. type 4 (on z/OS)



- DB2 10 brought improvements to T2 connections, and DB2 11 expanded them to stored procedures called by T2
  - Result returned using limited block fetch (LBF) and progressive streaming.
  - Result sets now implicitly closed after (SQLCODE +100).
  - CALL and DESCRIBE PROCEDURE statements are bundled into a single message reducing trips between the driver and DBM1.
  - ALLOCATE CURSOR and DESCRIBE CURSORs bundled
- DB2 11:
  - With 50-100 smallish rows T2 saves CL1 CPU time
  - With >100 rows and large rows, or ugly SQL, T4 w/ zIIP has lower CL1 CPU

*T2 Stored Proc and Limited Block Fetch*



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)





# Why Would I use a DB2 Connect Gateway

- Two Phase Commit
  - Many transaction managers still require a DB2 Connect Server for two phase commit if they use a dual transport model
    - For example, Tuxedo, Encina
  - DB2 and MS DTC are two transaction managers that do not require a DB2 Connect Server
- Federation
  - Homogeneous Federation is possible with DB2 Connect Server
    - Use of nicknames to access other DB2 and Informix data servers
- Licensing
  - DB2 Connect Enterprise Edition with concurrent user licensing requires DB2 Connect Server configurations only
  - Ensure you have license for MSUs and Host (based on # of subsystems or D.S. groups)

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# High Availability...not in the Connect Gateway



- Availability is the key advantage to the IBM Data Server Driver
  - Automatic client reroute
  - Transaction level Sysplex Workload Balancing
- Automatic Client Reroute (ACR) - only available in driver
  - Connection to DB2 drops, seamless reroute connection at transaction boundary to other DB2 in data sharing group
  - Next SQL statement re-driven under the covers to another member
    - JAVA applications will not get -30108 during re-connection to another DB2
  - If DB2 crashed then in-flight transaction is rolled back – app needs to re-drive transaction
- Sysplex WLB - transaction level balancing only available in driver
  - Every 10 seconds WLM refreshes list of optional DB2 members with relative weights
  - At transaction boundary next SQL can be routed to another member due to
    - Displaceable CPU on the LPAR
    - DB2 health
    - Enclave WLM Service class goal attainment and queuing
    - MAXCONQN / MAXCONQW

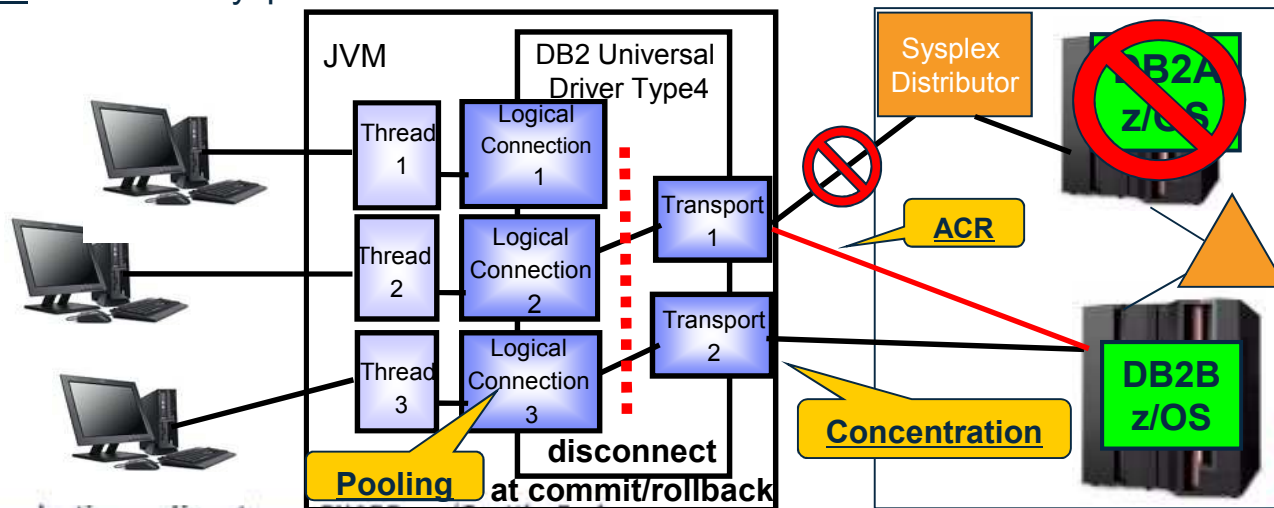
Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Sysplex Workload Balancing...failover



