The next step in big data starts with IBM.
Extreme speed. Increased flexibility. Enhanced performance.

IBM DB2 for z/OS
Distributed Strategy in the Mobile Age
DB2Night Show

DB2 11 for z/OS
The Enterprise Data Server for Business Critical Transactions and Analytics.

Adrian Burke DB2 for z/OS SWAT team
agburke@us.ibm.com

© 2014 IBM Corporation

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, [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
- 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.

Agenda

- High Performance DBATs
- High Availability
- Properties
- Tuning
  - Timeouts/connections
- Protecting DB2
  - Profiles
  - Mobile pricing

The next step in big data starts with IBM.

Distributed Environment Today

- Credit card company – 6-way data sharing group
  - 220M transactions through the DDF address space in a day
    - 2.2B DML statements on 1 member in 24 hours
  - Roughly 36 hours of CPU/zIIP in DDF address space on 1 member in 1 day
- Bank in South America
  - 123M ATM transactions routed to data sharing group in 1 day
- U.S. Bank – 1 member of a data sharing group
  - Running peak online day with 5,000 concurrent DDF threads
  - Roughly 15,000 concurrent connections bringing in work
- U.S. Insurer - 4 member data sharing group
  - Possibility of 220K concurrent to from remote application servers
- Large Asian bank – 12-way data sharing
  - Over 300M customers with cell phones able to check their balance concurrently

The next step in big data starts with IBM.

High Perf DBATs

- **DB2 10 high performance database access threads**
  - Provides additional CPU benefits in DB2 10 (10-20% CPU) reintroduces RELEASE(DEALLOCATE) to DDF
  - Requires rebinding of DB2 Connect client packages
    - Should use the set **currentPackageSet** in driver for 'good' applications to point to another collection ID – not NULLID
    - DB2Binder utility defaults to **releasePackageResourcesAtCommit(false)** for DB2 10 packages (9.7 FP 3a) – so you get DEALLOCATE by default
    - MAXDBAT should be increased by number of estimated HPDBATs because they will essentially be dedicated to those connections
  - POOLINAC will determine how long a high performance DBAT remains idle before being destroyed
    - Threads are recycled after 200 uses
  - WILL drive up number of concurrent DBATs → MAXDBAT because they are no longer available to other transactions at commit

The next step in big data starts with IBM.

Remote connection monitoring for High Perf. DBATs

- 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 accnt 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

Accounting Short Report

High Perf DBATs

- **DB2 10 high performance database access threads**
  - Can be altered on and off with -MODIFY DDF PKGREL command
    - COMMIT – no HPDBATs
    - BINDOPT – HPDBATs for DEALLOCATE packages, destroyed at disconnect time
    - BNDPOOL - added by PI31957 allows HPDBATs to be pooled when connection goes away (deallocation)
  - If HPDBATs are not used 200 times before application disconnects can result in DIST SRB (non-preemptable) overhead from thread creation/destruction
    - BNDPOOL option allows HPDBATs to be reused by ANYONE so you get a portion of the HPDBAT performance savings from DEALLOCATE packages, and once pooled any thread can pick up a HPDBAT
      - Even more reason to segregate by collection ID
      - Number of concurrent DBATs will still be higher than with COMMIT
    - Downside is that package and object intent locks remain on HPDBAT until deallocation even if a different application picks it up

The next step in big data starts with IBM.

Getting to High Availability

The next step in big data starts with IBM.

DB2 z/OS Inactive Thread Support (CMTSTAT=INACTIVE)

- Each inbound SQL request from distributed environments to the DB2 for z/OS server requires a DDF connection and a DB2 database access thread (DBAT). DB2 for z/OS inactive connection support (formerly referred to as Type 2 inactive thread support or thread pooling) is a mechanism to share few DBATs among many connected applications.

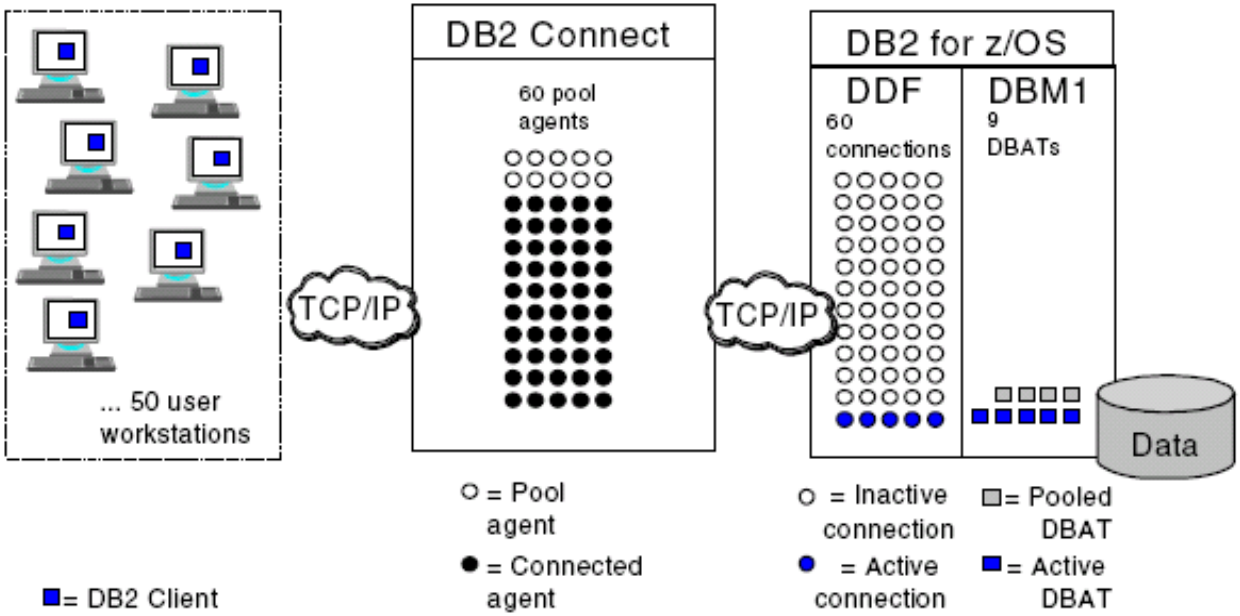

- The benefits for DB2 z/OS Thread Pooling are:
  - CPU savings in DB2, by avoiding repeated creation and destruction of DBAT
  - Real memory savings in z/OS, by reducing the number of DBATs
  - Virtual storage savings in DBM1, by reducing the number of DBATs
  - Greater capacity to support DRDA connections

