DB2 for z/OS Data Sharing: Configurations and Common Issues

Session 17008

Mark Rader
IBM – DB2 for z/OS
March 6, 2015

#SHAREorg

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

SHARE in Seattle 2015

Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. 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. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. 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 ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

Agenda

- **DB2 data sharing overview**
  - What is it? Why do our clients implement it?
- **DB2 data sharing concepts**
  - Parallel Sysplex ®
  - Coupling Facilities (CFs)
  - CF Structures
- **DB2 data sharing configurations**
- **Common issues**
  - Performance
  - Availability
  - Dynamic workload balancing
- **Resources and FAQs**

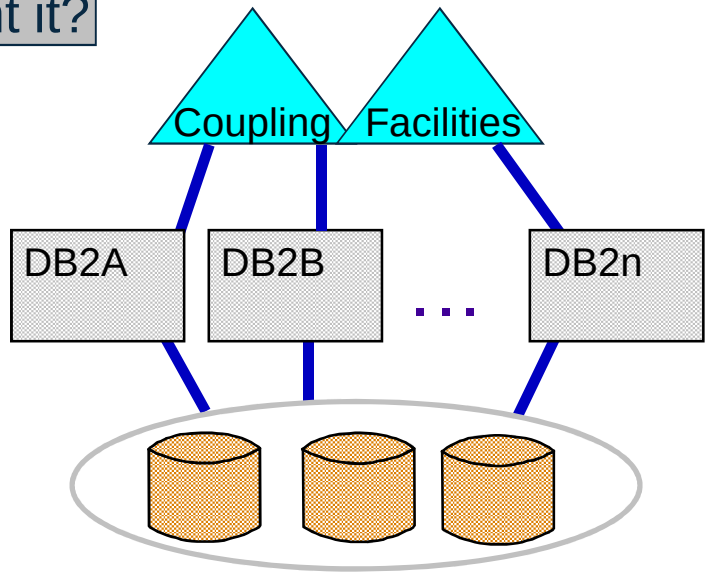

DB2 Data Sharing - Overview

- **DB2 data sharing and Parallel Sysplex (PSX) provide the on-demand infrastructure for:**
  - Increased availability
  - Non-disruptive scalability
  - Dynamic workload balancing
- **Definitions:**
  - Technical:
    - DB2 data sharing allows applications running on more than one DB2 subsystem to read and write to the same set of DB2 data concurrently.
  - Business:
    - DB2 data sharing allows our clients to provide the highest level of scalability, performance and continuous availability to enterprise applications that use DB2 data.

Why IBM Clients Implement to DB2 Data Sharing

- **Most common drivers:**
  - Higher availability requirements
    - Protection against planned and unplanned outages
  - Easier growth accommodation
    - Need for scalable, non-disruptive growth to handle business, market, or regulatory changes
  - Dynamic workload balancing
    - Effective utilization of available MIPS for mixed workloads
    - Handle unpredictable workload spikes
  - Capacity: outgrowing size of a single system
    - Avoid splitting the databases
  - Avoid ‘Cold Start’ of DB2 when approaching the end of the log RBA
  - System consolidation for easier systems management
- **Application investment protection**
  - SQL interface is unchanged for data sharing
  - Excellent scaling
  - Applications do not need to become "cluster aware" as processing nodes are added

Agenda

- **DB2 data sharing overview**
  - What is it? Why do our clients implement it?
- **DB2 data sharing concepts**
  - Parallel Sysplex ®
  - Coupling Facilities (CFs)
  - CF Structures
- **DB2 data sharing configurations**
- **Common issues**
  - Performance
  - Availability
  - Dynamic workload balancing
- **Resources and FAQs**