- Availability aspect is arguably more important because if we loose a DB2 or connection, we can re-drive the SQL request across another existing connection to another member of the group without the application's knowledge – Automatic Client Reroute only exists in driver, not Connect Gateway. Gateway will reroute at Connection level
- Even without failure WLM will feed back results so transactions, not just connections are managed across the sysplex



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Key points to Connection Reuse



- Connection pooling allows connections to be reused by other applications, when a Disconnect is issued
- Reusing connections after a commit was historically known as connection concentration
  - Restrictions:
    - If you declare global temporary tables, they must be closed explicitly at transaction or branch boundary
    - SAP does not support Connection Concentrator (KEEPDYNAMIC)
    - Only supports SSL for outbound connections
    - Dynamic prepare requests from embedded dynamic SQL applications will be rejected.
    - See DB2 Connect user guide for details on other restrictions
  - Observations:
    - Most customers use Connection Pooling
    - Thoroughly test usage of Connection Concentrator – often times applications are using some of the above restrictions which do not show up until in production
- If you have enableSysplexWLB=true then enableConnectionConcentrator has no effect on the behavior in the Data Server Driver

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)







## Sysplex WLB properties JAVA driver

- Java datasource properties - V9.7 FP6/ V10.1 FP2/ V10.5 FP4
  - EnableSysplexWLB enables sysplex workload balancing
  - MaxTransportObjects specifies maximum number of connections managed across the data sharing group. The default is 1,000 (previously unlimited)
  - maxRefreshInterval maximum amount of time between refreshes of the client copy of the WLM server list. The default is 10
  - maxTransportObjectIdleTime maximum elapsed time before an idle transport is dropped. The default is 60
  - maxTransportObjectWaitTime time waiting for a transport to become available. The default is 1 second or -4210 (previously unlimited)
  - minTransportObjects the lower limit for the number of transport objects in a global transport object pool. The default is 0

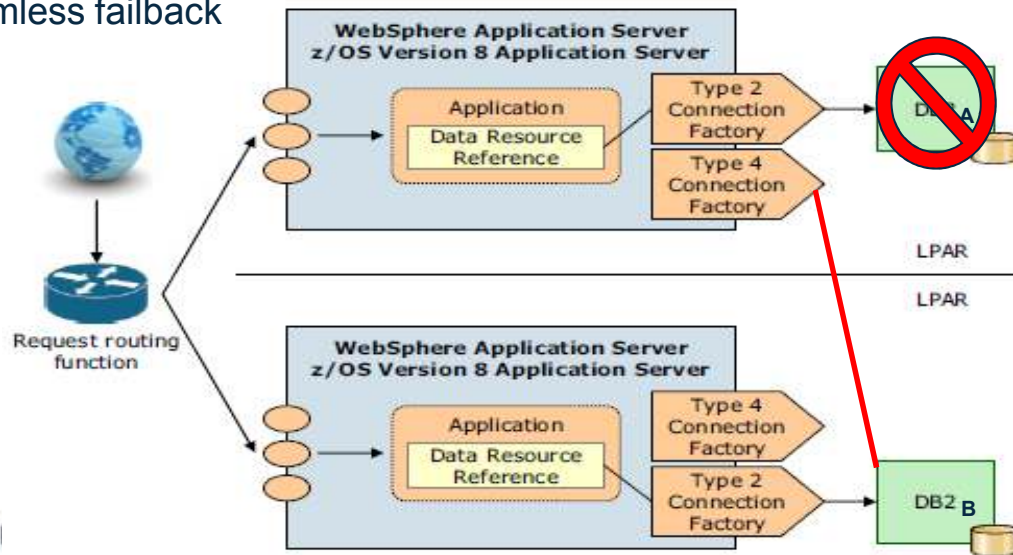
Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# T2 Driver High Availability



- Resource Adapter Failover feature starting in WebSphere Application Server V8
  - If local DB2 goes down it reverts to remote DB2 using T4 attach
    - Could also have another member of same group on that LPAR for T2 attach (specifying group attach)
  - Not seamless failover, applications need to re-drive connection
  - Non-seamless failback



