DB2® for z/OS: Data Sharing Update

Mark Rader
IBM Corporation

August 4, 2014
Session 15940

www.SHARE.org

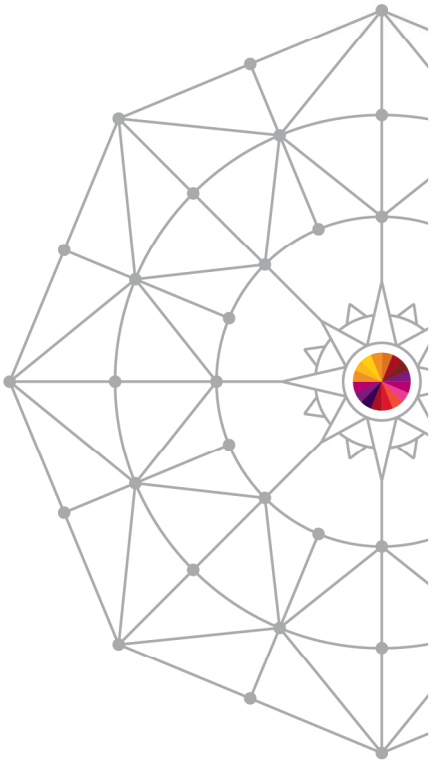

#SHAREorg

SHARE in Pittsburgh 2014

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.

Agenda

- DB2 Data Sharing Overview
- DB2 10 Data Sharing Enhancements
- DB2 11 Data Sharing Enhancements
- Resources

Assumptions

- DB2 Data Sharing and Parallel Sysplex provide the infrastructure for the highest levels of availability and scalability in the industry
  - High availability + continuous operations = Continuous availability
- The audience for this presentation is familiar with Parallel Sysplex and DB2 Data Sharing components and benefits
  - Please refer to ‘Resources’ if you feel you need more introductory material on DB2 data sharing

