

Exploiting Db2 DDF Without Fear or Revelations

Anthony Ciabattoni
Executive IT Specialist
Db2 for z/OS SWAT Team Leader


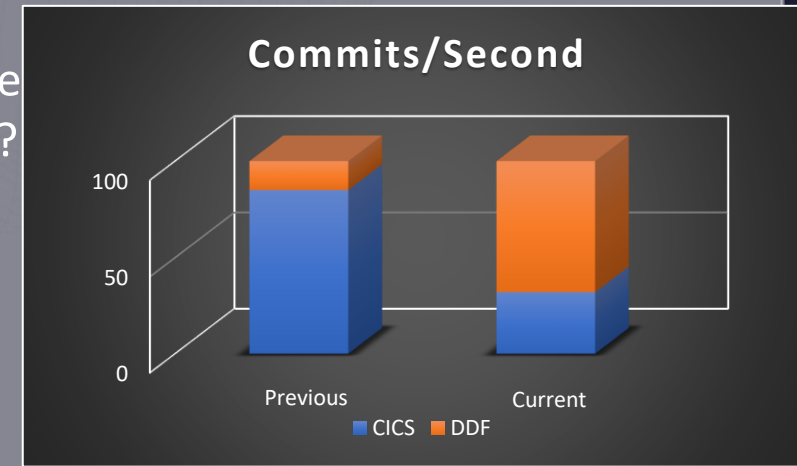


Topics

- Objectives
- Overview of Common Topologies and Terminologies
- Db2 Thread and Connection Pooling
- High Availability Workload Routing
- High Performance DBATs
- Transaction Identification and Monitoring
- Configure for Resiliency
- Db2 System Monitor Profiles
- Questions

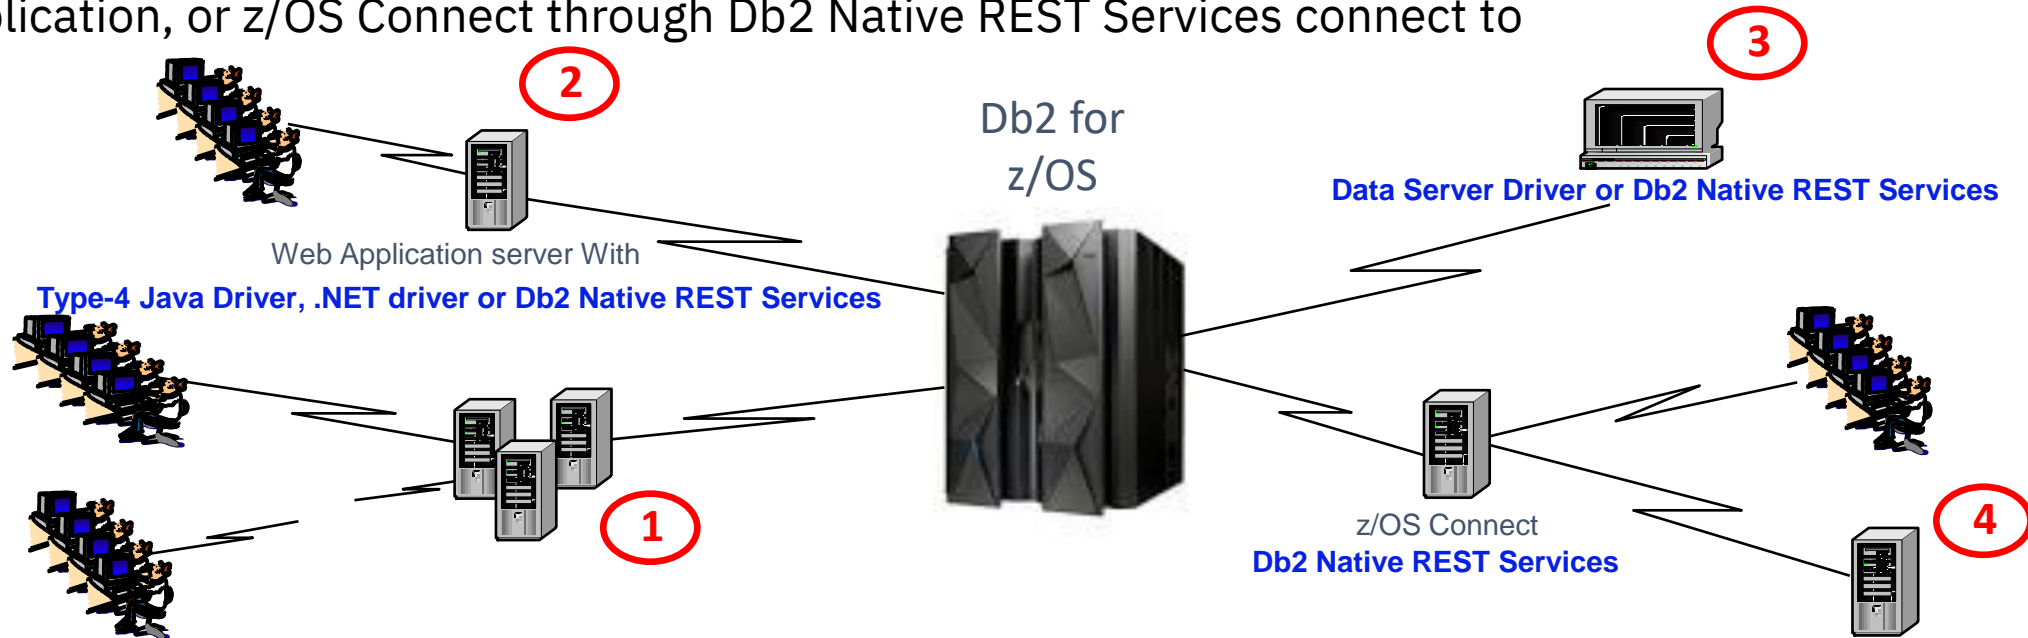
Objectives

- DDF online mainstream (OTLP) workload has significantly increased in adoption over the past several years
- Why is there an appetite from moving transactions from a traditional reliable CICS – Db2 infrastructure to a Db2 for z/OS DDF architectural infrastructure ?
 - Lower CPU cost per transaction (zIIP offload)
 - Simplify the architectural infrastructure 2-tier client server
 - Familiar coding techniques for developers accessing data residing on Db2 for z/OS
 - DDF maturity evolution
 - Focusing development on availability, redundancy and efficiencies
- Even with adoption, growth and availability successes and efficiency, customers have an apprehensiveness when discussing DRDA
 - What are the real and perceived concerns
 - Managing the unknown – Where are all the servers coming from ?
 - How do I identify and monitor DRDA transactions – “Black Box” ?
 - How do I control what is coming into my system – “Wild Wild West” ?
- The session objective is to address the common concerns, describe technology solutions to ultimately decrease the fear and anxiety of exploiting DRDA



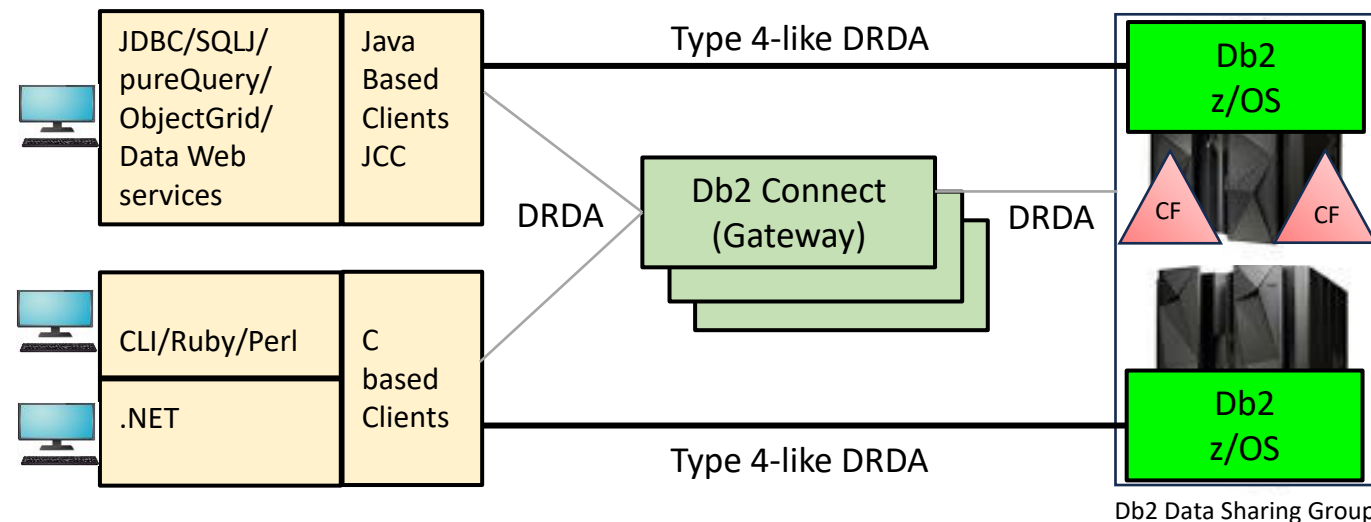
Remote access to Db2 – Common Topologies

- Common scenarios
 1. End client connects to an http/web server which then goes through a Db2 Connect Gateway, or cluster of Gateways to get to Db2 for z/OS
 2. End client/requestor connects to an application server where the IBM Data Server Driver resides which then connects to Db2 for z/OS
 - Application server could be WebSphere, .NET, Tomcat, etc.
 3. End user/developer connects directly to Db2 for z/OS from their workstation via the Data Server Driver or Db2 Data Server Driver
 4. End user, Application, or z/OS Connect through Db2 Native REST Services connect to Db2 for z/OS




Remote access to Db2 – Strategic Topology

- IBM's strategic direction is toward the IBM Data Server Driver embedded within the application servers for enterprise applications
 - Eliminates Gateway as a single point of failure and/or sprawling server farms of Gateways
 - Eliminates latency and another network 'hop'
 - Data Server driver has a much smaller footprint (~5MB as opposed to several 100MB)
 - Eliminates version-to-version migration of Gateways and coordination of release dependencies between client/Gateway/Db2 z
 - No need to catalog node/DCS entries or migrate them
 - Simplifies Db2 Connect licensing which can now be hosted in Db2 z itself (Enterprise license)