Key points to Pooling and Concentration

- Both can support Data Sharing
- Both can support DVIPA and Sysplex Distributor
- Connection Pooling shares connections at connection time
- Connection Concentration shares connections at commit time
  - 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.
    - Threads using RELEASE(DEALLOCATE) packages
    - 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

The next step in big data starts with IBM.

High Availability…reasons why NOT to use a 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
  - 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
      - RTPIFACTOR in IEAOPT can reduce the impact of PI>1, reduce connection sloshing
    - MAXCONQN / MAXCONQW

The next step in big data starts with IBM.

Sysplex Workload Balancing…ACR

- 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

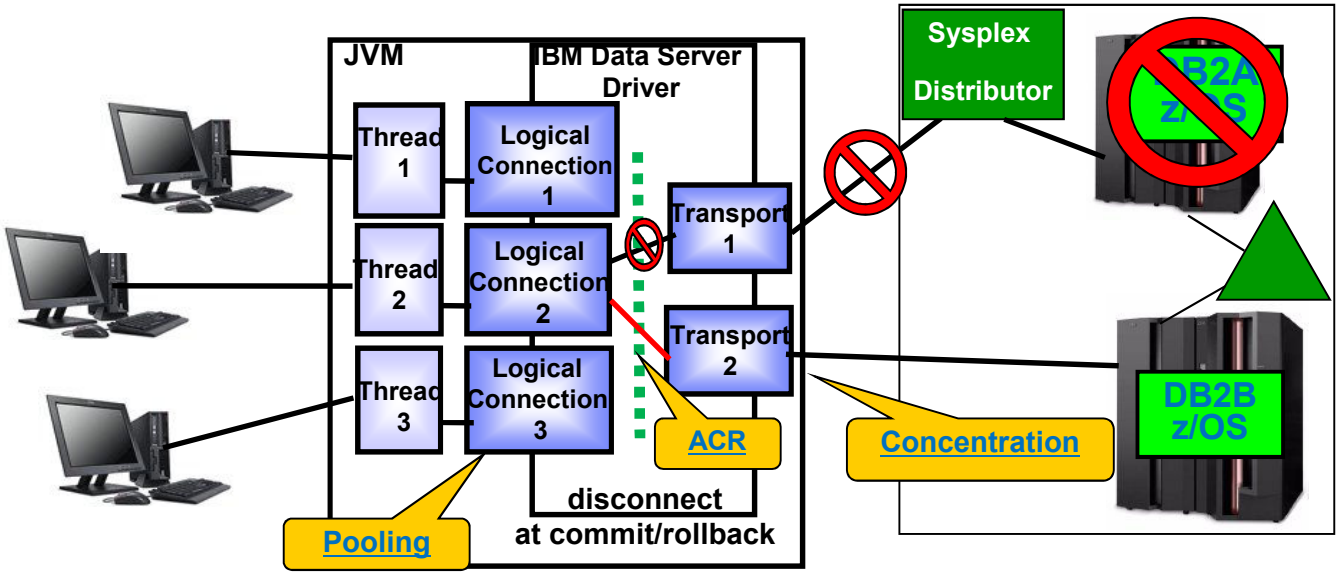

Sysplex WLB properties JAVA driver

- **Java datasource properties - V9.7 FP6/ V10.1 FP2/ V10.5 FP4**
  - **EnableSysplexWLB= TRUE** enables sysplex workload balancing
  - **MaxTransportObjects** specifies maximum number of connections managed across the data sharing group. The default is 1,000 (previously unlimited)
  - **maxRefreshInterval maxmum** 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 10
  - **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

The next step in big data starts with IBM.

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)
  - (Member's currently active transactions) / (Total active transactions for the group) <= (Member priority) / (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