DB2 Data Sharing Overview

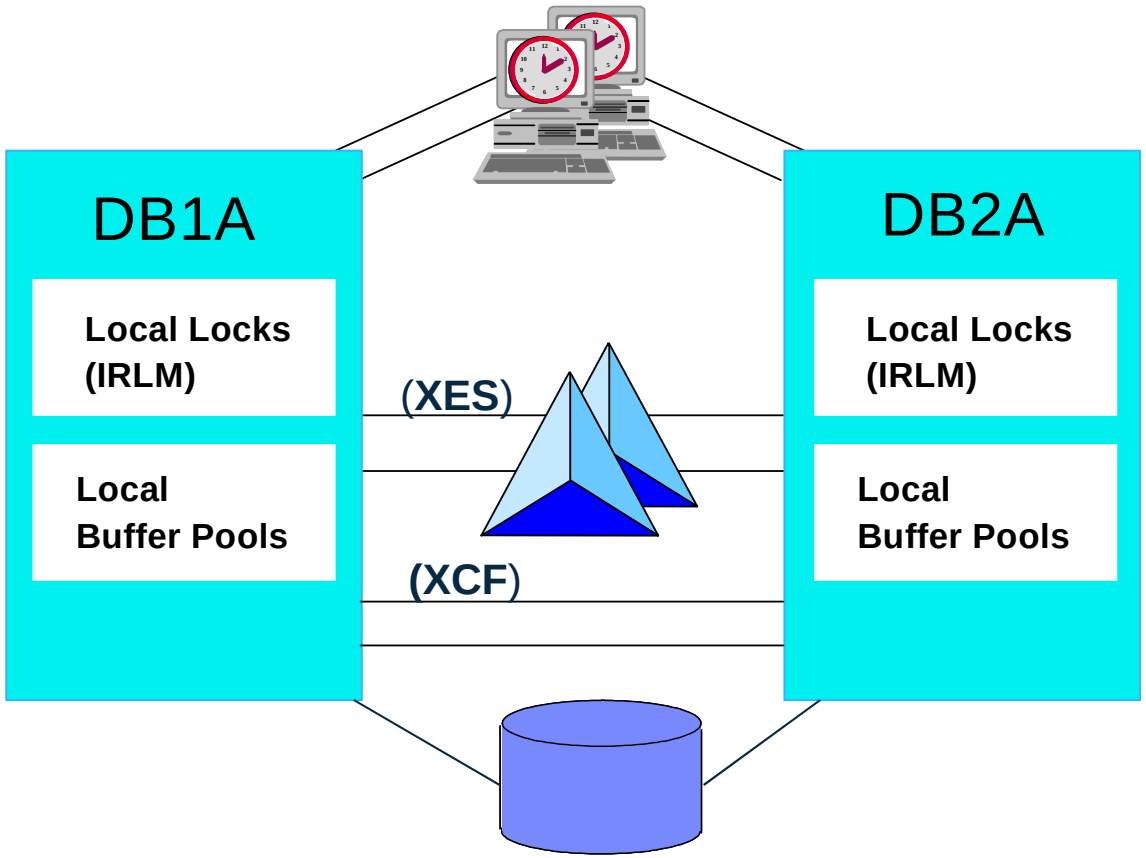

Sysplex Couple Data Set

Coupling Facility Resource Manager (CFRM) policy
Automatic Restart Management (ARM) policy
Sysplex Failure Management (SFM) policy

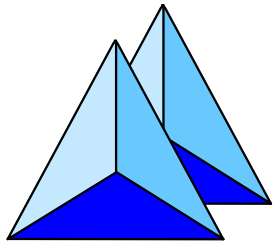

Coupling Facilities (CFs)

CFCC: Coupling Facility Control Code

Can be internal (ICF) or external

Duplexing options
- User managed (GBPs)
  Recommended
- System managed (LOCK1, SCA)

High speed links allow for synchronous operations

Cross System Coupling Facility (XCF)
Define groups, signaling between members, status monitoring.

Cross System Extended Services (XES)
for z/OS to access CF structures

DB2 Data Sharing Overview, cont.

- Sysplex Timers (now STP: Sysplex Timer Protocol)
- Coupling Facilities
  - Lock structures
  - Group Buffer Pools
  - Communications
- Coupling Links
- Up to 32 DB2 Members
- CF use based on Inter-DB2 interest
- LRSN used for log merge
- Shared Catalog/Directory
- Separate BSDS, Logs
- Shared disk
- Dynamic VIPA (DVIPA)
  - Best practice for TCP/IP access