Overview and Common terms

- What is a connection?
 - DRDA over TCP/IP = IP address (domain) and port resulting in an end-point or socket and a control block in DIST address space
 - Total number governed by ZPARM CONDBAT (default = 10,000)
 - At 80% of CONDBAT, DSNL074I message in MSTR log
 - If CONDBAT is exceeded, DSNL030I message in MSTR log and connection requests are rejected and Db2 cannot be accessed remotely
 - Socket not released until remote requestor closes connection
 - ZPARM TCPKPALV determines interval to ping socket to ensure connection is still there – avoid hung thread if connection drops
 - Default is 120 seconds
 - ENABLE means takes COMM Server default of 120 minutes
 - Legacy myth that connections are expensive
 - 2-4K in memory footprint above 2GB bar in DIST address space
 - Still have 210 bytes in Comm server ECSA
 - Max for CONDBAT has been 150,000



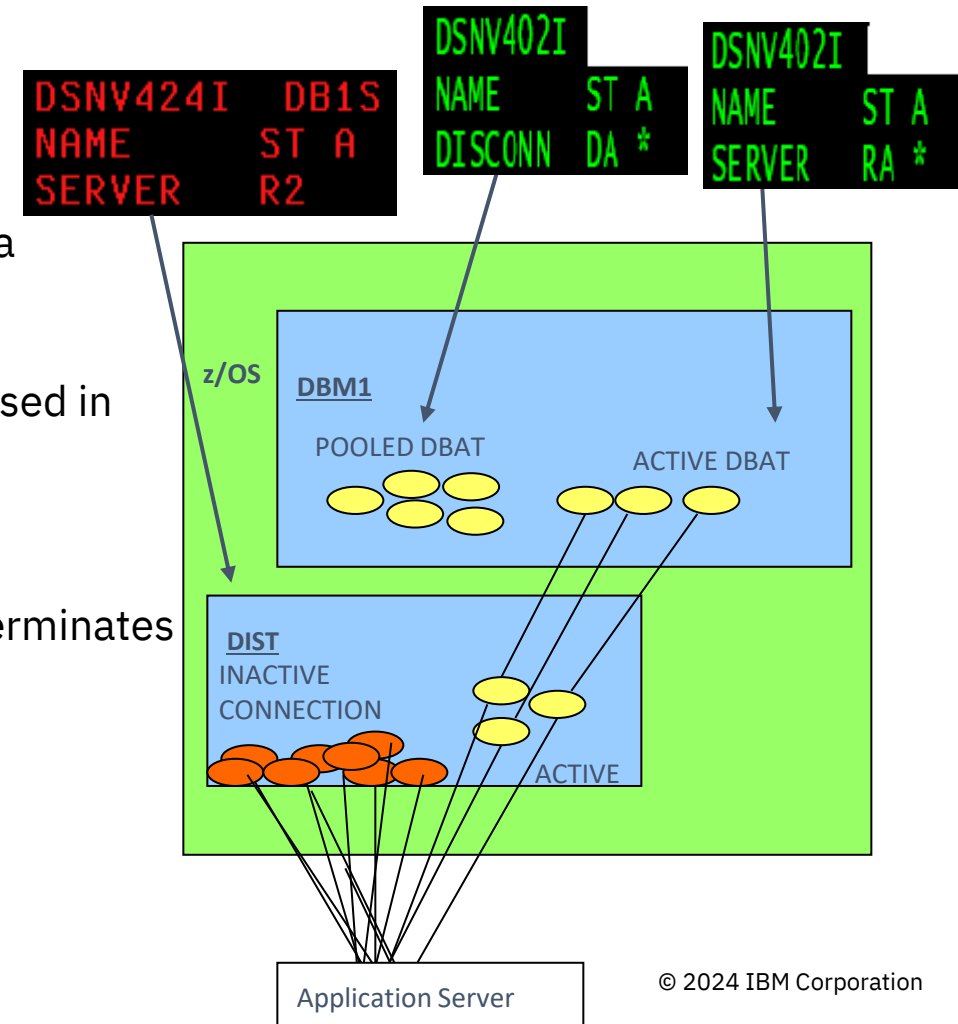
Overview and Common terms ...

- What is a database access thread (DBAT)?
 - Executed under WLM managed Enclave SRB mode
 - A DBAT is associated with a connection until the connection terminates
 - ZPARAM MAXDBAT still 200 by default
 - If MAXDBAT is hit, DSNL092I in –DIS DDF will increment and requests queue up to CONDBAT, then requests rejected
 - ZPARAM MAXCONQN limits the depth of the queue for a DBAT after MAXDBAT has been hit before being canceled/rerouted (default OFF)
 - ZPARAM MAXCONQW limits the time a DBAT request remains queued until it is canceled/rerouted (default OFF)
 - ZPARAM IDTHTOIN determines how long a thread remains active between issuances of SQL
 - Default 120 seconds
 - Memory in DBM1 (common range 12KB-2MB, or larger)
 - About 12KB in 31-bit private, the rest in 64-bit shared
 - -DIS THREAD(*) SERVICE(STORAGE) to see 31-bit usage
 - DBAT ECSA footprint
 - Db2 12 ~7k per thread
 - Db2 13 ~4k per thread



Db2 Thread pooling

- Thread pooling and inactive connection support
 - ZPARM CMTSTAT determines whether threads are disassociated with a connection at commit or at connection termination
 - CMTSTAT=INACTIVE (default)
 - At commit, threads are pooled to be reused by any new/resumed request
 - Connection becomes an inactive connection (formerly called a type 2 inactive DBAT)
 - Still holds a socket and counts against CONDBAT
 - ZPARM POOLINAC determines how long a DBAT remains unused in the pool before being terminated
 - Default 120 seconds
 - CMTSTAT=ACTIVE
 - Thread stays associated with a connection until connection terminates
 - Command -DIS THREAD(*) DETAIL for status
 - Both active and pooled DBATs count towards MAXDBAT

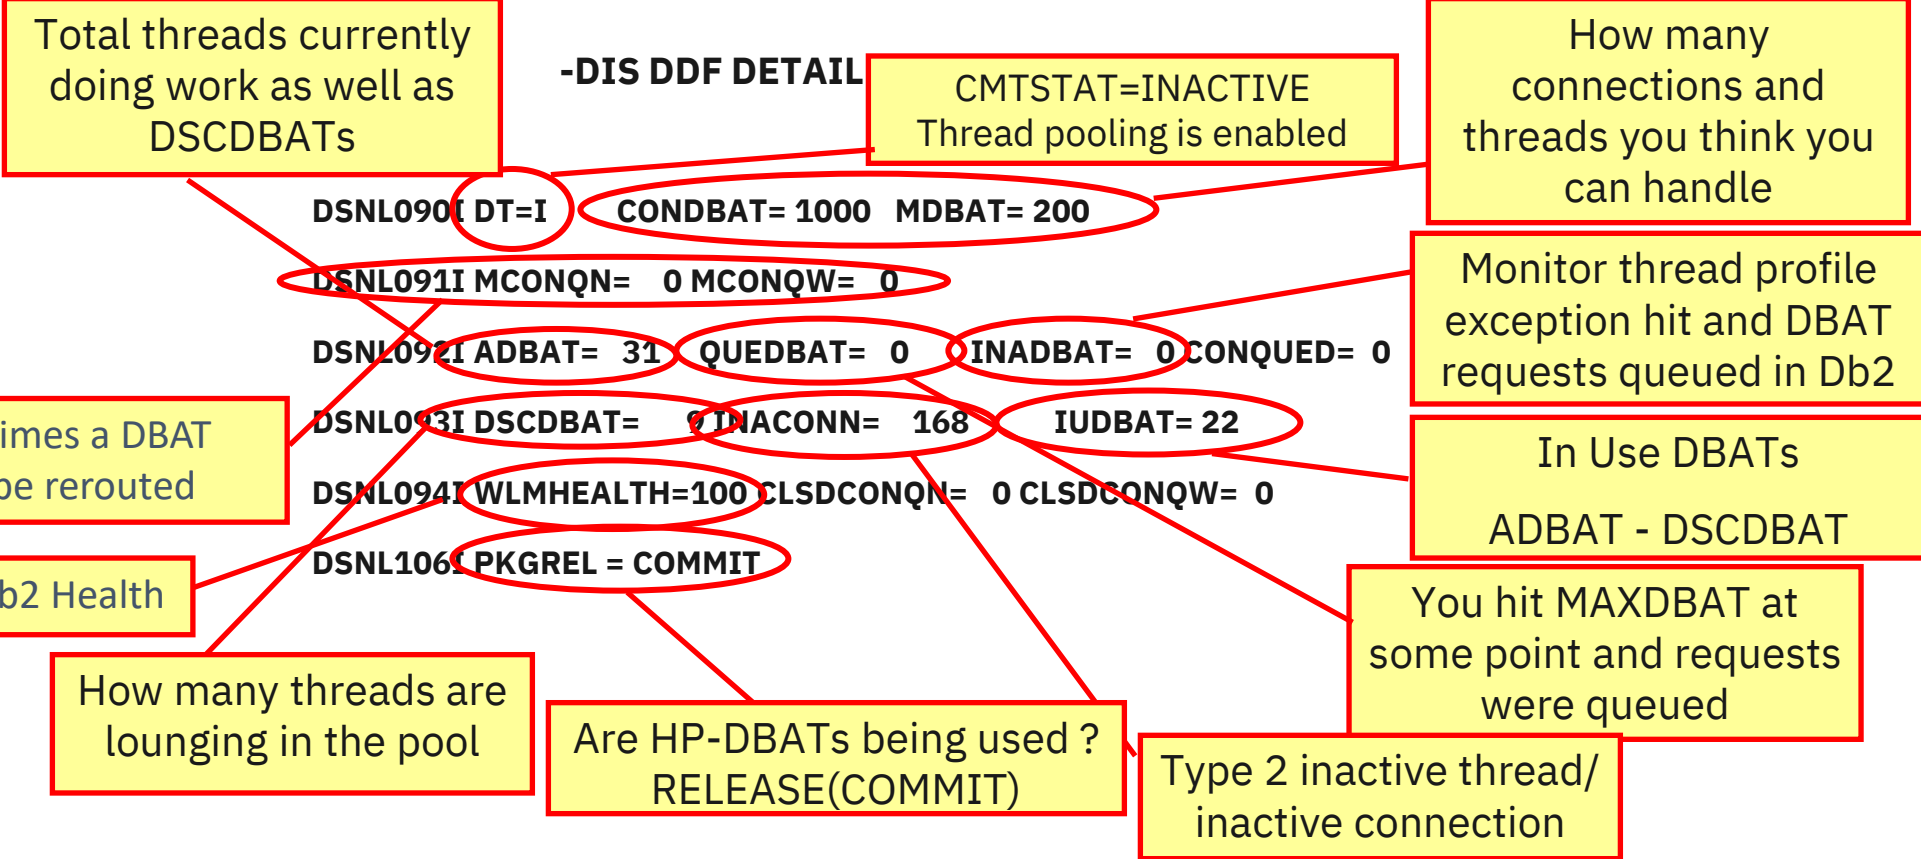


Db2 Thread pooling ...