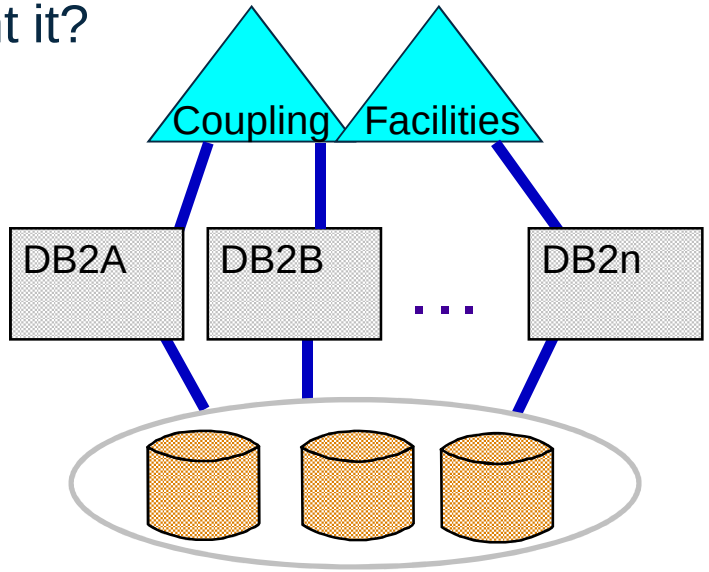

Alternative Parallel DBMS Architectures

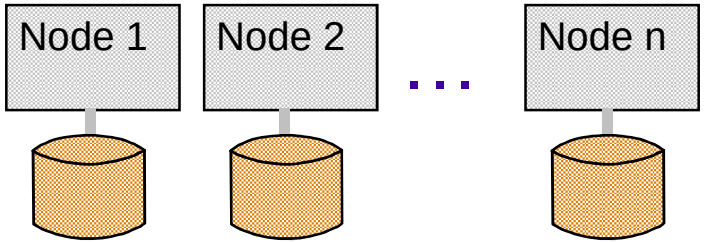

Shared Nothing (SN)

- Database is partitioned
- No disks are shared amongst the nodes
- Distributed commit is necessary
- Data repartioning necessary as nodes are added
- Susceptible to skewed access patterns

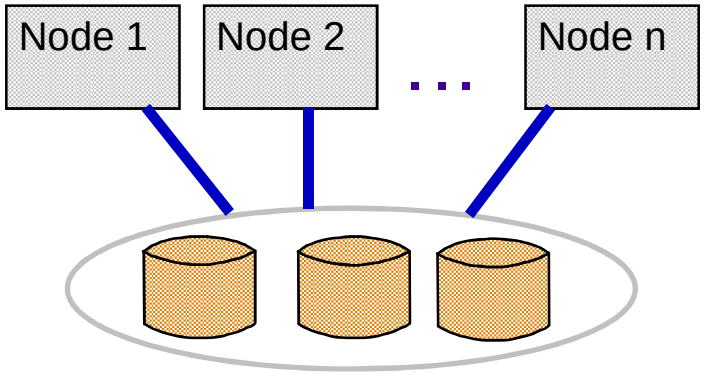

Shared Disks (SDi)

- No database partition necessary
  - But partitioning can give better performance
- Strong fail-over characteristics
- Dynamic load balancing
- Inter-node concurrency and coherency control mechanisms are needed
  - Messaging overhead limits scalability

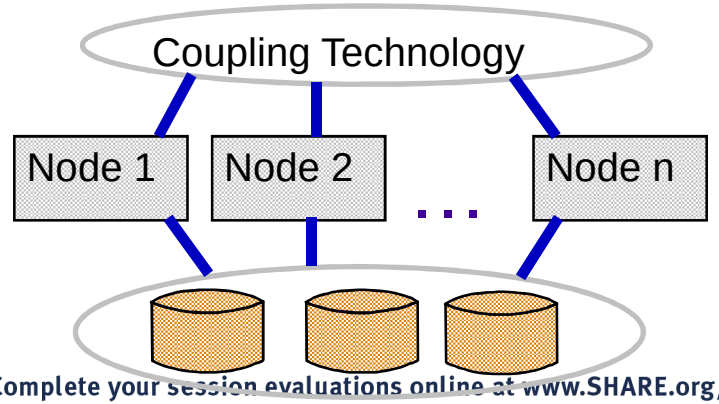

Shared Data (SDa)

- Adaptation of SDi
- Coupling facility is used as hardware assist for efficient concurrency and coherency control
- Strong fail-over and load balancing as with SDi
- Flexible growth
- Messaging overhead minimized, excellent scalability

Parallel Sysplex (PSX)

- PSX Components:
  - Sysplex Timer Protocol (STP)
    - ‘Sysplex Timers’
  - Coupling Facility (CF) LPARs
    - High-speed shared memory
    - CF Control Code (CFCC)
    - Structures (Lock, Cache, List)
  - CF Links [solid blue lines]
  - CF Resource Management (CFRM) Policy
  - Cross-System Extended Services (XES), part of z/OS
- Scalable Capacity
- Flexible Configuration
- Workload Balancing
- 7x24 Availability
- Single System Image