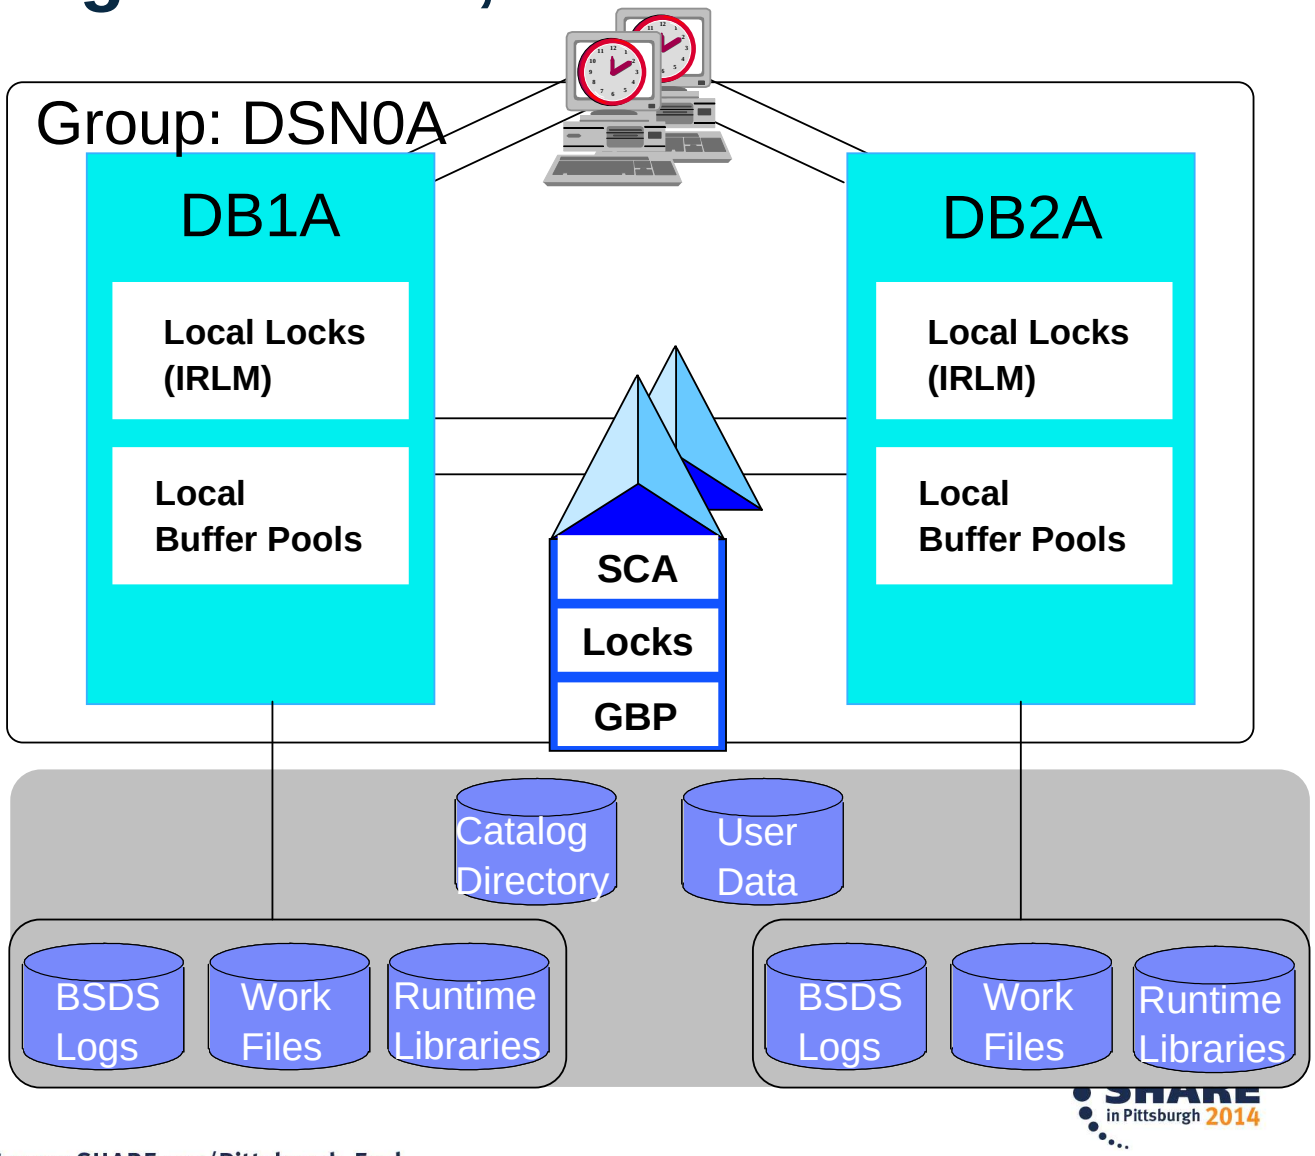

DB2 Data Sharing Overview: Local Attach

- Local attach is to a DB2 member on the same LPAR
  - CICS, IMS, TSO, Batch, CAF, Websphere, RRSAF, JDBC Type2, Local ODBC
  - CICS or IMS trans can be routed to any LPAR
- Can specify member or use group attach name
  - DB1A vs DB0A
- What if two members of same group on 1 LPAR?