- Recommendation is to utilize thread pooling and inactive connection support
 - Benefits for Db2 z/OS thread pooling:
 - CPU savings in Db2, by avoiding repeated creation and destruction of DBAT
 - Real memory savings in z/OS, by reducing the number of concurrent DBATs
 - Virtual memory savings in DBM1, by reducing the number of DBATs
 - Small number of threads can typically be used to service a large number of connections
 - Greater capacity to support DRDA connections avoiding need for re-authentication (SSL)
 - WLM granularity at a unit of work boundary for enclaves
 - What prevents a thread from going inactive?
 - Packages bound with KEEP DYNAMIC(YES)
 - Statement is held in local statement cache across commit boundary (SAP)
 - Held cursors and/or open LOB streams
 - Application server connection verification query (test-on-borrow) defaults to SYSLH200 package
 - Declared Global Temporary Tables result sets held beyond commit (no Drop on Commit)
 - High Performance DBATs (JCC packages bound DEALLOCATE and –MODIFY DDF issued)
 - DBAT reused 200 times (Db2 12) and 500 times (Db2 13) then recycled
 - ZPARM POOLINAC controls duration of inactivity
 - There should be a ‘separate’ NULLID collection for applications using High Performance DBATs

Db2 Thread pooling ...

- The display DDF detail command provides a snapshot of the remote work going on
 - The presence of DSCDBAT proves that thread pooling is enabled
 - Ideally QUEDBAT and CONQUED will always equal 0
 - PKGREL BNDOPT or BNDPOOL allows the use of high performance DBATs
 - INADBAT is no longer related to Private Protocol, but indicates Monitor Threads profile exceptions



Application Connection Pooling

- Optimal transaction behavior and performance relies on Db2 thread pooling in partnership with application connection pooling for a complete end to end solution
 - Pooling throughout the infrastructure will deliver
 - Efficiency
 - Pool integrity
 - Resiliency
 - Availability

WAS connection pool properties	Description
enableSysplexWLB	Specifies whether transaction-level workload balancing is in effect
Max Connections	Max connections from JVM instance
Min Connections	Lazy minimum number of connections in pool
Reap Time	How often cleanup of pool is scheduled in seconds
Unused Timeout	How long to let a connection sit in the pool unused (inactive conn in Db2)
Aged Timeout	How long to let a connection live before recycling (active or inactive conn in Db2)
Purge Policy	After StaleConnection, does the entire pool get purged or only individual connection

Db2dsdriver.cfg	Description
enableWLB parameter	Specifies whether transaction-level workload balancing is in effect
MaxTransports	Max number of connections that a requested can make to a data sharing group
maxRefreshInterval	Max elapsed time in seconds before the server list is refreshed
MaxTransportIdleTime	Max elapsed time in seconds before a transports is dropped
MaxTransportWaitTime	Number of seconds that the client waits for a transport to become available

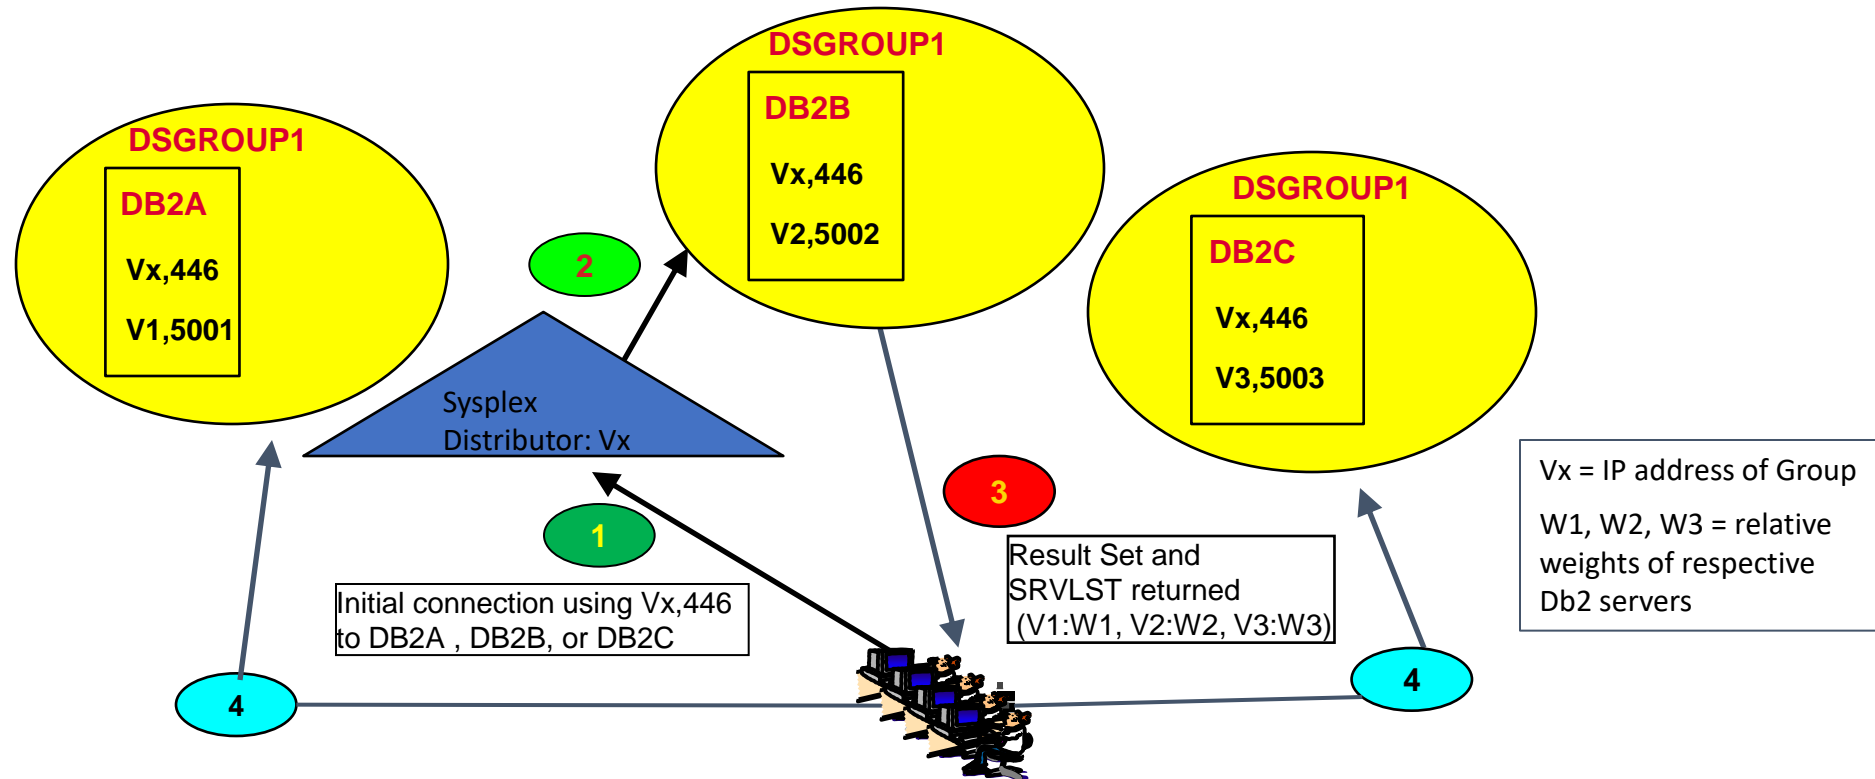
High Availability Workload Routing

- Installations are not always properly configured for high availability with regards to remote connections to a Db2 data sharing group
 - Need the ability to mask single points of failure both inside and outside of Db2
 - Planned maintenance/unplanned outages in a 24x7 environment
 - Rolling maintenance and release migrations
 - Guaranteed connection to Db2 on the first attempt
 - Applications should not need to know which member is down or when
 - Remove the need to for a “ping” transaction
 - Avoid stale connection exceptions in connection pools when Db2 bounces but app server is up
 - Distribution of connections/threads across Db2 data sharing group
 - In the event an LPAR is lost, need ability for connections to ‘follow’ Db2 to another LPAR or seamlessly route to a surviving Db2 member
- High availability requires a combination of z/OS and driver functions
 - Sysplex Distributor to guarantee the initial connection
 - Sysplex Workload Balancing to provide:
 - Automatic client reroute
 - Intelligent workload distribution at transaction level

High Availability – Sysplex Distributer

- Exploit Sysplex Distributor
 - Function of z/OS which serves to route incoming **connections** to the target LPARs
 - Configure a **Distributing Dynamic Virtual IP Address** (DDVIPA) for the Db2 data sharing group
 - IP address owned by Sysplex Distributor
 - All members listen to this IP address for the SQL port
 - Connections are distributed across all members
 - Connections are successful as long as one member is up
 - Used by clients to access the group providing a single image
 - DISTMETHOD for Sysplex Distributor determines method of routing (here are 3 basic ones)
 - ROUNDROBIN – based on order of IEFSSNxx member of PARMLIB
 - BASEWLM – based on displaceable capacity
 - SERVERWLM – based on displaceable capacity (including lower imp. work) and performance index
 - Configure a **Dynamic Virtual IP Address** (DVIPA) for each member
 - Member DVIPA is not distributed: routes connections to a specific member
 - Allows routing even if a member fails over to another LPAR using VIPA takeover
 - WLM weight and member DVIPA are provided to Db2 client driver for all registered subsystems

High Availability – Sysplex Distributer ...