Complete your session evaluation





# Tuning Data Connections

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Key DB2 DDF Parameters

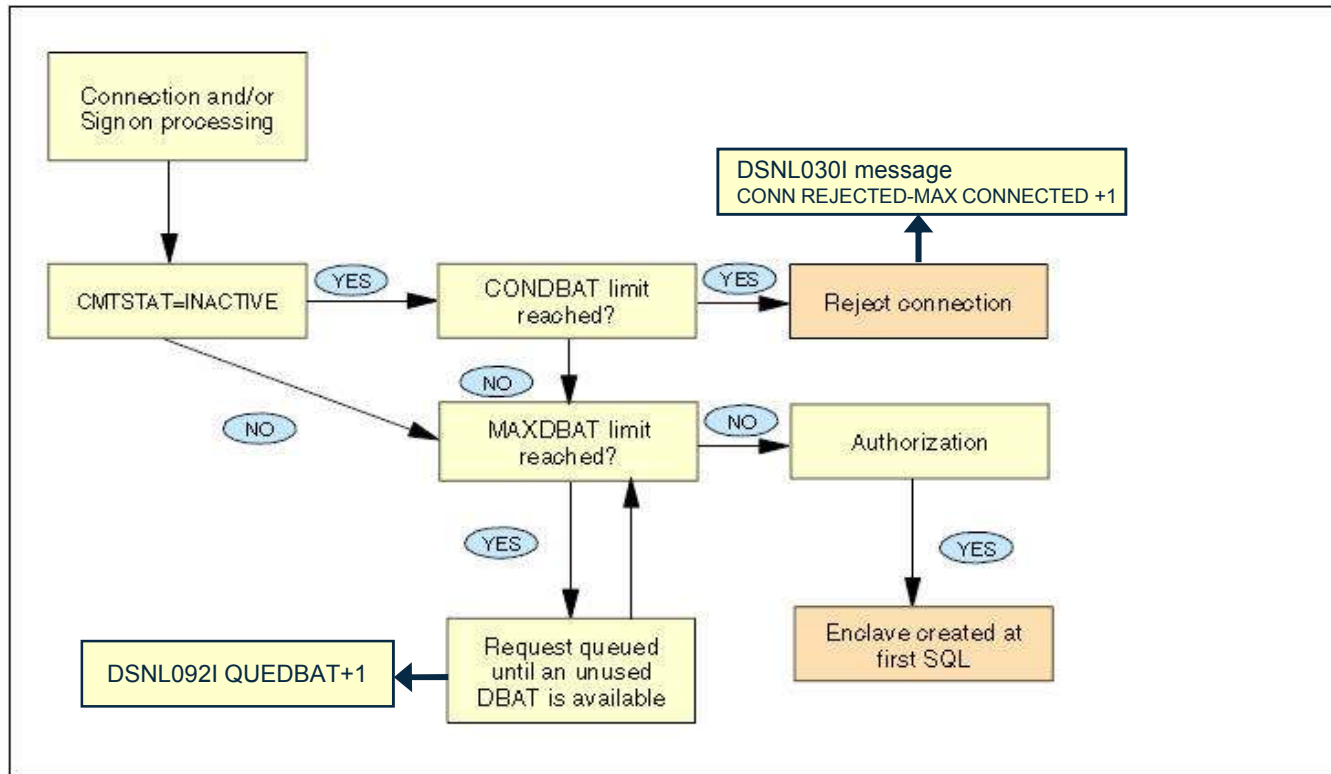


Parameter	Possible values	Default - V8*	Description
DDF	NO / AUTO / COMMAND	No	DDF Startup
CMTSTAT	ACTIVE / INACTIVE	INACTIVE*	Thread Pooling - Pool inactive threads
CTHREAD	1-2,000 (V10 – 10,000)	200*	Max users - allied (local) threads RRSAF or CAF (CICS, IMS, TSO, Batch attach, SPUFI, Classic QMF, etc)
MAXDBAT	0-1,999 (V10 – 10,000)	200*	Max remote active DDF Threads - DBM1 Address Space
CONDBAT	0-15,000 (V10 – 150,000)	10,000*	Max remote connections - DDF Address Space
MAXTYPE1	0-CONDBAT	0	Max inactive DBATs, these are used for private protocol. DRDA uses inactive connections.
POOLINAC	0-9,999	120	Approximate time, in seconds that an inactive/unused DBAT can remain idle in the pool before it is terminated. DBAT deleted after being used 200 times also.
IDTHTOIN	0-9,999	120*	The IDLE THREAD TIMEOUT (IDTHTOIN) parameter specifies the time (in seconds) that an active DBAT can remain idle before it is canceled. – Should set a “little” higher than TCPKPALV
TCPKPALV	ENABLE / DISABLE / 1-65534	120*	TCP/IP keep alive (Goes hand-in-hand with IDTHTOIN)
TCPALVER	NO (server), YES, CLIENT, SERVER, SERVER_ENCRYPT	NO	whether DB2 is to accept TCP/IP connection requests that contain only a user ID (no password, RACF PassTicket, or Kerberos ticket).