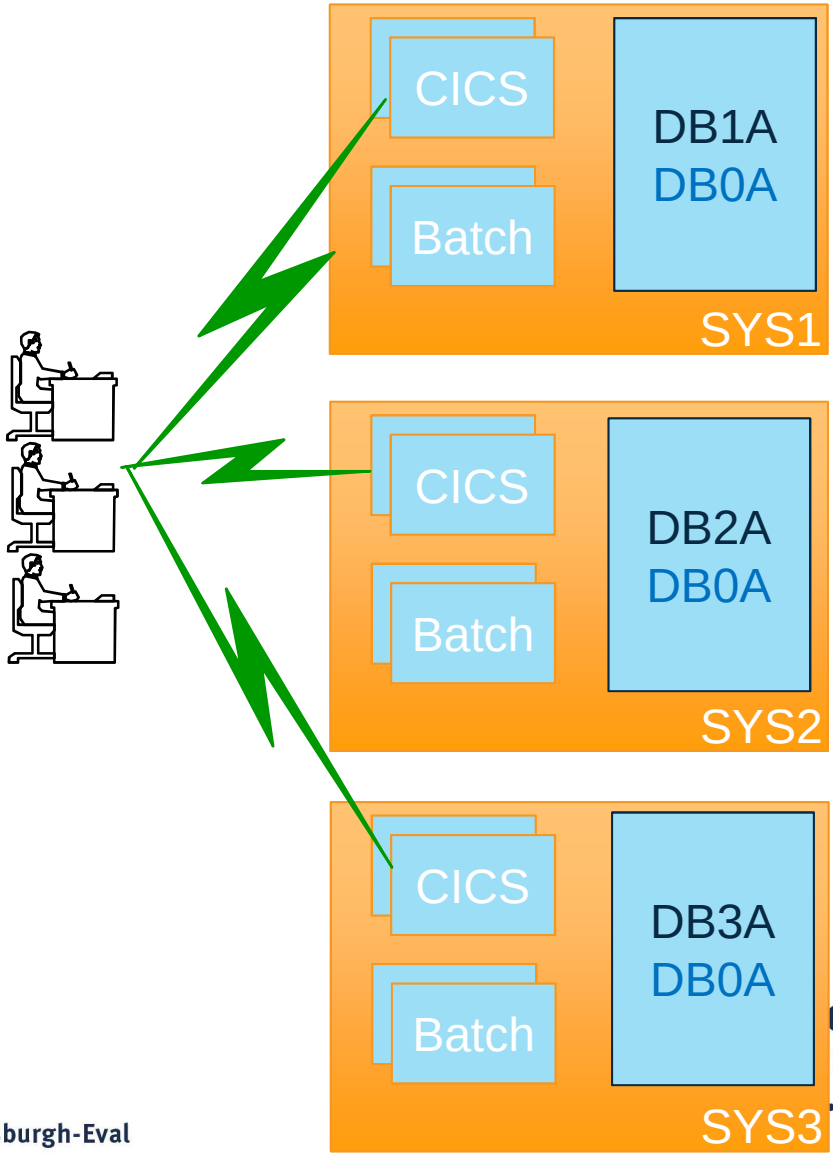

DB2 Data Sharing Overview: Distributed

- Clients or app servers use IP Address or Domain Name, Port and Location Name
  - Clients or app servers can connect to any member via Location Name of the group
  - Member subsetting with Location Alias
    - Static alias defined in BSDS
- Sysplex Workload Balancing can switch between members on transaction boundary
  - Based on weighted server list sent by DB2 members
  - Connect 9.7 FP9