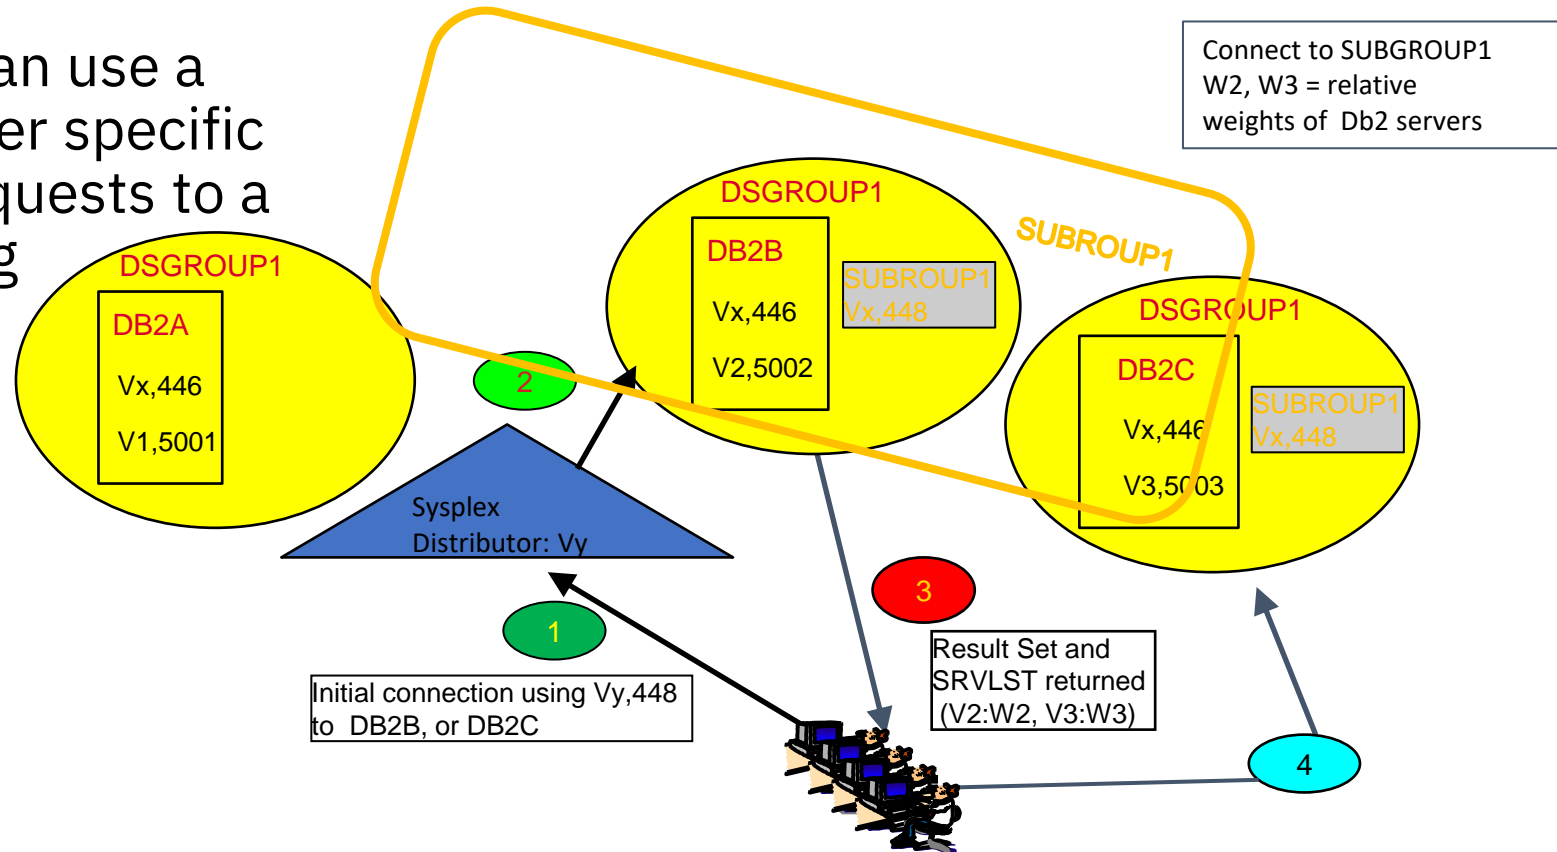


Vx = IP address of Group
W1, W2, W3 = relative weights of respective Db2 servers

- 1 First connection goes through Sysplex Distributer (DVIPA) connecting to the group DVIPA
- 2 Sysplex Distributer routes based on DISTMETHOD
- 3 Result set (and if Sysplex WLB is enabled) a list of available Db2 subsystems and their respective weightings are returned to the driver
- 4 Based on the respective weighting subsequent units of work are routed to the data sharing members (if Sysplex WLB is enabled) otherwise Sysplex Distributer decides

High Availability – Sysplex Distributor and Subgroup

- Similar to subgroup attach you can use a dynamic location alias for member specific access or to point connection requests to a subset of the members accepting DDF work

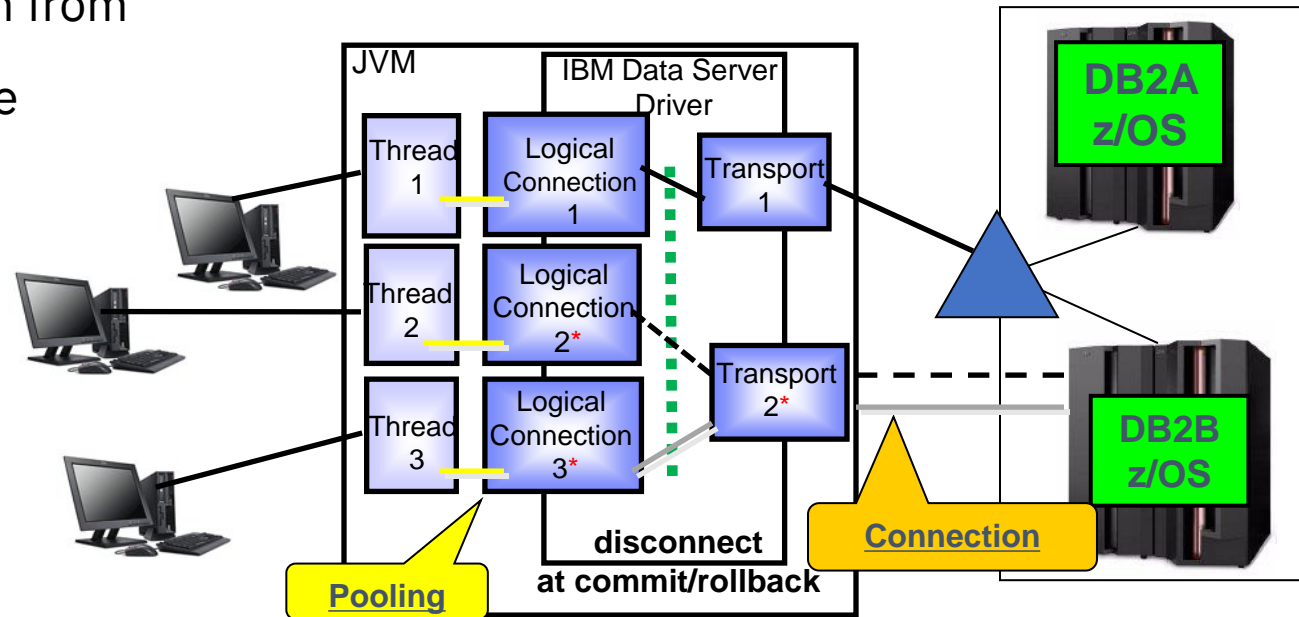


- 1 First connection goes through Sysplex Distributor (DVIPA) connecting to a subset of the group
- 2 Sysplex Distributor routes based on DISTMETHOD
- 3 Result set (and if Sysplex WLB is enabled) a list of available Db2 subsystems in that subset and their respective weightings are returned to the driver
- 4 Based on the respective weighting subsequent units of work are routed to the data sharing members (if Sysplex WLB is enabled) otherwise Sysplex Distributor decides

High Availability – Sysplex WLB

- Sysplex Workload Balancing
 - Enables transport pooling (connection concentration)