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
```


The next step in big data starts with IBM.

Tuning Data Connections

The next step in big data starts with IBM.

Connection flowchart on DB2 for z/OS

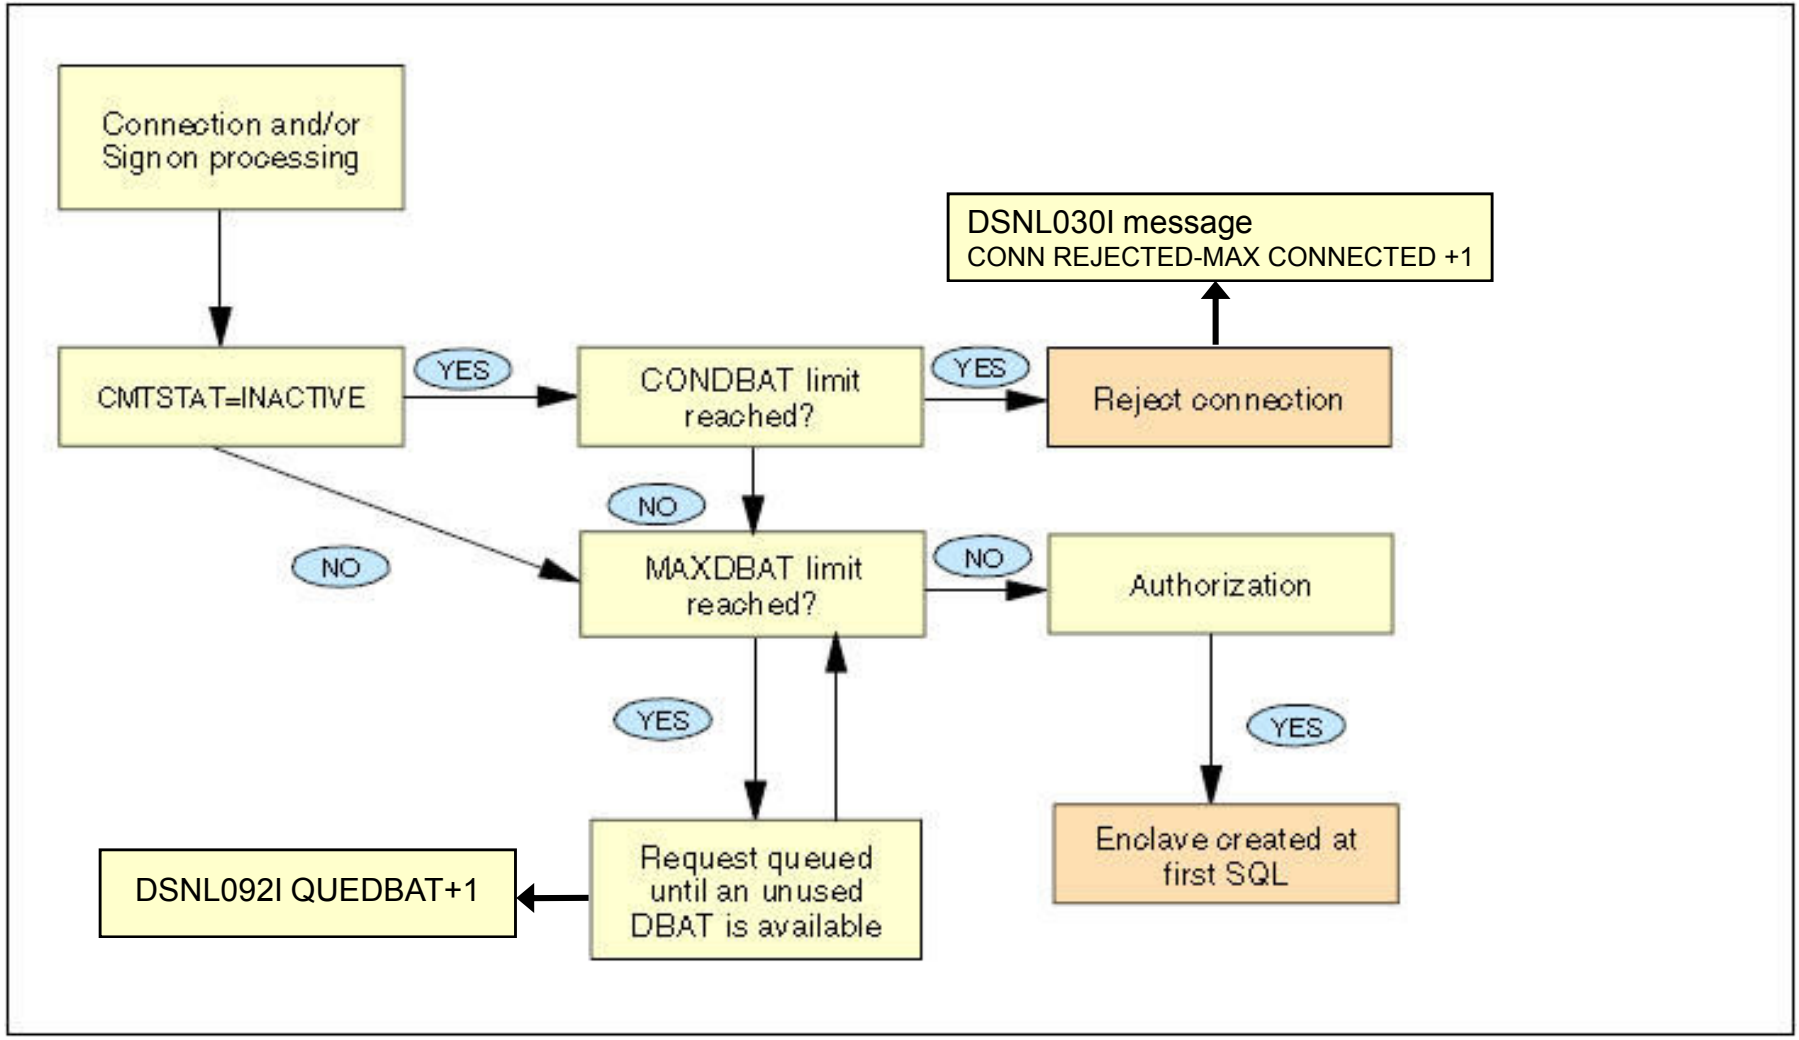

The next step in big data starts with IBM.

Remote Connections

- With relief of virtual storage constraint MAXDBAT has been raised but many customers ignored the new default of CONDBAT in V8 (10,000)
  - Much better to queue DBAT requests outside DB2 between MAXDBAT and CONDBAT, than to have the connection request be rejected
- With more remote apps/app servers no control over connections without a Gateway
- CONDBAT needs to reflect total MAX CONNECTIONS from the application servers
- If inactive connections continue to grow and reach CONDBAT they could cause a service outage to other applications – NO timeout value in DB2 to cancel inactive connections
  - Connection timeouts exist in the application servers (AGED timeout) or Driver with enableSysplexWLB=YES (maxTransportObjectIdleTime)
  - Conn 3 is inactive but holding IP socket hostage

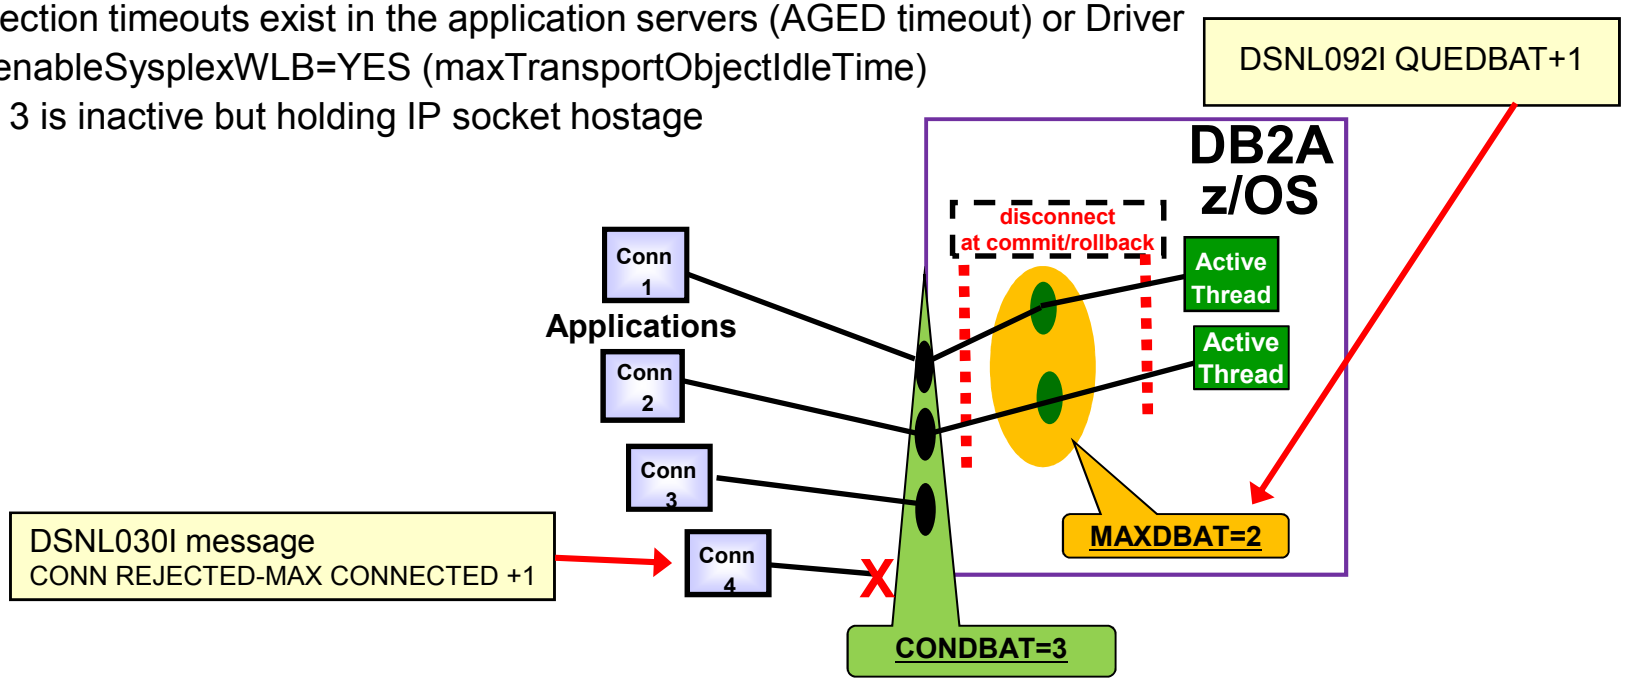

The next step in big data starts with IBM.

Connection Pool Properties WebSphere

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

- Statement Cache Size (Default: 10)
  - Number of statements to keep in prepared Statement Cache
- Connection Timeout (Default: 180)
  - How long to attempt connection creation before timeout
- **Max Connections (Default: 10)**
  - **max connections from JVM instance – determining factor CONDBAT**
- Min Connections (Default: 0)
  - lazy 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

The next step in big data starts with IBM.

Tomcat / WAS/ JCC driver

Tomcat	WebSphere Equivalent	Driver
Remove abandoned Timeout (seconds)	Aged Timeout	maxTransportObjectIdleTime
maxActive – Tomcat also has settings for maxIdle connections	maxConnections – WAS does not distinguish between active and idle in number	maxTransports
initialSize – Tomcat creates all these at start up	Min Connections- WAS will wait until they are needed then keep them around	minTransportObjects