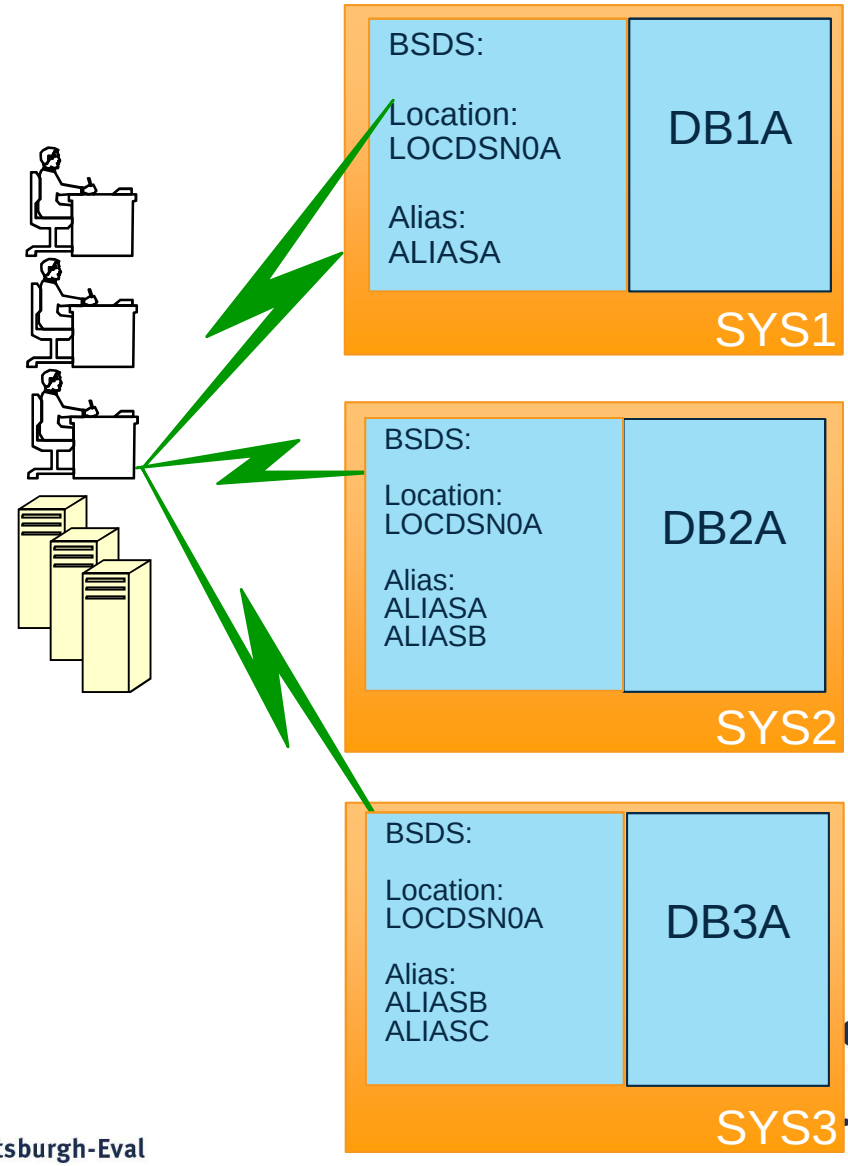

DB2 10 Data Sharing Enhancements

- Dynamic location alias
- Managing DBAT queues
- Subgroup attach name
- Delete data sharing member
- MEMBER CLUSTER for UTS
- RESTART LIGHT handles DDF units of recovery

Action required

- LRSN spin reduction when same data and index are updated
- Avoid excessive cross-invalidations (XIs) on conversion to non-GBP dependent
- Avoid local BP scan on change of GBP dependency
- Auto rebuild CF lock structure on long IRLM waits during restart

DB2 10: Dynamic Location Alias

- Use the MODIFY DDF command with ALIAS option
  - Define and manage up to 40 location aliases dynamically
    - Start, stop, cancel, modify, and delete dynamic location aliases without stopping either DDF or DB2.
    - Dynamic aliases cannot be defined or managed by the DSNJU003 utility
    - DSNJU004 utility does not print any information about dynamic location aliases
    - Use DISPLAY DDF command to find information about these aliases
  - DB2 must be started before you can define dynamic location aliases
    - DDF may or may not be started

DB2 10: Dynamic Location Alias, 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.

DB2 10: Managing DBAT Queues

- What if DBATs start queuing on one member but not on the other?
- Two DSNZPARMs introduced with DB2 10
  - MAXCONQN
    - Maximum connection queue depth
    - ON = value of MAXDBAT
    - Number (less than or equal to MAXDBAT)
  - MAXCONQW
    - Maximum connection queue wait time
    - ON = value of IDTHTOIN (seconds)
    - Number (between 5 and 3600 seconds)
  - Exceeding either value causes connections to be closed
    - Auto Client Reroute (ACR) takes effect
    - Sysplex WLB required in client or app server

DB2 10: Subgroup Attach Name

- Subgroup attach name
  - Can be used by local attaches: CICS, TSO, CAF, RRSAF, JDBC (Type 2), ODBC (local) and DB2 utilities
  - Example in IEFSSNxx member of PARMLIB:

```
DB1A,DSN3INI,'DSN3EPX,-DB1A,S,DB0A'
DB2A,DSN3INI,'DSN3EPX,-DB2A,S,DB0A,SBG1'
DB3A,DSN3INI,'DSN3EPX,-DB3A,S,DB0A,SBG1'
DB4A,DSN3INI,'DSN3EPX,-DB4A,S,DB0A,SBG2'
```