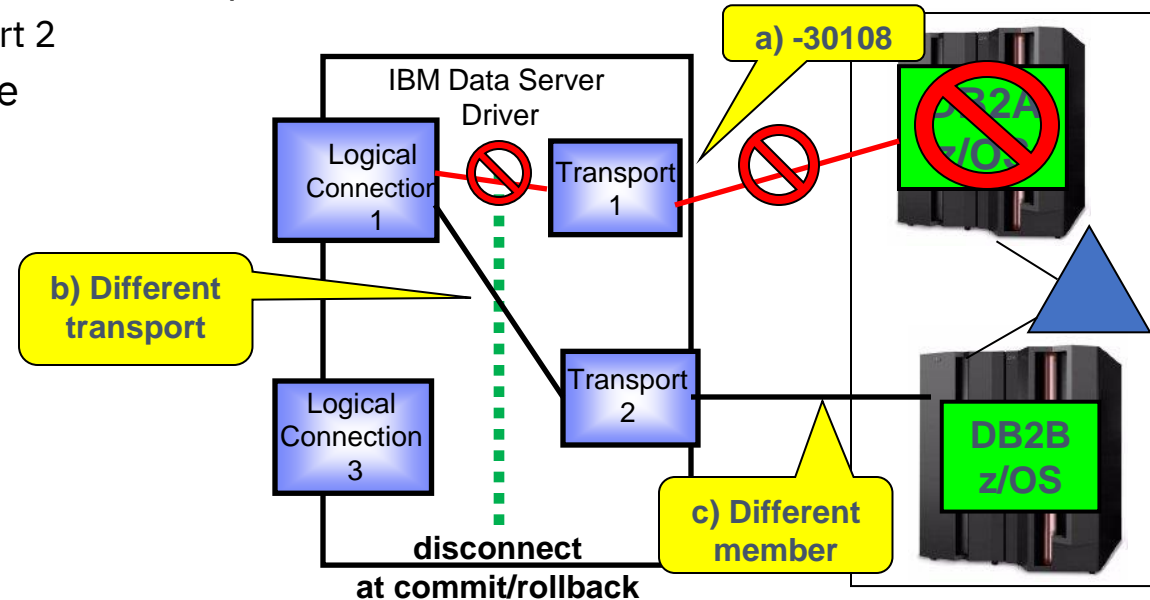
- Separates the logical connection from the application to the driver and the physical connection from the driver to Db2 for z/OS
- Global pool of Transport objects per instance
- Initial connection goes thru Sysplex Distributor, then driver relies on its WLM server list to route to the member directly
- Connection reuses/changes transport at transaction boundary (commit)



- Provides Automatic Client Reroute (ACR)
 - Masks connectivity issue with Db2 including Db2 failure/shutdown
 - See next slides....

High Availability – Sysplex WLB ...

- With ACR, application connectivity is *NOT* lost if DB2A is brought down or crashes
 - The intent is to route around connection failures and mask them from the applications
 - This also avoids stale connection exceptions in app server connection pools
 - If connectivity to DB2A is lost i.e., taken down for maintenance or crashes
 - If the first SQL statement of a transaction encounters this situation:
 - a) Driver receives network failure (-30108)
 - That transport would then be recycled to avoid being used by other requests
 - b) Driver seamlessly routes transaction to another transport on another member
 - Logical connection 1 moves to transport 2
 - c) Application receives no negative SQL code and continues processing



High Availability – Sysplex WLB ...

- Enable Sysplex Workload Balancing
 - Provides intelligent workload distribution at transaction level
 - Every 10 seconds driver polls WLM to refresh the weighted list of data sharing members
 - Driver does the routing, Sysplex Distributor only provides the first connection to the group
 - At transaction boundary next SQL can be routed to any member of the group based on 3 main factors:
 - Enclave WLM Service class goal attainment (Performance Index) and queuing
 - Displaceable CPU (GCP + zIIP) on the LPAR
 - Db2 health
 - Based on 31-bit virtual memory constraint
 - Also lowered as the number of connections in use approaches CONDBAT
 - Can be monitored with `-DIS DDF DETAIL` and `-DIS THREAD(*) TYPE(SYSTEM)`

High Availability – Sysplex WLB ...

- Db2 13 Sysplex WLB Enhancements
 - Installations require appropriate workload distribution but find it a difficult effort to get application teams to enable Db2's Sysplex high availability function for client drivers
 - Requires client configuration on multiple systems
 - Application developers are not leveraging z-specific configuration settings
 - V13R1M502 enables Db2 Sysplex Workload Balancing capabilities without requiring any application changes
 - Applications connecting to a DDF location alias can seamlessly enable or disable the SysplexWLB behavior by issuing MODIFY DDF command
 - RQSTWLB
 - MODIFY DDF RQSTWLB (location-name|alias-name)
 - Db2 returns to the remote client a request to enable Sysplex Workload Balancing
 - Client decides where to honor the request or not
 - [Driver level 11.5 FP0/JCC driver level 4.26/SQL11050](#)
 - DFLTWLB
 - MODIFY DDF DFLTWLB (location-name|alias-name)
 - Db2 honors the Sysplex Workload Balancing option requested by the client
 - Default

High Performance DBATs

- High Performance DBATs (HP-DBATs)
 - Enables DDF-using applications to achieve the CPU efficiency benefit that that comes from combining thread reuse with the RELEASE(DEALLOCATE) package bind option
 - CICS-Db2 workloads have exploited this performance tuning action successfully for many years
 - HP-DBATs have potential for significant Db2 CPU reduction (~10%+) milage may vary
 - Once established a DBAT will be dedicated to a connection until it is deallocated or reused maximum number of times
 - Db2 12 - 200
 - Db2 13 - 500
 - HP-DBAT will not go into the DBAT pool when the thread completes
 - Number of DBATs will increase, potentially substantially
 - ZPARM MAXDBAT will probably need to increase
 - SysplexWLB will continue to provides intelligent workload distribution at transaction level
 - Enablement
 - MODIFY DDF PKGREL (BNDOPT|BINDPOOL)
 - Access one package that has been bound with RELEASE(DEALLOCATE) bind option
 - Any subsequent reuse will result in the DBAT being a HP-DBAT
 - Even if packages are not RELEASE(DEALLOCATE)

High Performance DBATs ...

- High Performance DBATs (HP-DBATs) ...
 - Carefully plan and design for selective introduction of HP-DBATs
 - Incremental select additional candidates after initial success
 - Exploitation
 - Before enabling HP-DBATs all potential DDF accessed packages (Driver and Db2 Stored Procedures) should be bound with RELEASE(COMMIT) as a defensive mechanism
 - RELEASE(COMMIT) should be the designed default for Db2 Stored Procedure packages
 - Implement a new collection where a copy of the driver package are bound with RELEASE(DEALLOCATE)
 - Identify application workloads with long-lived connections from a combination of accounting trace summary and Db2 statistic trace data
 - (COMMIT RECEIVED+ABORTS RECEIVED/INTIATED FROM REMOTE SITE) > 200|500
 - Use Db2 System Monitor Profile(s) to selectively route some application workloads to the alternative collection
 - Monitor and adjust MAXDBAT upward as the demand for DBATs is very likely to increase
 - Aggressively monitor ECSA usage
 - Be prepared to use –MODIFY DDF PKGREL(COMMIT) to switch off HP-DBATs at first signs of DBAT congestion i.e., overuse of DBATs
 - WLM Service classes for HP-DBATs should be defined using execution velocity goals (PH34378)
 - Existing service classes utilizing response time goals will generally not meet their performance objectives
 - After implementing HP-DBATs Service class is no longer achieving desired Performance Index (PI) goal = .81 - 1.2
 - Effective response time of the enclave will be longer than the response time to process a single transactions when using HP-DBATs

Transaction Identification and Monitoring

- Historically a common problem has been the inability to properly identify Db2 for z/OS DRDA transactions
 - How does one differentiate remote applications?
 - User IDs were often created and distributed to application teams to access Db2
 - Application teams would use the same User ID for all workload executing from their application servers
 - Eliminating the ability to properly identify business specific transactions
 - Without proper classification numerous customers referred to Db2 connect applications as attempting to monitoring the “black box”
 - Resulting in not embracing the infrastructure
 - CICS application transactions identification has been well planned and evolved over time
 - Transaction names
 - CICS Application Owning Regions (AOR) are specific for an application
 - If a proper naming strategy was initially designed a User ID could have been an alternative to a CICS transaction
- As a result, customers had limited ability to properly monitor and tune
 - Db2 Accounting reports grouping and filtering on a User ID
 - Db2 DDF DISPLAY LOCATION DETAIL
 - Db2 Statistic – subsystem level reporting

Transaction Identification and Monitoring ...

- Db2 Accounting reports triggering off of a User ID
 - Lack of granularity due to User IDs be reused for all applications
- Db2 DDF DISPLAY LOCATION DETAIL
 - Connection level information
 - Additional parsing and formatting needed
 - Difficult to join back to application User IDs

Requester IP	Total Commits	Commits/Sec	% Trans/Sec
::xxx.xxx.xx.16	1,068,620	594	4.85%
::xxx.xxx.xx.20	1,064,755	592	4.84%
::xxx.xxx.xx.15	1,014,738	564	4.61%
::xxx.xxx.xx.34	1,064,708	592	4.84%
::xxx.xxx.xx.14	1,064,581	591	4.84%
::xxx.xxx.xx.12	1,064,545	591	4.84%
::xxx.xxx.xx.18	1,064,525	591	4.84%
::xxx.xxx.xx.17	1,064,029	591	4.83%
::xxx.xxx.xx.90	1,055,101	586	4.79%
::xxx.xxx.xx.11	1,048,961	583	4.76%
::xxx.xxx.xx.125	1,046,318	581	4.75%
::xxx.xxx.xx.122	1,028,769	572	4.67%
::xxx.xxx.xx.124	1,028,276	571	4.67%
::xxx.xxx.xx.123	1,028,223	571	4.67%
::xxx.xxx.xx.129	1,028,135	571	4.67%
::xxx.xxx.xx.127	1,028,106	571	4.67%
::xxx.xxx.xx.128	1,027,974	571	4.67%
::xxx.xxx.xx.126	1,016,620	565	4.62%
::xxx.xxx.xx.131	1,016,159	565	4.62%
::xxx.xxx.xx.130	1,015,981	564	4.62%
::xxx.xxx.xx.118	610,310	339	2.77%
::xxx.xxx.xx.109	565,159	314	2.57%
Grand Total	22,014,593	12,230	100.00%

Auth ID	Total Commits	Commits/Sec	% Trans/Sec
ONLMAIN	11,889,124	6,605	54.01%
ONLPPP	4,183,889	2,324	19.01%
ONLMAN	2,069,390	1,150	9.40%
ONLXAP	1,100,582	611	5.00%
ONLALC	1,025,728	570	4.66%
ONLYEP	616,775	343	2.80%
ONLTRE	489,097	272	2.22%
ONLIOC	217,873	121	0.99%
ONLGMC	102,824	57	0.47%
ONLPGA	101,675	56	0.46%
ONLMLB	81,398	45	0.37%
ONLCIC	62,061	34	0.28%
ONLNBA	58,254	32	0.26%
ONLNCAA	10,456	6	0.05%
ONLMOUNT	5,467	3	0.02%
Grand Total	22,014,593	12,230	100.00%

```

DSN
  -DISPLAY LOCATION DETAIL
DSNL200I  -D21P DISPLAY LOCATION REPORT FOLLOWS-
1LOCATION                                PRDID      T  ATT  CONNS
::xx.xxx.xxx.35                          SQL10055  G           0
::xx.xxx.xxx.38                          SQL10055  S           0
::xx.xxx.xxx.40                          SQL10055  G           2
                                           WLB         2
::xx.xxx.xxx.41                          SQL10055  S           9
::xx.xxx.xxx.2                            SQL11013  G           2
L209-FUNCTIONAL LEVEL=s1803021700
                                           WLB         2
::xx.xxx.xxx.3                            SQL11013  G           2
L209-FUNCTIONAL LEVEL=s1803021700
                                           WLB         2
::xx.xxx.xxx.78                          SQL11054  G           4
L209-FUNCTIONAL LEVEL=s2006161200
                                           TLS         4
::xx.xxx.xxx.81                          SQL11054  G           0
L209-FUNCTIONAL LEVEL=s2006161200
::xx.xxx.xxx.82                          SQL11054  G           1
L209-FUNCTIONAL LEVEL=s2006161200
                                           TLS         1
    
```

Transaction Identification and Monitoring ...