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Connection flowchart on DB2 for z/OS



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Connection Pool Properties (*WebSphere*)



- Statement Cache Size (**Default: 10**)
  - Number of statements to keep in prepared Statement Cache (different from DB2's stmt cache)
- Connection Timeout (**Default: 180**)
  - How long to attempt connection creation before timeout
- Max Connections (**Default: 10**)
  - max connections from JVM instance
- Min Connections (**Default: 1**)
  - minimum number of connections in pool
- Reap Time (**Default: 180**)
  - How often cleanup of pool is scheduled in seconds
- Unused Timeout (**Default: 1800**)
  - How long to let a connection sit in the pool unused
- Aged Timeout (**Default: 0**)
  - How long to let a connection live before recycling – controls inactive connections seen on host
- Purge Policy (**Default: EntirePool**)
  - After StaleConnectionException, does the entire pool get purged or only individual connection

PercentUsed	connectionPoolModule.percentUsed	9	The average percent of the pool that is in use	Per connection pool
PercentMaxed	connectionPoolModule.percentMaxed	10	The average percent of the time that all connections are in use	Per connection pool



Connection Pool statistics in WebSphere should be used to determine the applications need for concurrent connections



# Basic strategy for Connection pools

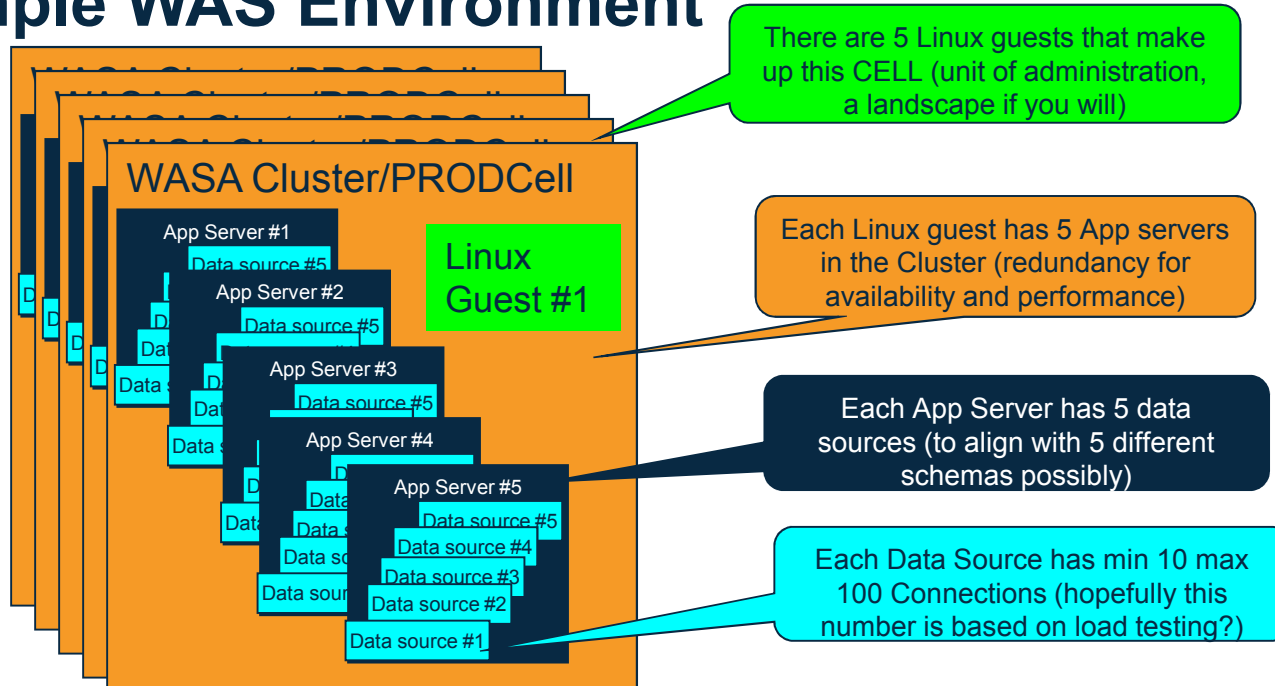


- Ideally an application would be tested for throughput and an appropriate Max Connections set by monitoring pool stats... but generally the defaults are taken
- Multiply the Max Connection values by the number of application servers going against that particular DB2 member or group. This gives you the real number of total connections that could be made to the database at anytime.
  - This number should relate to the sum of maxTransports in the drivers, since the transport is the physical connection
- If the total max connections going to the database (DB2 on z/OS) is larger than CONDBAT, then CONDBAT should be reassessed to absorb the overage, and avoid the Connection Timeout popping –
  - Should WAS be allowed to hold all the available connections?
  - CONDBAT +1 means –SQLCODE and DB2 appears unavailable
- What about DB2? – connections are not a big resource consumer, we are concerned with threads
  - MEMU2 or IFCID 225 information to determine max number of threads based on 31-bit virtual constraint
  - Do you have enough REAL storage to back more threads? \*\*