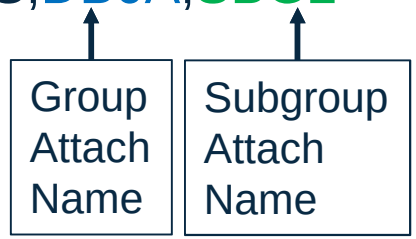

DB2 10: Subgroup Attach Name, example

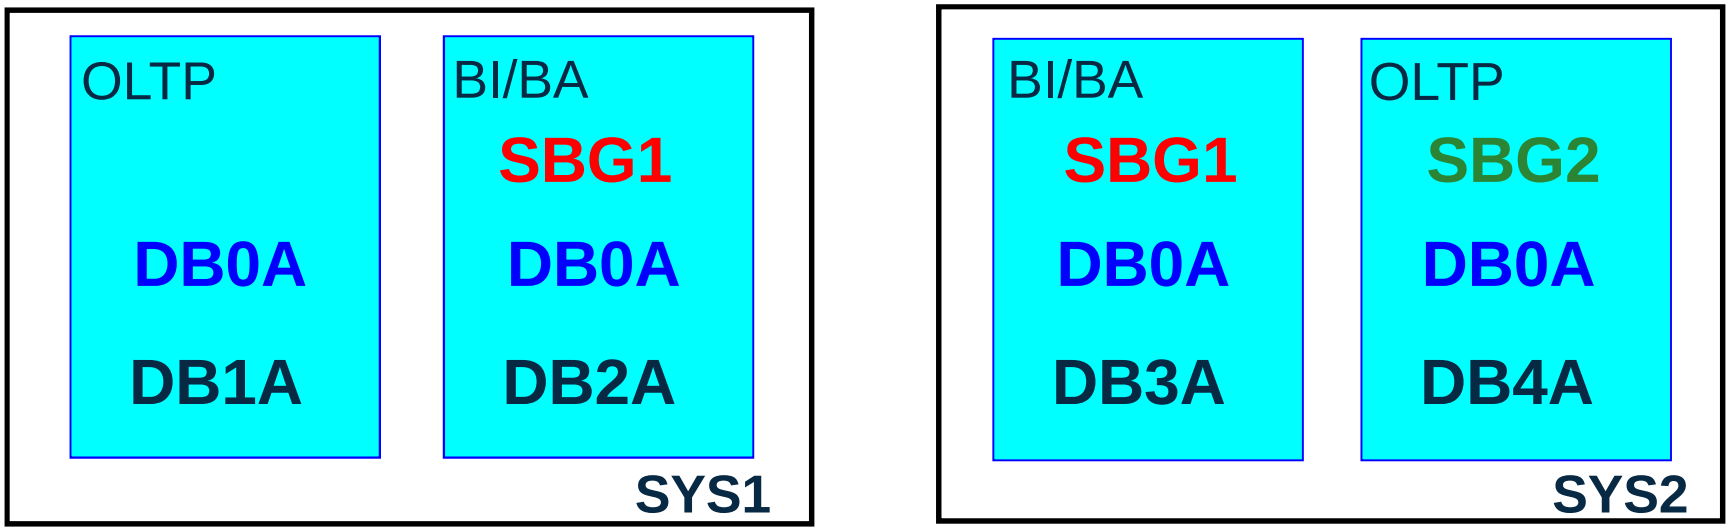

- DB2 members can be specified by
  - Member name: DB1A, DB2A, DB3A, DB4A
  - Group attach name: DB0A
  - Subgroup attach name:
    - SBG1 for DB2A and DB3A
    - SBG2 for DB4A