- Db2 Statistic – Subsystem level reporting

CUR ACTIVE AND DISCON DBATS

Current # of DBATs (assigned and pooled)

CUR DISCON DBATS NOT IN USE

Current # of pooled DBATs available for reuse

CUR INACTIVE CONNS (TYPE 2)

Current # of inactive connections

DBAT/CONN QUEUED-MAX ACTIVE

of times MAXDBAT was reached

CONN REJECTED-MAX CONNECTED

of times CONDBAT was reached

CONN CLOSED - MAX QUEUED

of times MAXCONQN was reached

CUR QU INACT CONNS (TYPE 2)

Current # of connections requests queued for DBAT

ACC QU INACT CONNS (TYPE 2)

of resumed connection requests

DBATS CREATED vs. DISCON (POOL) DBATS REUSED

Indicator of DBAT pooling efficiency

CUR ACTIVE DBATS-BND DEALLC

Current # of High Performance DBATs available for reuse

CUR ACTIVE DBATS-BND KEEPDPY

Current # of KDR DBATs

GLOBAL DDF ACTIVITY	QUANTITY	/SECOND ...
DBAT/CONN QUEUED-MAX ACTIVE	0.00	0.00
CONN REJECTED-MAX CONNECTED	0.00	0.00
CONN CLOSED - MAX QUEUED	0.00	0.00
CONN CLOSED - MAX WAIT	0.00	0.00
COLD START CONNECTIONS	0.00	0.00
WARM START CONNECTIONS	0.00	0.00
RESYNCHRONIZATION ATTEMPTED	0.00	0.00
RESYNCHRONIZATION SUCCEEDED	0.00	0.00
CUR TYPE 1 INACTIVE DBATS	0.00	N/A
HWM TYPE 1 INACTIVE DBATS	4.00	N/A
TYPE 1 CONNECTIONS TERMINAT	0.00	0.00
CUR INACTIVE CONNS (TYPE 2)	7726.93	N/A
HWM INACTIVE CONNS (TYPE 2)	9744.00	N/A
ACC QU INACT CONNS (TYPE 2)	1712.5K	481.86
CUR QU INACT CONNS (TYPE 2)	21.38	N/A
MIN QUEUE TIME	0.000006	N/A
MAX QUEUE TIME	54.04610	N/A
AVG QUEUE TIME	0.138253	N/A
HWM QU INACT CONNS (TYPE 2)	660.00	N/A
CUR ACTIVE AND DISCON DBATS	879.34	N/A
HWM ACTIVE AND DISCON DBATS	1855.00	N/A
HWM TOTL REMOTE CONNECTIONS	9964.00	N/A
CUR DISCON DBATS NOT IN USE	290.98	N/A
HWM DISCON DBATS NOT IN USE	1658.00	N/A
DBATS CREATED	4581.00	N/A
DISCON (POOL) DBATS REUSED	1708.9K	N/A
DBATS TERM SINCE DDF START	0.00	N/A
DBATS TERM-POOLINAC	N/A	N/A
DBATS TERM-REUSE LIMIT	N/A	N/A

Transaction Identification and Monitoring ...

- Db2 Statistic – Subsystem level reporting ...

CUR ACTIVE DBATS-BND DEALLC

Current # of High Performance DBATs available for reuse

CUR ACTIVE DBATS-BND KEEPDY

Current # of Keep Dynamic DBATs

GLOBAL DDF ACTIVITY	QUANTITY	/SECOND
...		
CUR ACTIVE DBATS-BND DEALLC	84.28	N/A
HWM ACTIVE DBATS-BND DEALLC	169.00	N/A
CUR ACTIVE DBATS-BND KEEPDY	0.00	N/A
HWM ACTIVE DBATS-BND KEEPDY	0.00	N/A
ILOS CANCELS DECLINED	6681.00	N/A
CUR DBATS SUSPND PROF EXCEP	N/A	N/A
HWM DBATS SUSPND PROF EXCEP	N/A	N/A

ACTIVE CONNS FROM LOC

Snapshot when the record is cut of # of connections

ACTIVE CONNS FROM LOC – IN. HWM

HWM in this interval of # of connections

ACTIVE DBATS FOR LOC – SNAP

Snapshot when the record is cut of # DBATs

ACTIVE DBATS FOR LOC – INT. HWM

HWM in the interval of #DBATs

Great information for monitoring spiky behavior

DRDA Remote Location Statistics

DRDA REMOTE LOCS	CLIENT	CONDITIONS	QUANTITY	DRDA REMOTE LOCS	REQUEST TERMINATION	QUANTITY
WITH HOLD CURSOR NOT CLOSED			1988.00	CONDBAT REACHED		0.00
DGTT NOT DROPPED			1.00	PROFILE EXCEPTION		0.00
KEEPDYNAMIC PACKAGES USED			24.00	MAXCONQN REACHED		0.00
HIGH PERF DBATS USED			21004.00	MAXCONQW REACHED		0.00
HELD LOB LOCATORS EXIST			0.00			
SP COMMIT PERFORMED			0.00			
DRDA REMOTE LOCS	THREAD	TERMINATION	QUANTITY	<div style="border: 1px solid black; padding: 5px;"> <p>Any Limits were reached</p> <p>Thread were recycled/cancelled</p> <p>Why a DBAT could not be pooled</p> </div>		
QUEUED – PROFILE EXCEPTION			0.00			
TERMINATED – PROF EXCPTION			0.00			
ABENDED			0.00			
CANCELED			108.00			
TERMINATED – POOLINAC TIME			N/A			
TERMINATED – REUSE LIMIT			N/A			

DRDA REMOTE LOCS	DDF CONN. DETAILS	QUANTITY
ACTIVE CONNS FROM LOC – SNAP		7996.00
ACTIVE CONNS FROM LOC – INT. HWM		8104.00
ACTIVE DBATS FOR LOC – SNAP		204.00
ACTIVE DBATS FOR LOC – INT. HWM		347.00
REST SERVICE REQUESTS		0.00
PROFILE SET SPECIAL REGS		28.00
PROFILE SET GLOBAL VARS		0.00
SYSPLEX WKLD BAL. USED		131.00
TLS/SSL USED		0.00
TRUSTED CONTEXT USED		0.00
AES ENCRYPTION USED		0.00
XA GLOBAL TRANSACTION USED		0.00
DRDA ENCRYPTION USED		0.00
UID/PWD USED		1256.00
KERBEROS USED		0.00
MFA USED		0.00
CLIENT CERTIFICATE USED		0.00
FAILED SECURITY AUTH		34.00

Transaction Identification and Monitoring ...

- How does one differentiate remote applications coming in when User IDs are saturated and many applications executing on the same servers? ...
 - The application can set these client information fields using JAVA APIs, set methods for .NET and CLI, data source properties, driver connection properties, and Db2 special registers
 - More granular workload management classifications for relative importance levels as well as software pricing discounts (i.e., Mobile workload pricing) as well as internal chargeback
 - Simpler problem determination and exposure of tuning possibilities
 - Break-down of what the most common fields are and where their values are found

English	JAVA	.NET or .CFG keyword	WLM Classification	Monitor Profile	Special Register
Client Accounting Information	ClientAccountingInformation	ClientAccountingString	CAI		CLIENT_ACCTNG
Correlation ID	ClientProgramName		CI		
Client CorrelationToken	ClientCorrelationToken	ClientCorrelationToken			CLIENT_CORR_TOKEN
Client End User ID	ClientUser	ClientUserID	CUI	Y	CLIENT_USERID
Client End User Workstation Name	ClientHostName	ClientWorkstationName	CWN	Y	CLIENT_WRKSTNNAME
Client End User Transaction Name	ApplicationName	ClientApplicationName	CTN	Y	CLIENT_APPLNAME
Process Name (client tran or appl name)	ApplicationInformation	ClientApplicationName	PC	Y	CLIENT_APPLNAME

Transaction Identification and Monitoring ...