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Example WAS Environment



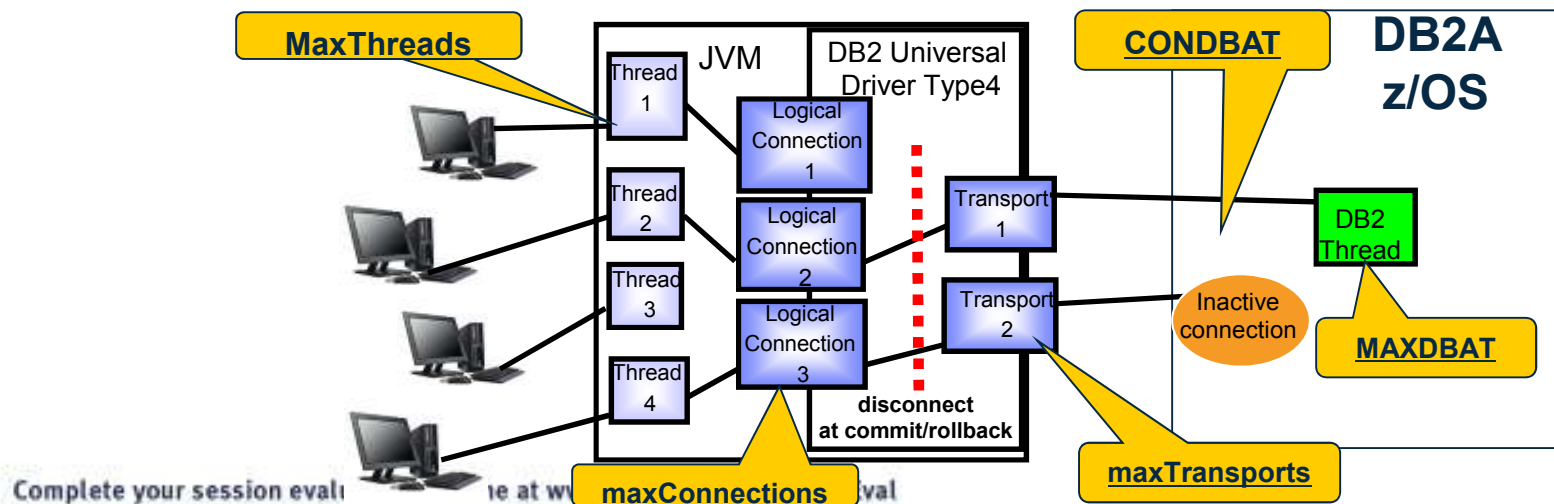
- For this example: Min connections of 10 and Max connections of 100 is defined at the CELL level which is the overarching administrative unit of control, meaning EVERY data source will have that min and max, they are not divvied up amongst all the servers
  - 5 data sources x 10 connections x 5 app servers x 5 clusters = min of 1,250 connections
  - 5 data sources x 100 connections x 5 app servers x 5 clusters = max of 12,500 connections
- This may be just one application landscape consuming 12,500 connections in the data sharing group!!





# So when you say connections.....??

- Scope of defined resources = CELL → NODE → CLUSTER → SERVER
  - More granular setting overrides more general one – need to account for resources SERVERs consume
    - So if maxConnections at cell is 50 with 5 servers, then total connections is really 250
- General logic: I have more app server threads than I need connections to DB2
  - I have more app server ‘connections’ than transports (real connection to DB2)
    - I have more connections to DB2 than active threads in DB2
      - Thus we funnel the most abundant resources down the most finite::
      - How many concurrent threads can I support in DB2??





## Recommendations WITHOUT Sysplex WLB

- WAS settings relative to DB2 timeouts, use <,> suggestions to adjust defaults
  - TCPKPALV (Enable=720) < IDTHTOIN (120) - we want to ping the socket prior to having DB2 cancel a thread for being idle incase we have lost the TCP/IP connection
  - IDTHTOIN (120) > Aged Timeout (0) - This may be the most controversial one, having aged timeout lower will recycle the connections after so many seconds (minutes) in order to: 1) reduce the possible number of inactive connections seen in the DB2 member 2) Allow the connection to be re-driven to another member when Sysplex WLB is being used
    - The downside to this would be the small overhead in re-establishing a new connection, and thread reuse when the connection is closed
  - Unused Timeout (1800) > Reap Time (180) - means we kill unused connections at the server to free up unused resources, and that no setting should be less than the Reap time, since it is the daemon that goes out to determine if other timeouts have popped
    - But if you have a minimum connection count then those will not be affected by Unused timeout
  - Connection Timeout (Network Timeout) – how long can the application handle not getting a connection, this should be a very small number of seconds with sysplex distributor connecting to a data sharing group
  - **\*\*IDTHTOIN > TCPKPALV > Aged Timeout > Unused Timeout > Reap > ConnectionTimeout - this the end-to-end picture with all 4 of the major timeout values**