DB2 10: Member Consolidation

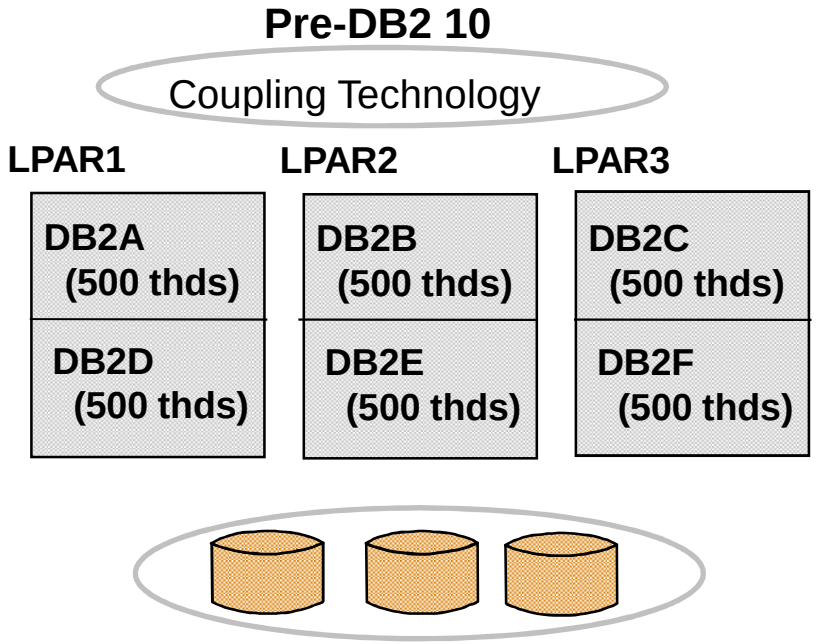

- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s per LPAR

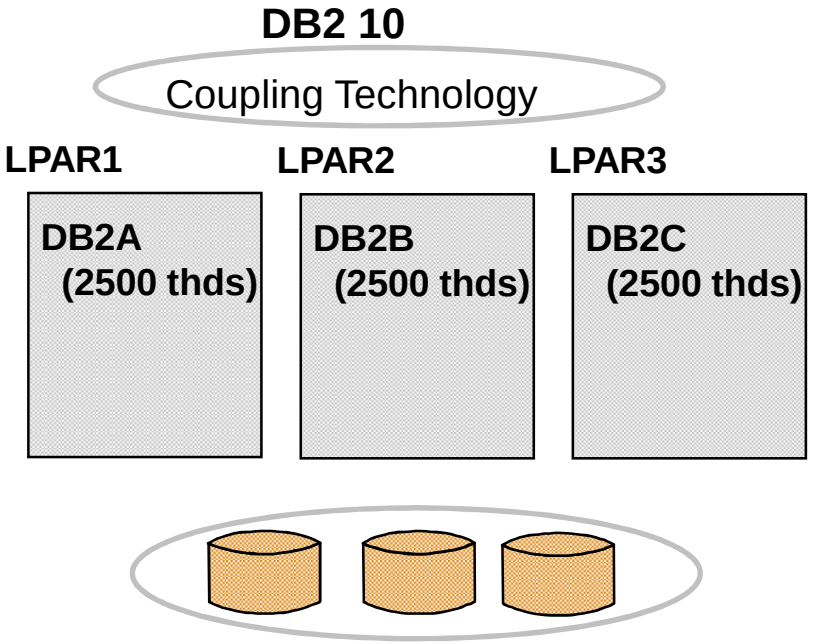

- More threads per DB2 image
- More efficient use of large n-way processors
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for very high availability and scale
- Rule of thumb: save ½% CPU for each member reduced
  - Also save on memory

DB2 10: Member Consolidation

- Delete data sharing member
  - Significant DBM1 virtual storage relief enables data sharing member consolidation (previous slide)
  - DSNJU003, Change Log Inventory, to delete a member, therefore offline function: Group outage required
    - DELMBR DEACTIV
      - *Space in BSDS remains in use*
      - *Member can be restored*
    - DELMBR DESTROY
      - *Reclaim space in BSDS*
      - *Member cannot be restored*
    - RSTMBR – restore deactivated member

DB2 10: MEMBER CLUSTER for UTS

- Assign a space map page, and the pages it manages, to a single member
  - Avoid page P-lock contention for space map pages
- Greater throughput, especially on insert
- Less clustering may impact some workloads
  - REORG may be necessary more frequently

DB2 10: START LIGHT Enhancement

- Restart LIGHT handles DDF units of recovery
  - DB2 in LIGHT mode stays up until DDF UoRs are resolved
  - Case where DB2 DDF is the participant, not the coordinator
  - For XA transactions, another member must be active
    - Other member must be available through DVIPA
    - Returns member DVIPA and resync port for indoubt work (XIDs) owned by ‘light’ member

DB2 11 Data Sharing Enhancements

- Improved CASTOUT processing
- RESTART LIGHT(CASTOUT) option
- LRSN spin elimination
- GBP write-around protocol
- Locking enhancements

Action required

- Improved DELETE_NAME performance
- GBP write performance
- Index availability and performance
- Automatic LPL recovery at end of restart