- Additional Db2 Accounting granularity
 - Provides enhanced reporting above and beyond the use of a User ID
 - Application function
 - Enhanced reporting
 - Remote IP can be used to join with statistic data

ENDUSER TRANS NAME	Commits Received	Commits/Sec	% Trans/Sec
BANK_ACCT_QUERY	2,119,185	1,177	9.63%
ACCOUNT_PAYMENTS	1,658,805	922	7.54%
CUSTOMER_EBILLS	1,653,078	918	7.51%
BALANCE	1,509,468	839	6.86%
AUTO_PAYMENTS	1,234,288	686	5.61%
ACCOUNT_HISTORY	1,126,776	626	5.12%
BILL_PAY_HISTORY	750,419	417	3.41%
ONLINE_PAYMENTS	746,468	415	3.39%
AUTO_DEPOSITS	721,048	401	3.28%
REPEATING_PAYMENTS	710,699	395	3.23%
DIRECT_PAYMENTS	567,572	315	2.58%
CONSUMDER_NAME	481,934	268	2.19%
CC_PAYMENTS	464,846	258	2.11%
CC_HISTORY	425,516	236	1.93%
DEBIT_CARD_PAYMENTS	412,821	229	1.88%
AUTO_WITHDRAWLS	403,720	224	1.83%
AUTO_DEPOSITS	339,270	188	1.54%
BANK_ACCT_2YEAR_QUERY	329,388	183	1.50%
Grand Total	22,014,593	12,230	100.00%

USER_ID - ONLMAIN				
ENDUSER TRANS NAME	Tot CL2CPU	Avg CL2CPU	COMMITs/sec	zIIP Offload
BANK_ACCT_QUERY	311.12	0.000072	1,193.31	58.16%
::xxx.xxx.xx.126	16.30	0.000072	62.57	58.87%
::xxx.xxx.xx.130	16.27	0.000073	62.28	57.75%
ACCOUNT_PAYMENTS	1,350.88	0.000407	943.59	59.06%
::xxx.xxx.xx.12	70.52	0.000395	49.52	58.79%
::xxx.xxx.xx.16	70.46	0.000395	49.41	58.95%
CUSTOMER_EBILLS	597.63	0.000178	930.92	59.14%
::xxx.xxx.xx.129	30.16	0.000180	46.60	58.52%
::xxx.xxx.xx.123	30.16	0.000180	46.60	59.60%
AUTO_PAYMENTS	499.22	0.000206	672.51	58.77%
::xxx.xxx.xx.125	27.38	0.000203	37.45	59.90%
::xxx.xxx.xx.12	27.19	0.000205	36.76	58.46%
BILL_PAY_HISTORY	417.88	0.000282	411.98	60.44%
::xxx.xxx.xx.130	21.85	0.000285	21.39	59.89%
::xxx.xxx.xx.125	21.68	0.000282	21.25	58.05%
CC_HISTORY	1,272.96	0.001382	256.80	58.52%
::xxx.xxx.xx.20	67.49	0.001475	12.82	60.83%
::xxx.xxx.xx.14	66.03	0.001445	12.74	63.51%
AUTO_WITHDRAWLS	394.28	0.000486	225.09	59.62%
::xxx.xxx.xx.125	20.12	0.000492	11.34	60.12%
::xxx.xxx.xx.15	19.99	0.000490	11.29	59.43%
DEBIT_CARD_PAYMENTS	262.93	0.000319	228.59	59.67%
::xxx.xxx.xx.130	13.41	0.000324	11.49	57.55%
::xxx.xxx.xx.131	13.34	0.000323	11.40	60.16%

Transaction Identification and Monitoring ...

- Db2 DRDA Statistical Instrumentation Enhancements (Db2 SMF 102 records)
 - IFCID 365 - Statistics Class 7
 - Provides detailed statistics about the remote locations with which a Db2 subsystem communicates using the DRDA protocol

TIMESTAMP	REMOTE LOCATION	PRDID	INITIATED FROM	COMMITTS	ABORTS	REST	SysplexWLB	XA TRANS	CSR WITH HOLD	DGTT NOT	KEEPDYNAMIC	HP-DBATS	PROFILE	MAXCONQN	CONNECTIONS	THREADS
			REMOTE SITE	RECEIVED	RECEIVED	REQUESTS			NOT CLOSE	DROPPED			EXCEPTIONS	REACHED	HWM	HWM
2023-09-15-10.00.00.542139	::xxx.xxx.xx.11	SQL11055	10,000,145	35,468,515	23,173,988	0	0	0	195,849	0	0	0	0	0	15	2
2023-09-15-10.00.00.542139	::xxx.xxx.xx.12	SQL11055	10,001,456	34,924,033	22,722,438	0	0	0	187,570	0	0	0	0	0	12	2
2023-09-15-10.00.00.542139	::xxx.xxx.xx.122	SQL11055	10,100,357	36,383,864	23,815,120	0	0	0	196,028	0	0	0	0	0	13	4
2023-09-15-10.00.00.542139	::xxx.xxx.xx.123	SQL11055	10,030,555	34,276,691	22,442,466	0	0	0	191,602	0	0	0	0	0	9	2
2023-09-15-10.00.00.542139	::xxx.xxx.xx.124	SQL11055	10,000,569	39,837,342	26,119,479	0	0	0	213,017	0	0	0	0	0	6	0

- IFCID 411 - Statistics Class 10
 - Records detailed statistics about remote applications that connect to a local Db2 subsystem, based on the CURRENT CLIENT_APPLNAME special register

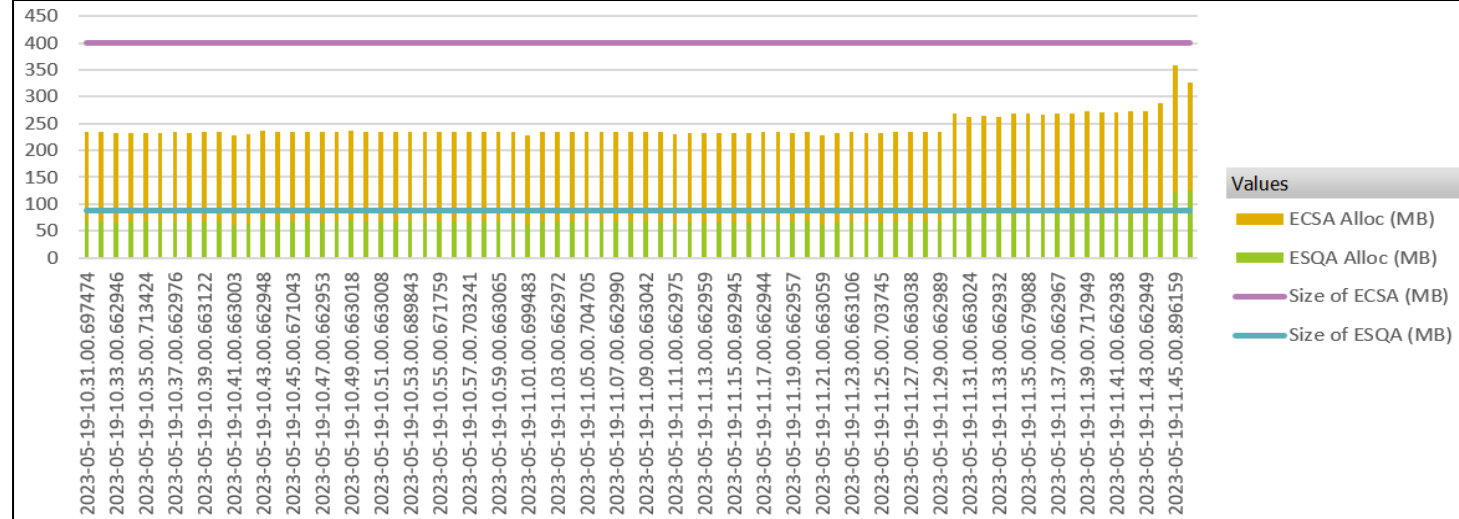
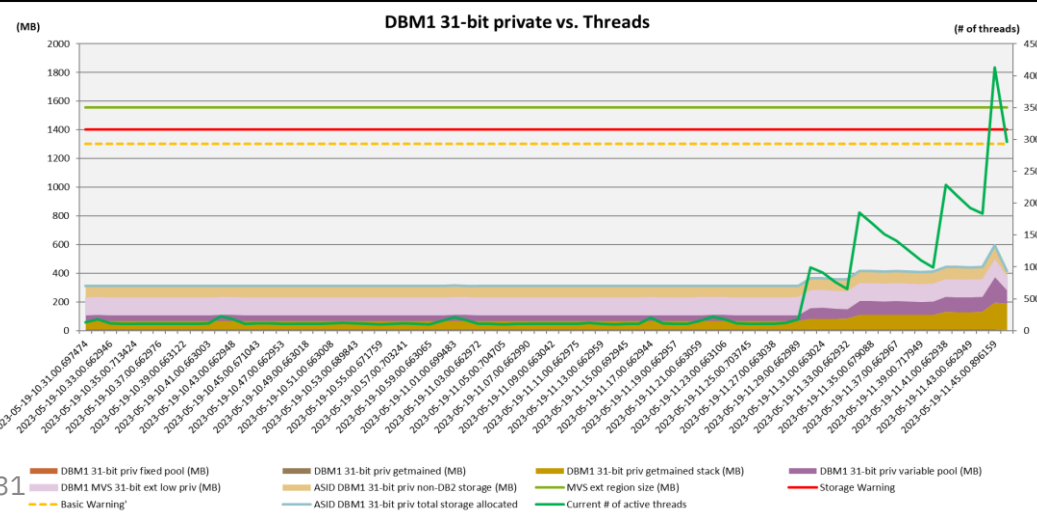
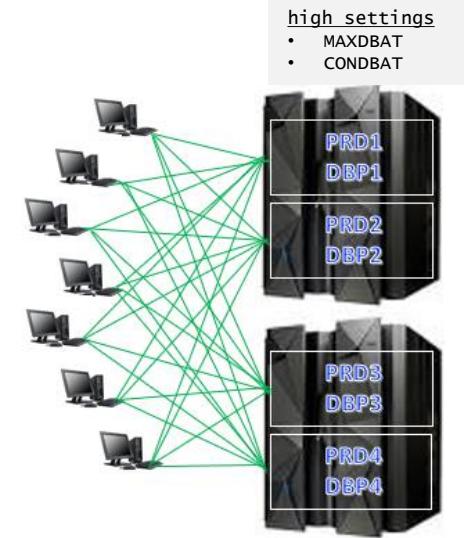