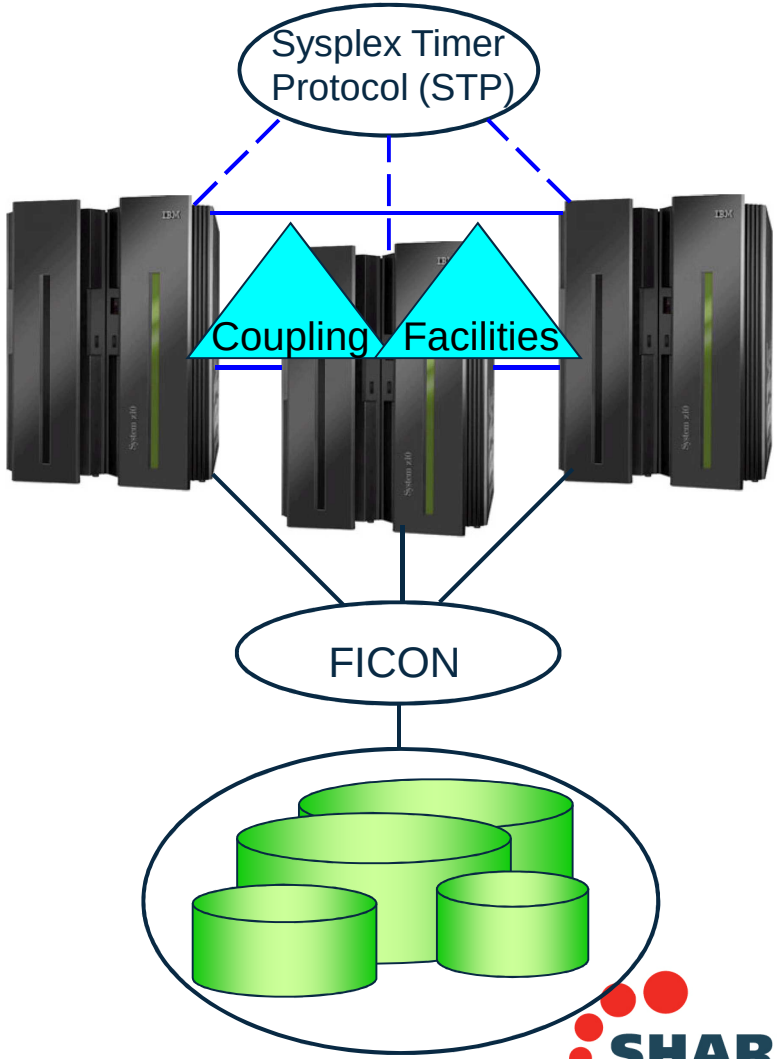

DB2 Data Sharing Concepts

- **A DB2 Data Sharing Group consists of:**
  - 2 or more DB2 members with a single Catalog/Directory
    - 1-way data sharing initially, for migration
  - Active and archive logs for each member
    - Use log record sequence number (LRSN) primarily, instead of RBA
  - DB2 and User data on shared disk
  - A DB2 Data Sharing Group is a single location for distributed access
- **For DB2 use, Coupling Facilities (CFs) contain:**
  - 1 Lock structure per data sharing group *
  - 1 Shared Communications Area (SCA) per group *
  - Multiple Group Buffer Pools (GBPs) per group
    - 1 GBP per Buffer Pool containing shared data
    - GBP0, GBP8K0, GBP16K0, GBP32K required
  - * = Required structures