****validationQuery** – string for query to submit (leave it out) – pops on query timeout in driver	Connection validation timeout - SQL validation deprecated, use JDBC validation (JDBC 4.0)	Sysplex WLB should take care of this, re-drive failed connections through syspelx distributor, and automatic client reroute
testOnBorrow- check it when connection used from pool testOnReturn –check in going back to pool	Validate new connections Validate existing connections	** *Turn Validation off to avoid overhead..*

- Many open source or vendor application servers do not take advantage of availability enhancements in z/OS
  - Even our IBM Data Server Driver has different settings for z/OS vs. LUW (DB2)
  - As an example Tomcat still has the concept of a validation query
    - Enabling this function means the connection object in the pool is tested before the application acquires the connection, as well as when the connection is returned to the pool
  - The JDBC 4.0 specification made this feature irrelevant as the driver itself (IBM Data Server Driver) guarantees the connection given to the application server from the pool

The next step in big data starts with IBM.

.NET / WAS/ JCC driver

IBM .NET provider/ Default Connection Pooling = true	WebSphere Equivalent	JDBC Driver
**N/A	Aged Timeout	maxTransportObjectIdleTime (10)
Max Pool Size = No Maximum	maxConnections – WAS does not distinguish between active and idle in number	maxTransports (1,000)
ConnectionLifetime	Unused Timeout	N/A
N/A	REAP time	N/A
Min Pool Size = 0	Min Connections- WAS will wait until they are needed then keep them around	minTransportObjects
Connection timeout	Connection timeout	maxTransportObjectWaitTime (1 sec)

**** This means there is no way to timeout the connections coming from the driver**

**** Settings from PMR →**

	Default
Max Pool Size	No max
Min Pool Size	0
Connection Lifetime	60

The next step in big data starts with IBM.

Basic strategy for Connection pools

- Ideally an application would be tested for throughput and an appropriate Max Connections set by monitoring pool stats… but otherwise start by setting the Max Connections value to the sum of all of the max threads possible in the app server thread pools used by the application which might access DB2