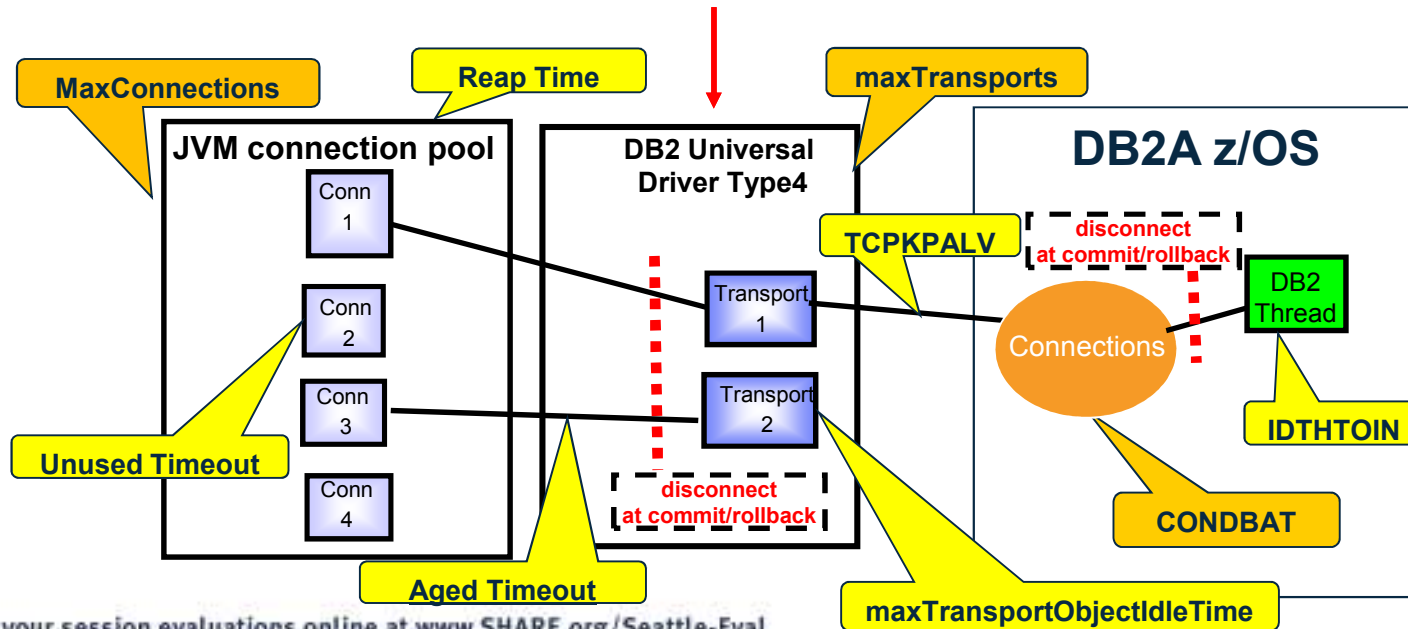
\*\*These settings are to eliminate stale, inactive, and hung connections in DB2

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Timeout Values

- Timeout values can be used to avoid hung threads and stale connections to DB2
- If inactive connections continue to grow and reach CONDBAT they could cause a service outage to other applications
  - These timeouts can be used to keep them to a minimum
  - **Without** Sysplex WLB the Conn's go directly to DB2, no Transports to manage them or free up connections



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

# Recommendations WITH Sysplex WLB



- Zero out application server parameters that correspond to Transport object settings
  - Aged timeout = 0 back to default / Unused = 0
  - Reap Time = 0 (also disables aged and unused timeouts)
  - Purge Policy = FailingConnectionOnly
    - Applications should know how to handle stale connection exceptions :
      - <http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg21063645>
- Use the Driver settings to protect DB2 if WAS pool settings are not optimal
  - maxTransportObjects (1,000 default)
    - Max number of connections at driver level (not data source level)
    - TOTAL # transports < CONDBAT \* DB2 Members
      - But MAXDBAT limits concurrency of execution
  - maxTransportObjectIdleTime = 60 seconds (default)
    - This is how you get rid of inactive connections (like Aged timeout in WAS)
    - As long as this is less than IDLTHTOIN then we are okay
  - maxTransportObjectWaitTime, default is 1 second
    - SQLCODE -4210, SQLSTATE 57033 for connection waiting for a transport
  - maxRefreshInterval = 10 seconds (for WLM)