TIMESTAMP	APPLICATION_NAME	PRDID	INITIATED FROM	COMMITTS	ABORTS	REST	SysplexWLB	XA TRANS	CSR WITH HOLD	DGTT NOT	KEEPDYNAMIC	HP-DBATS	PROFILE	MAXCONQN	CONNECTIONS	THREADS
			REMOTE SITE	RECEIVED	RECEIVED	REQUESTS			NOT CLOSE	DROPPED			EXCEPTIONS	REACHED	HWM	HWM
2023-09-15-10.00.00.542139	BANK_ACCT_QUERY	SQL11055	10,000,145	35,468,515	23,173,988	0	0	0	195,849	0	0	0	0	0	15	2
2023-09-15-10.00.00.542139	ACCOUNT_PAYMENTS	SQL11055	10,001,456	34,924,033	22,722,438	0	0	0	187,570	0	0	0	0	0	12	2
2023-09-15-10.00.00.542139	CUSTOMER_EBILLS	SQL11055	10,100,357	36,383,864	23,815,120	0	0	0	196,028	0	0	0	0	0	13	4
2023-09-15-10.00.00.542139	BALANCE	SQL11055	10,030,555	34,276,691	22,442,466	0	0	0	191,602	0	0	0	0	0	9	2
2023-09-15-10.00.00.542139	AUTO_PAYMENTS	SQL11055	10,000,569	39,837,342	26,119,479	0	0	0	213,017	0	0	0	0	0	6	0

- IFCID 412 - Statistics Class 11
 - Records statistics about the client user IDs that are associated with remote connections to a local Db2 subsystem, based on the CURRENT CLIENT_USERID special register

TIMESTAMP	USER_ID	PRDID	INITIATED FROM	COMMITTS	ABORTS	REST	SysplexWLB	XA TRANS	CSR WITH HOLD	DGTT NOT	KEEPDYNAMIC	HP-DBATS	PROFILE	MAXCONQN	CONNECTIONS	THREADS
			REMOTE SITE	RECEIVED	RECEIVED	REQUESTS			NOT CLOSE	DROPPED			EXCEPTIONS	REACHED	HWM	HWM
2023-09-15-10.00.00.542139	ONLMAIN	SQL11055	10,000,145	35,468,515	23,173,988	0	0	0	195,849	0	0	0	0	0	15	2
2023-09-15-10.00.00.542139	ONLPPP	SQL11055	10,001,456	34,924,033	22,722,438	0	0	0	187,570	0	0	0	0	0	12	2
2023-09-15-10.00.00.542139	ONLMAN	SQL11055	10,100,357	36,383,864	23,815,120	0	0	0	196,028	0	0	0	0	0	13	4
2023-09-15-10.00.00.542139	ONLXAP	SQL11055	10,030,555	34,276,691	22,442,466	0	0	0	191,602	0	0	0	0	0	9	2
2023-09-15-10.00.00.542139	ONLALC	SQL11055	10,000,569	39,837,342	26,119,479	0	0	0	213,017	0	0	0	0	0	6	0

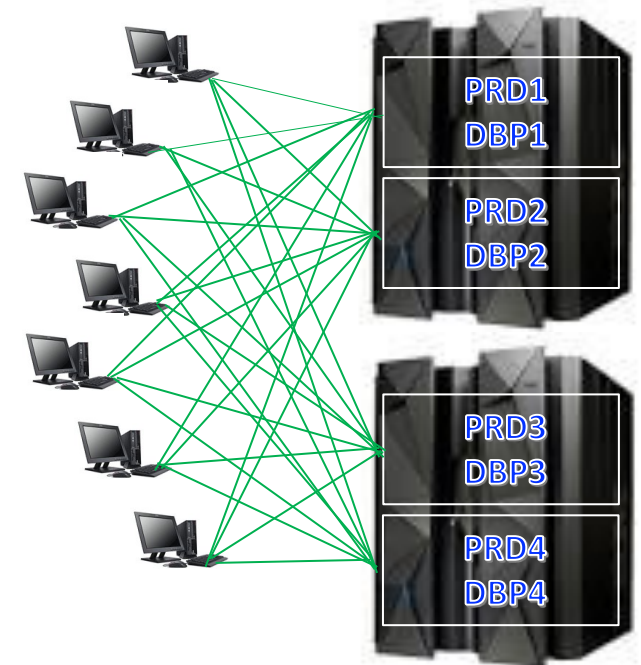
Configure for Resiliency

- Commonly, installations do not set any throttles to limit the number of DBATs from entering into Db2 – “Wild West”
 - Applications expect full access without limitations
 - Without limitations Db2 for z/OS can be compromised
 - Unlimited number of threads allowed into the system
 - Eventually memory will be exhausted
 - Elongated transactions
 - Potential consequences
 - Db2 member crash
 - z/OS LPAR crash
 - Db2 group wide outage



Configure for Resiliency ...

- Protect the System
 - How many DBATs can concurrently execute in a single Db2 member?
 - Calculate the DBAT budget for each member
 - 31-bit below the line thread agent footprint
 - ESCA footprint
 - Set MAXDBAT slightly lower than the maximum threads for the calculated budget
 - Memory in DBM1 (common range 12KB-2MB, or larger)
 - About 12KB in 31-bit private, the rest in 64-bit shared
 - DIS THREAD(*) SERVICE(STORAGE) to see 31-bit usage
 - ECSA
 - Db2 12 (7kb * MAXDBAT(5000)) - example (7kb * 5,000) = ~35mb
 - Db2 13 (4kb * MAXDBAT(5000)) - example (4kb * 5,000) = ~20mb
 - Set CONDBAT
 - (total remote IPs x max connection pool) x 1.25
- Strategic Topology is a type 4 or type 4-like DRDA access
 - Central Db2 Connect Gateway Server “Hub” is now removed
 - How do I manage all of the connections?
 - How to I protect mission critical applications from less important workload?



Db2 System Monitor Profiles

- Introduced as long ago as Db2 9 and is continually enhanced in each release and in the service stream since from customer feedback and usage
- Provides a central control point to set various ZPARMs at a more granular level than what is set at the subsystem level
 - Use Cases
 - Reserve connection and threads for critical applications
 - Prioritize and assign resources to critical workload
 - Place limits/throttles on less important workload
 - Protect the system from any unexpected rogue access
 - Denial-of-service attack characteristics
 - Ability to supply appropriate settings to support a variety of applications
 - Manage migration to a new driver
 - Resources to manage
 - Connections
 - Threads
 - Special Registers
 - Db2 13 support of local threads
 - Modes
 - Warnings
 - Exceptions

Db2 System Monitor Profiles ...

- Profile Tables

- SYSIBM.DSN_PROFILE

- Defines the profile and filtering combinations

SYSIBM.DSN_PROFILE_TABLE								
AUTHID	LOCATION	PRDID	COLLID	CLIENT_APPLNAME	CLIENT_USERID	CLIENT_WRKSTNNAME	PROFILEID	PROFILE_ENABLED
ONLMAIN							13	Y
ONLMLB							15	Y
	::xxx.xxx.xx.11						23	Y
	::xxx.xxx.xx.118						25	Y
				BANK_ACCT_QUERY			33	Y
				CC_PAYMENTS			35	Y

- SYSIBM.DSN_ATTRIBUTES

- Defines actions and settings

SYSIBM.DSN_PROFILE_ATTRIBUTES				
PROFILEID	KEYWORDS	ATTRIBUTE1	ATTRIBUTE2	Attribute Timestamp
13	MONITOR THREADS	WARNING	200	2023-05-19-06.00.00.542139
15	MONITOR THREADS	EXCEPTION	20	2023-05-19-06.00.00.542139
23	MONITOR IDLE THREADS	WARNING	180	2023-05-19-06.00.00.542139
25	MONITOR CONNECTIONS	WARNING	400	2023-05-19-06.00.00.542139
33	MONITOR THREADS	WARNING	150	2023-05-19-06.00.00.542139
35	MONITOR THREADS	EXCEPTIONS	5	2023-05-19-06.00.00.542139

- Wildcarding examples

Wildcard Support: Examples								
PROFILEID	LOCATION	PRDID	AUTHID	COLLID	PKGNAME	CLIENT_APPLNAME	CLIENT_USERID	CLIENT_WRKSTNNAME
100		JCC03*						
101	::xxx.xxx.xx.118/24							
102			ONL*					
103						BANK_ACCT_QUERY		
104							*	
105	0.0.0.0							

USER_ID	CONNECTIONS	THREADS
	HWM	HWM
ONLMAIN	15	2
ONLPPP	12	2
ONLMAN	13	4
ONLXAP	9	2
ONLALC	6	0
ONLYEP	8	2
ONLTRE	14	2
ONLIOC	15	3
ONLGMC	10	2
ONLPGA	14	2
ONLMLB	14	2
ONLCIC	11	2
ONLNBA	4	0
ONLNCAA	15	2
ONLMOUNT	6	3

REMOTE LOCATION	CONNECTIONS	THREADS
	HWM	HWM
::xxx.xxx.xx.11	15	2
::xxx.xxx.xx.12	12	2
::xxx.xxx.xx.122	13	4
::xxx.xxx.xx.123	9	2
::xxx.xxx.xx.124	6	0
::xxx.xxx.xx.125	8	2
::xxx.xxx.xx.128	14	2
::xxx.xxx.xx.129	15	3
::xxx.xxx.xx.130	10	2
::xxx.xxx.xx.131	14	2
::xxx.xxx.xx.14	14	2
::xxx.xxx.xx.15	11	2
::xxx.xxx.xx.16	4	0
::xxx.xxx.xx.17	15	2
::xxx.xxx.xx.18	6	3

- How do I determine the values

- Intellectual knowledge
 - IFCID 365, 411, 412
 - Use Db2 System Monitoring Profiles in warning mode
 - Db2 AI for z/OS – Distributed Connection Control (DCC)

Questions



Thank You

The image features the words "Thank You" in a bold, three-dimensional, metallic grey font. The text is centered and sits on a highly reflective, white surface that creates a clear mirror image of the letters below. The background is a light grey and white geometric pattern of overlapping squares and rectangles, some of which are semi-transparent, creating a layered, architectural effect. The lighting is soft and even, highlighting the 3D nature of the text.