- Multiply the Max Connection values by the number data sources and then 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 maxTransports in the driver if Sysplex WLB enabled, 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 deal, we are concerned with threads or MAXDBAT
  - 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? **

The next step in big data starts with IBM.

Example WAS Environment

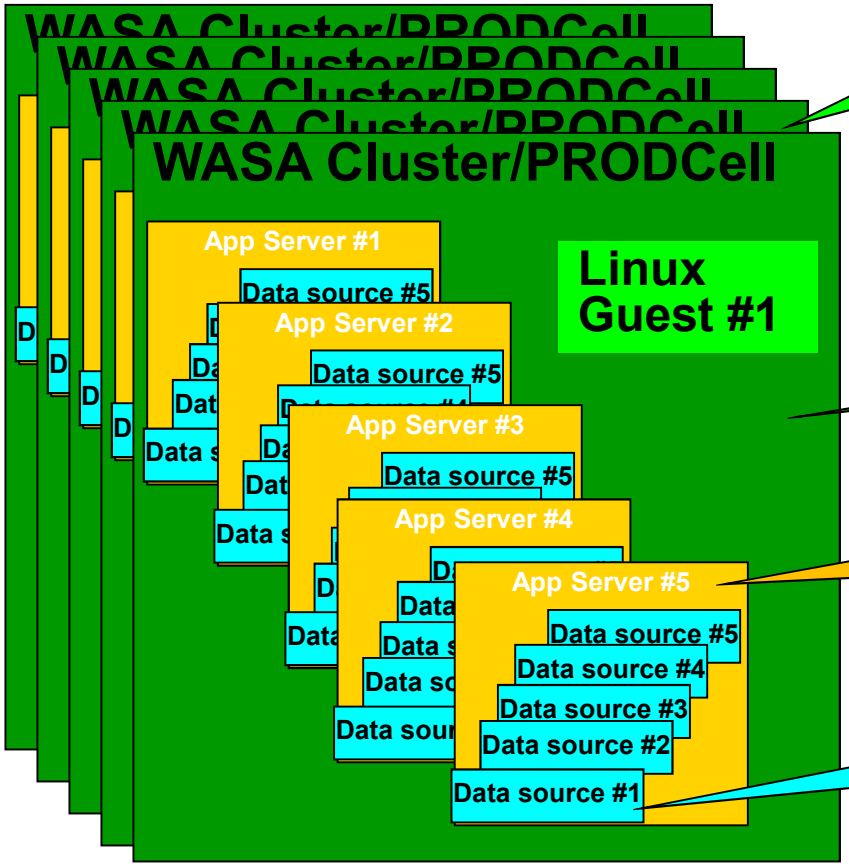

There are 5 Linux guests that make up this CELL (unit of administration, a landscape if you will)

Each Linux guest has 5 App servers in the Cluster (redundancy for availability and performance)

Each App Server has 5 data sources (to align with 5 different schemas possibly)

Each Data Source has min 10 max 100 Connections (hopefully this number is based on load testing?)

- 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!!

The next step in big data starts with IBM.

So when you say connections…..??

- Scope of defined resources = CELL → NODE → CLUSTER→ SERVER → DS
  - 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??*

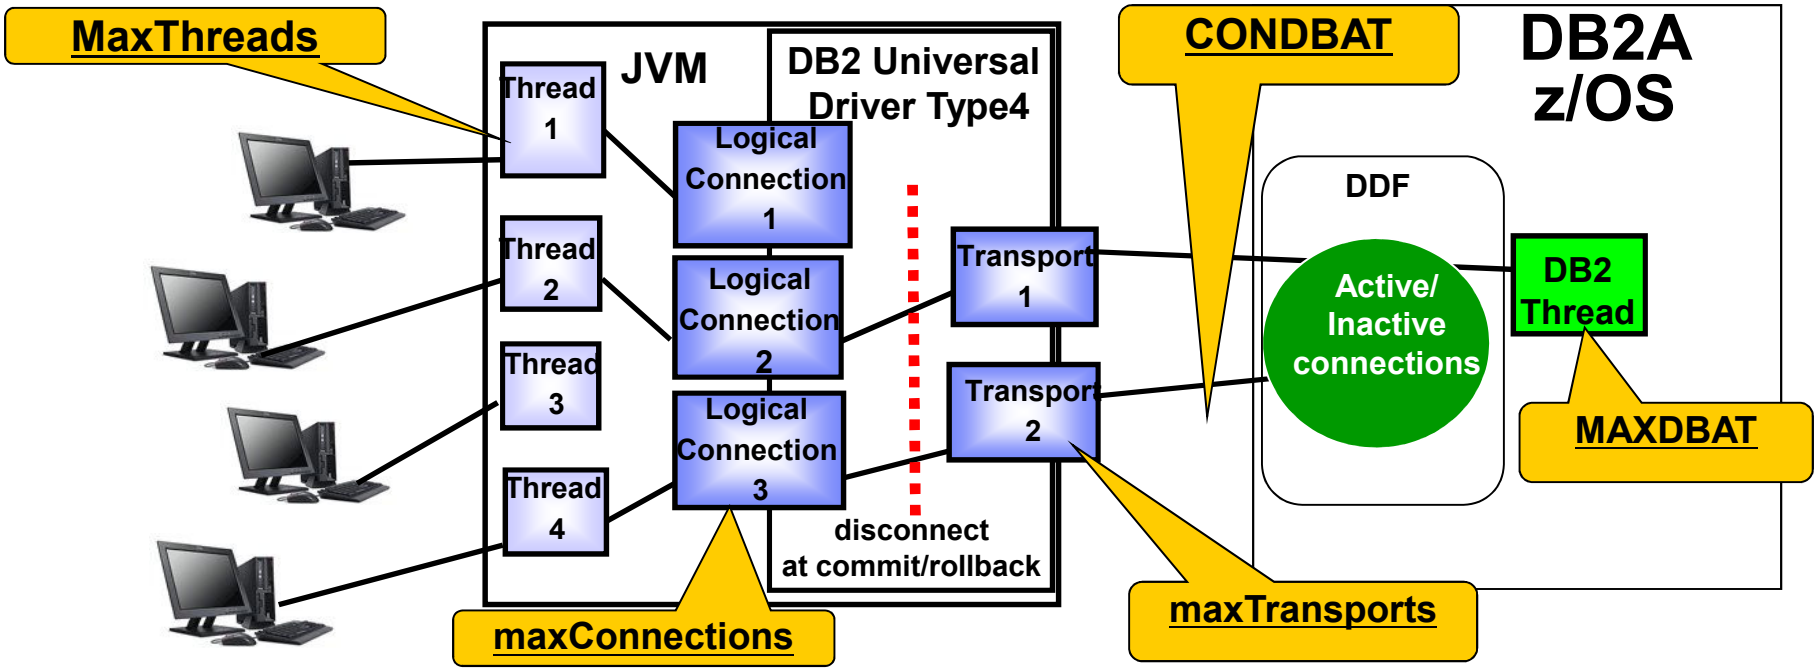

The next step in big data starts with IBM.

Timeout Values

- Timeout values should be used to avoid stranded resources
- 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 inactive connections to a minimum
    - AgedTimeout in the WebSpere data source
    - maxTransportObjectIdleTime in the JDBC properties if enableSysplexWLB=True

»Refer back to slide 19-20 for non-WAS timeout values…

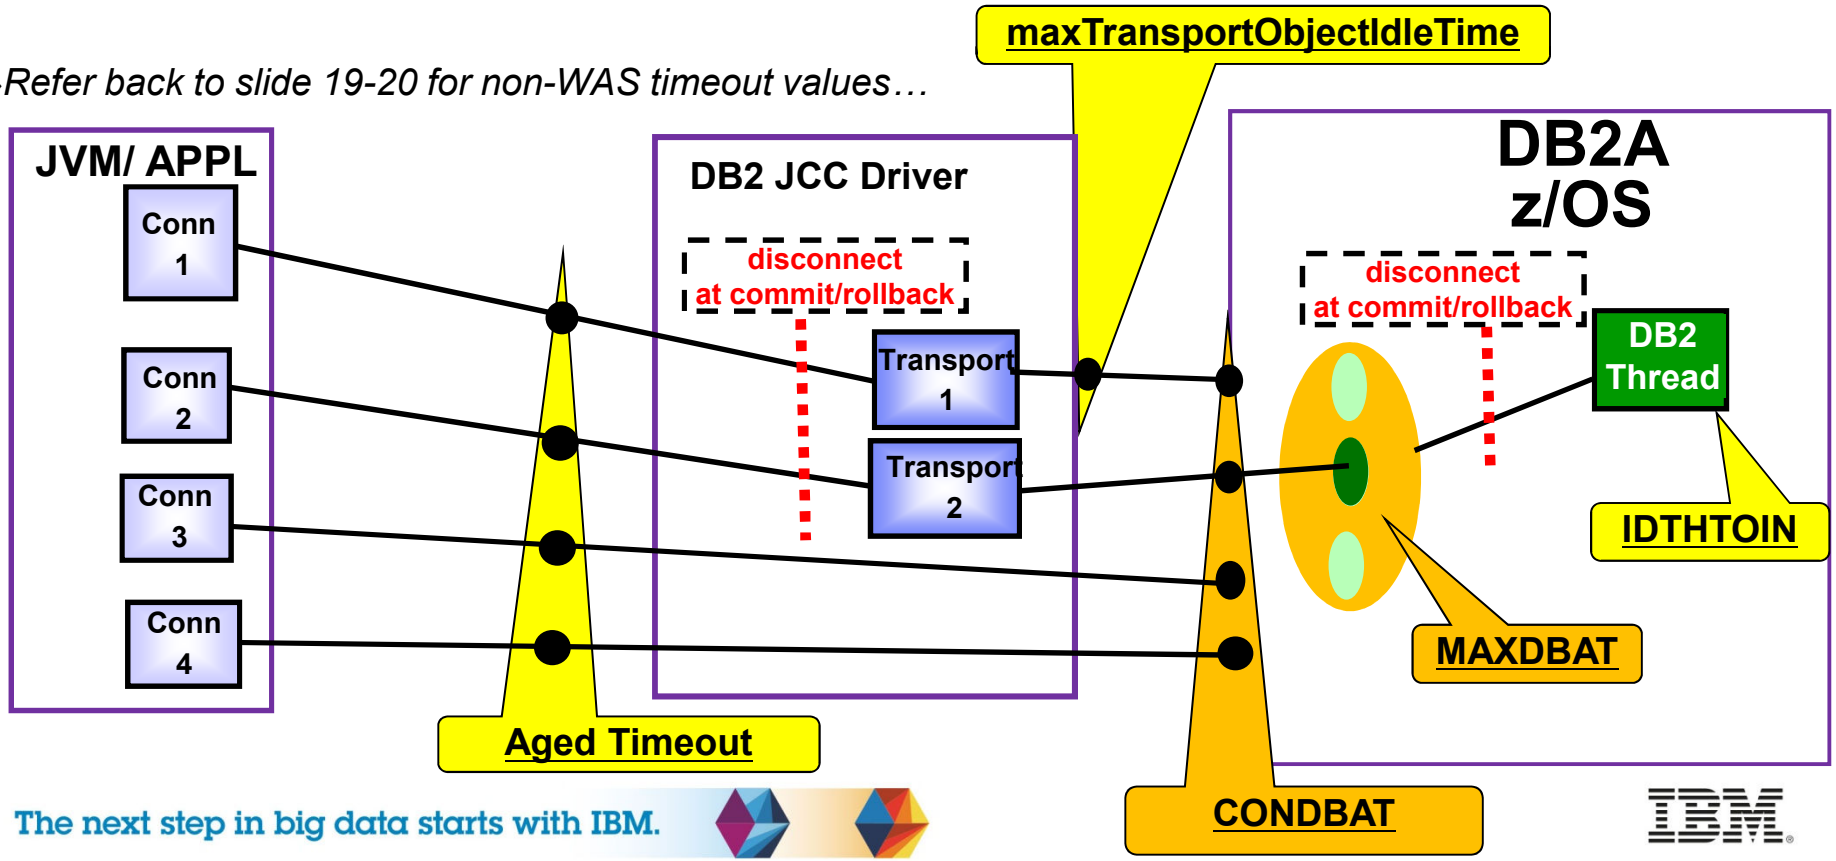

The next step in big data starts with IBM.

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**

The next step in big data starts with IBM.

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 further tune connections/timeouts if there are issues
  - maxTransportObjects (1,000 default)
    - Max number of connections at driver level (not data source level)
    - TOTAL # transports * #drivers < CONDBAT * DB2 Members of DS Group
      - But MAXDBAT limits concurrency of execution
  - maxTransportObjectIdleTime = 10 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)

The next step in big data starts with IBM.

Thread Monitoring:

-DIS DDF DETAIL