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)





# Sysplex Workload Balancing routing

- Think of it as 'contingency management' not balancing \*\*
- Where would the next transaction be assigned when a 2 member group has Member A weight =30 / Member B =10 (3:1 ratio)
  - $(\text{Member's currently active transactions}) / (\text{Total active transactions for the group}) \leq (\text{Member priority}) / (\text{Total of all member priorities})$
  - Member A has 25 active connections (in a UoW)
  - Member B has 10 active connections (in a UoW)
  - Current ratios –
    - Member A = active connections / total connections = 25 / 35 = 0.71 (which is < 0.75 target ratio)
    - Member B = 10 / 35 = 0.28 (which is > 0.25 target ratio)
  - Next transaction would be assigned to Member A
- If the app servers were recycled the first 30 requests would go to A, then next 10 to B

DSNL102I  
displays the  
relative weights of  
the servers  
returned by WLM

```
-DIS DDF DETAIL  
DSNL101I WT IPADDR  
DSNL102I 30 ::10.138.48.221  
DSNL102I 10 ::10.138.48.218
```



# Thread Monitoring:



-DIS DDF DETAIL

How many threads are currently doing work as well as DSCDBATs

How many open sockets (connections) possible

How many threads you think you can handle.

```

=I CONDBAT= 1000 MDBAT= 200
ONQN= 0 MCONQW= 0
DSNL092I ADBAT= 31 QUEDBAT= 0 INADBAT= 0 CONQUED= 0
DSNL093I DSCDBAT= 9 INACONN= 168
DSNL094I WLMHEALTH=100 CLSDCONQN= 0 CLSDCO
DSNL106I PKGREL= COMMIT
    
```

How many threads are lounging in the pool

You hit MAXDBAT at some point and requests were queued

Type 2 inactive thread/ inactive connection

RELEASE(COMMIT)

-DIS THREAD(\*) DETAIL → V482-WLM-INFO=DB2DDFHI:1:2:50 (need OA45996 with z/OS 2.1)

- Shows WLM service class of enclave
- Shows what period it is executing in
- Shows the WLM Importance
- Shows performance index \* 100 (50 = P.I. of 0.5) of the service class period





# Protecting DB2 from Runaway Connections

- DBAs want to control the connection queue depth and the wait time for a queued connection to get a DBAT, exceeding which the connection should failover to another member, AND have health decrease as connections approach CONDBAT
  - Two new zPARMs: ON means MAXCONQN reacts when queue = MAXDBAT
    - MAXCONQN (depth) and MAXCONQW (wait time) are supported. -DISPLAY DDF
  - DSNL074I issued when > 80% CONDBAT reached (DB2 health cut to 50%)
    - DSNL075I issued when condition is relieved
  - Otherwise it takes WLM 10 seconds to lower DB2 member health
- -DIS DDF DETAIL will show the current configured values and connections closed because limit exceeded. Message DSNL030I also issued with reason code and stats updated
- APAR number is PM43293 (V9 & V10).
  - <http://www-01.ibm.com/support/docview.wss?crawler=1&uid=swg1PM43293>

DSNL094I WLMHEALTH=100 CLSDCONQN= 0 CLSDCONQW= 0

CUR QU INACT CONNS (TYPE 2)	0.00
MIN QUEUE TIME	0.000008
MAX QUEUE TIME	0.000384
AVG QUEUE TIME	0.000038
HWM QU INACT CONNS (TYPE 2)	13.00

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Dynamic Location Aliases



- Scenario : *Block incoming traffic from certain app servers (connection storm). Setting MAXDBAT=0 blocks all inbound traffic.*
  - Solution: You can associate a location alias with an app server and use the –MODIFY DDF command to dynamically stop the alias on a member which blocks incoming traffic from only those app servers that are connected to that alias without affecting other inbound or any outbound traffic.
- Scenario : *Need the ability to manage poorly performing applications by removing them from the Sysplex.*
  - Solution: You can associate a location alias with an app server and use the –MODIFY DDF command to dynamically exclude poorly behaved applications from sysplexWLB (change PORT)
- Scenario : *Stopping DDF causes all the app servers connected to a member to failover, and stops all outbound DB2 traffic. DBAs are looking for a more granular failover where only certain app servers fail over while others don't.*
  - Solution: You can associate a location alias with an app server and use the –MODIFY DDF command to dynamically stop the alias which only causes the app servers connected to that alias to failover, without affecting others.