DB2 11: Improved Castout Processing

- New ALTER GROUPBUFFERPOOL syntax
  - CLASST=(0, integer)
    - Integer = # of changed pages in class castout queue
    - Similar to VDWQT

- Faster castout – read/write overlap
- Reduced NOTIFY message size

DB2 11: Restart LIGHT(CASTOUT)

- Failed DB2 member holds retained locks
  - Including page set p-locks
  - Fast restart to release retained locks vital to availability
  - Not all systems have resources to restart ‘full’ DB2
- START LIGHT(YES) and LIGHT(NOINDOUBTS)
  - Did not address castout, therefore could not clear page set p-locks in IX or SIX mode
  - Utilities can be blocked by these retained page set p-locks
- START LIGHT(CASTOUT)
  - When castout processing complete, all page set p-locks released

DB2 11: LRSN Spin Elimination

- DB2 11 NFM Large RBA / LRSN EXTENDED format (10 bytes)
  - Uses more granular STCK (Storeclock) precision value
  - Therefore LRSN will be unique for consecutive updates, inserts, deletes

DB2 11: GBP Write Around Protocol

- GBPs can be filled rapidly during heavy batch type processing
  - GBP full conditions, transaction response time degradation
- Avoid write of ‘new’ pages to GBP
  - Conditionally enables GBP write around (pages will be written directly to DASD)
    - Enabled at 50% GBPOOL OR 20% CLASSQ pages
    - Disabled at 40% GBPOOL OR 10% CLASSQ pages
    - DSNB777I will provide statistics (-DIS GBPOOL MDETAIL)
    - If page already in GBP, must be written to GBP
  - CFCC 17 and 18; z196 and later
  - z/OS 1.13 with OA37550 or z/OS 2.1
- Should reduce GBP performance problems

DB2 11: Locking Enhancements

- IRLM 2.3
  - Increased maximum CF lock table entries
    - IRLM now supports 2 G entries

  - Improved performance handling lock waiters
  - Conditional propagation of child U locks to CF
    - Improved performance for SELECT FOR UPDATE
  - Throttle batched unlock requests
    - Avoid overwhelming processor
    - IRLM 2.2 and 2.3; PM60449
  - More efficient deadlock cycle processing
    - Reduced CPU time for deadlock processing
    - Reduced contention on LOCK1 structure access

DB2 11: Delete Name

- CF DELETE NAME enhancement
  - Uses CFCC17 new option to avoid XI
    - z114 or z196
  - CFCC 18 for zEC12
  - z/OS APAR OA38419
  - Significant performance benefit for remote CF
- PM67544 retrofit to DB2 9 and DB2 10

DB2 11: GBP Write Performance

- DB2 11 allocates fixed storage for GBP batch write
- Eliminate page fix / page free instruction
- Reduced path length for COMMIT

DB2 11: Index Availability and Performance

- Avoid placing indexes in RBDP during Group Restart
  - Removes necessity to rebuild indexes
  - NFM only
- Reduction of log force write during Index tree structure modification
  - Reduced to 1
  - Throughput improvement in Insert and Delete by reducing log force write per index modification event
    - Index page split or index page delete
  - Elapsed time reduction and minor CPU reduction

DB2 11: Automatic LPL Recovery

- Pages added to Logical Page List (LPL) if they cannot be written to or read from GBP
  - Pages on LPL cannot be accessed by applications
  - START DB(xx) SPACE(yy) for every object in LPL
    - Time consuming and potentially error prone
- DB2 11 adds automatic LPL recovery at end of normal restart and restart light
  - Auto LPL recovery not triggered for some circumstance
    - For example:
      - *Indoubt or postponed abort recovery*
      - *Member started in ACCESS MAINT*
      - *Others*
  - Retrofit to DB2 10: PM78128

mrader@us.ibm.com

Resources

- Redbooks
  - *DB2 10 for z/OS Technical Overview*, SG24-7892
  - *DB2 10 for z/OS Performance Topics*, SG24-7942
  - *DB2 11 for z/OS Technical Overview*, SG24-8180
  - *DB2 11 for z/OS Performance Topics*, SG24-8222

  - For introductory information on DB2 data sharing
    - *DB2 for z/OS: Data Sharing in a Nutshell*, SG24-7322


Thank
You