```
DSNL090I DT=I CONDBAT=   1000 MDBAT=  200
DSNL091I MCONQN=    0 MCONQW=   0
DSNL092I ADBAT=   31 QUEDBAT=     0 INADBAT=     0 CONQUED=     0
DSNL093I DSCDBAT=     9 INACONN=   168
DSNL094I WLMHEALTH=100 CLSDCONQN=     0 CLSDCON
DSNL106I PKGREL = COMMIT
```

How many open sockets (connections) possible

How many threads are currently doing work as well as DSCDBATs

How many threads you think you can handle.

You hit MAXDBAT at some point and requests were queued

How many threads are lounging in the pool

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**

The next step in big data starts with IBM.

Customer Conns Rejected

- A connection storm or incremental growth can occur when database access threads are not even consumed
- Need to monitor the number of inactive connections and ensure CONDBAT is never encroached upon
  - Sum of inactive connections and active DBATS
  - DSNL047I (CONDBAT at 80%) messages generally come too late to react

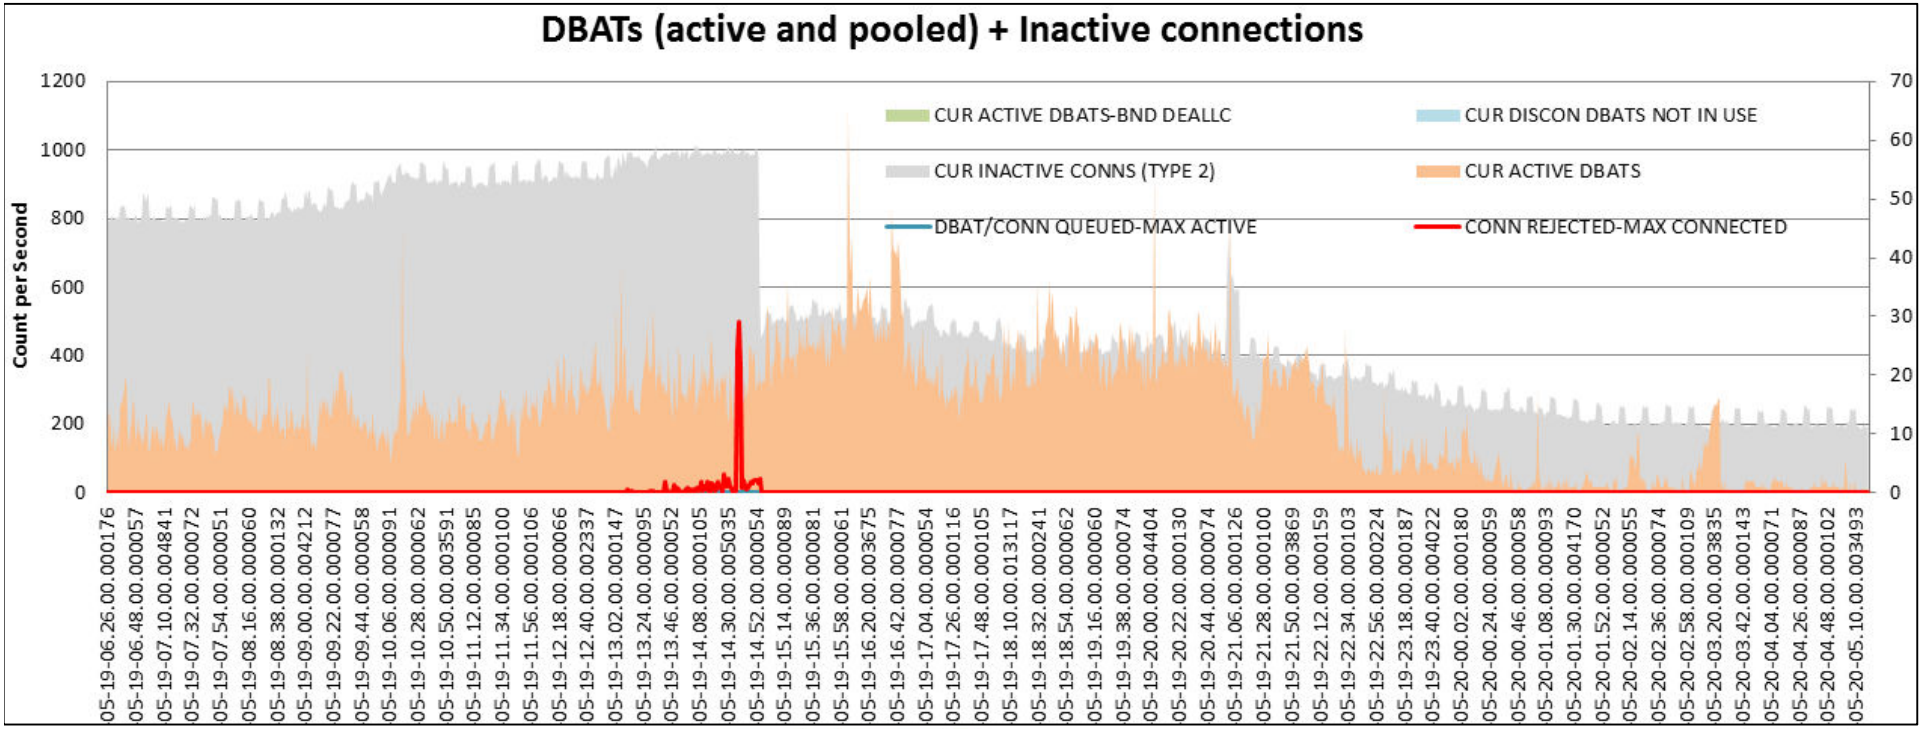

The next step in big data starts with IBM.

Who is connecting to DB2?

- *-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**

http://www-01.ibm.com/support/docview.wss?uid=swg21363866

Example 2-2 DISPLAY LOCATION report