## Dynamic Location Aliases - Continued

### ■ *Examples :*

-MODIFY DDF ALIAS(alias1) ADD

- Alias1 is created and is stopped by default.

-MODIFY DDF ALIAS(alias1) PORT(9000)

- Alias1 is associated with port 9000.

-MODIFY DDF ALIAS(alias1) IPv4(2.2.2.2)

- Alias1 is associated with IP address 2.2.2.2

-MODIFY DDF ALIAS(alias1) START

- DDF will accept requests for alias1 on port 9000.  
When a client connects to alias1, IP address 2.2.2.2 is returned in the server list.

-MODIFY DDF ALIAS(alias1) STOP

- Alias1 is stopped and will not accept new requests.  
Existing requests will be allowed to complete.

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Displaying location alias-specific information



## -DISPLAY DDF ALIAS(*alias1*) DETAIL

DSNL080I @ DSNLTDDF DISPLAY DDF (alias1) REPORT FOLLOWS:

DSNL087I ALIAS PORT SECPORT STATUS

DSNL088I ALIAS1 9000 5005 **STARTED**

DSNL089I MEMBER IPADDR=::2.2.2.2

DSNL089I MEMBER IPADDR=2002:91E:610::1

DSNL096I **ADBAT** = 100 **CONQUED** = 1000 **TCONS** = 1000

DSNL100I LOCATION SERVER LIST:

DSNL101I WT IPADDR IPADDR

DSNL102I 32 ::2.2.2.2 2002:91E:610::1

DSNL102I 32 ::1.2.3.4

DSNL099I DSNLTDDF DISPLAY DDF REPORT COMPLETE

**ADBAT** : The number of active data base access threads (DBATs) that are currently processing requests on behalf of the specified alias.

**CONQUED** : The number of connection requests that are currently queued and waiting to be serviced on behalf of the specified alias.

**TCONS** : The total number of remote connections that are currently associated with the alias.



# Profile enhancements in DB2 10



- Catalog table holds profiles
- Can limit:
  - Active threads
  - Idle thread timeout
  - Connections
- Limit based on
  - LOCATION only
  - PRDID only
  - AUTHID, ROLE, or both.
  - COLLID, PKGNAME, or both
  - One of CLIENT\_APPLNAME, CLIENT\_USERID, CLIENT\_WORKSTNNAME
- Warning or Exception
  - 1 attribute could be warning (just once or for each threshold), another could be exception where incoming connections are queued or failed
  - DSNT771I reason code 00E3050x

PRO-FILEID	KEYWORDS	AT-TRIB-UTE1	AT-TRIB-UTE2	AT-TRIB-UTE3	AT-TRIB-UTE TIMES TAMP	RE-MARKS
1	MONITOR THREADS	EXCEPTION_D IAG-LEVEL2	10		2008-12-19...	
2	MONITOR CON-NEC-TIONS	WARN-ING	50		2008-12-17...	
3	MONITOR IDLE THREADS	EXCEPTION_D IAG-LEVEL1	300		2009-01-25...	

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)



# Limiting Total Connections

- Could use Profiles to protect DB2 from connections spawned by application servers in a loop, poorly behaving application, or a denial of service attack
  - Profile #1: Wildcard location so **NO** remote IP address or domain can only use 1,000 connections before they are rejected
  - Profile #2: that IP address can only use up to 500 connections



SYSIBM.DSN\_PROFILE\_TABLE

PROFILEID	LOCATION
1	*
2	9.76.193.254

SYSIBM.DSN\_PROFILE\_ATTRIBUTES

PROFILEID	KEYWORDS	ATTRIBUTE1	ATTRIBUTE2
1	MONITOR CONNECTIONS	WARNING	1000
2	MONITOR CONNECTIONS	EXCEPTION	500

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)





# References

- Techdoc for V10 and V11 MEMU2 with spreadsheet sample
  - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS5279>
- DB2 for z/OS and WebSphere Integration for Enterprise Java Applications
  - <http://www.redbooks.ibm.com/abstracts/sg248074.html?Open>
- Subsystem and Transaction Monitoring and Tuning with DB2 11 for z/OS
  - <http://www.redbooks.ibm.com/redpieces/abstracts/sg248182.html?Open>
- DB2 and WebSphere Synergy Workshop to help manage connections
  - Contact me at [agburke@us.ibm.com](mailto:agburke@us.ibm.com)

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)





# Thank You (??'s)



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