```
DSNL200I  -DOZ2 DISPLAY LOCATION REPORT FOLLOWS-
LOCATION                                  PRDID    T ATT CONNS
::9.12.4.142                              JCC03640 S        50
                                                     WLB    50
                                                     XA     50
```


The next step in big data starts with IBM.

Cancelling Inactive Connections

- If you have 'stale' connections you can get rid of them *one-at-a-time*
- Display the inactive connections

  -DISPLAY THREAD(*) TYPE(INACTIVE) DETAIL

  V437-WORKSTATION=9.76.193.254 ← IP address

  V448--(1) 446:53637 ← local and foreign ports

- Find the connection identifier (CONN) associated with the IP addresses for that foreign local port combination:

  D TCPIP,,NETSTAT,conn,ipaddr=9.76.193.254

  USER ID CONN STATE

  DB1SDIST 162724E6 ESTBLSH ← get the CONN identifier

- Drop the inactive connection which corresponds to that identifier:

  V TCPIP,,DROP,CONN=162724E6

The next step in big data starts with IBM.

Monitor profiles to limit threads/connections

- 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

PRO-FILEID	KEYWORDS	AT-TRIB-UTE1	AT-TRIB-UTE2	AT-TRIB-UTE3	AT-TRIB-UTE TIMESTAMP	RE-MARKS
1	MONIT-OR THREADS	EXCEP-TION_DIAG-LEVEL2	10		2008-12-19...	
2	MONIT-OR CON-NEC-TIONS	WARN-ING	50		2008-12-17...	
3	MONIT-OR IDLE THREADS	EXCEP-TION_DIAG-LEVEL1	300		2009-01-25...	

- 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
  - DSNT77xI reason code 00E3050x with various degrees of detail

The next step in big data starts with IBM.

Limiting Total Connections

- Could use Profiles to protect DB2 from connections spawned by application servers in a loop (garbage cleanup), poorly behaving application, or a denial of service attack
  - Profile #1: Wildcard location so any remote IP address or domain using more than 1,000 connections to any member writes a warning to the master log (DSNT772I)
  - Profile #2: That IP address can only use up to 500 connections to a member where the profile is started, and the 501st is rejected

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

http://www.ibm.com/support/knowledgecenter/SSEPEK_11.0.0/com.ibm.db2z11.doc.perf/src/tpc/db2z_monitorconnectionprofiles.dita?lang=en

The next step in big data starts with IBM.

WLM and MOBILE Pricing

- I can define a WLM service classification rule to subset transactions with client accounting information out from those using a specific Authorization ID on the server
  - You could also use client IP address to distinguish them; this IP address would be the same one you would see in the accounting string if you did a –DIS THREAD
- Using a sub-rule for the classification where the accounting string is 'WASMOB' I can push it down into a lower service class if needed or just report on it
- The CPU consumption will then be reported as well

```
                    Modify Rules for the Subsystem Type        Row 1 to 4 of 4
Command ===> ____________________________________________ Scroll ===> PAGE

Subsystem Type . : DDF          Fold qualifier names?   Y  (Y or N)
Description  . . . Distributed DDF work

Action codes:   A=After     C=Copy        M=Move      I=Insert rule
                B=Before    D=Delete row  R=Repeat    IS=Insert Sub-rule
                                                         <=== More
         -------Qualifier--------      Storage   Reporting  Manage Region
Action   Type      Name     Start      Critical  Attribute  Using Goals Of

 ____   1 UI        WAS*     ___        N/A       NONE       N/A
 ____   2   CAI       WASMOB  ___       N/A       MOBILE     N/A
```

TRANSACTION APPL% :	TOTAL :	CP	N/A	AAP/IIP ON CP	N/A	AAP/IIP	N/A
	MOBILE :	CP	N/A	AAP/IIP ON CP	N/A	AAP/IIP	N/A

The next step in big data starts with IBM.

References

- PI63531 – Improve granularity of monitor profile messages
  - http://www-01.ibm.com/support/docview.wss?crawler=1&uid=swg1PI63531
- Techdoc for V10 and V11 MEMU2 with spreadsheet sample
  - https://www-01.ibm.com/support/docview.wss?uid=tss1prs5279
- 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
- World of DB2
  - http://www.worldofdb2.com/events/event/listUpcoming

The next step in big data starts with IBM.

Thank You
(??'s)

The next step in big data starts with IBM.