DB2 Data Sharing – Implementation Steps

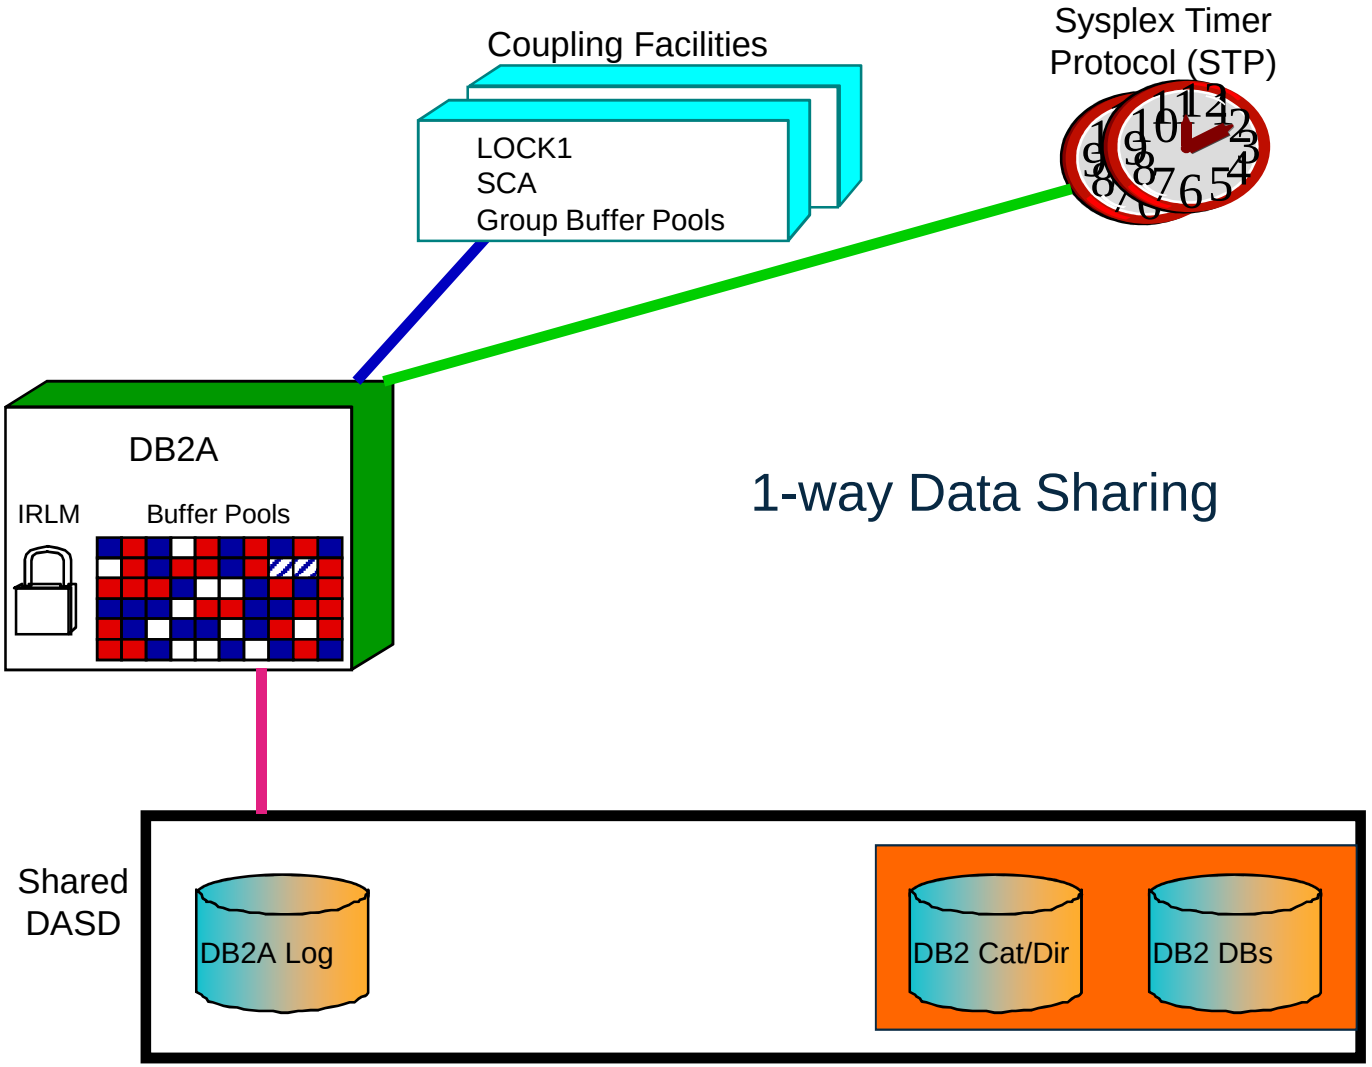

DB2 Data Sharing – Implementation Steps

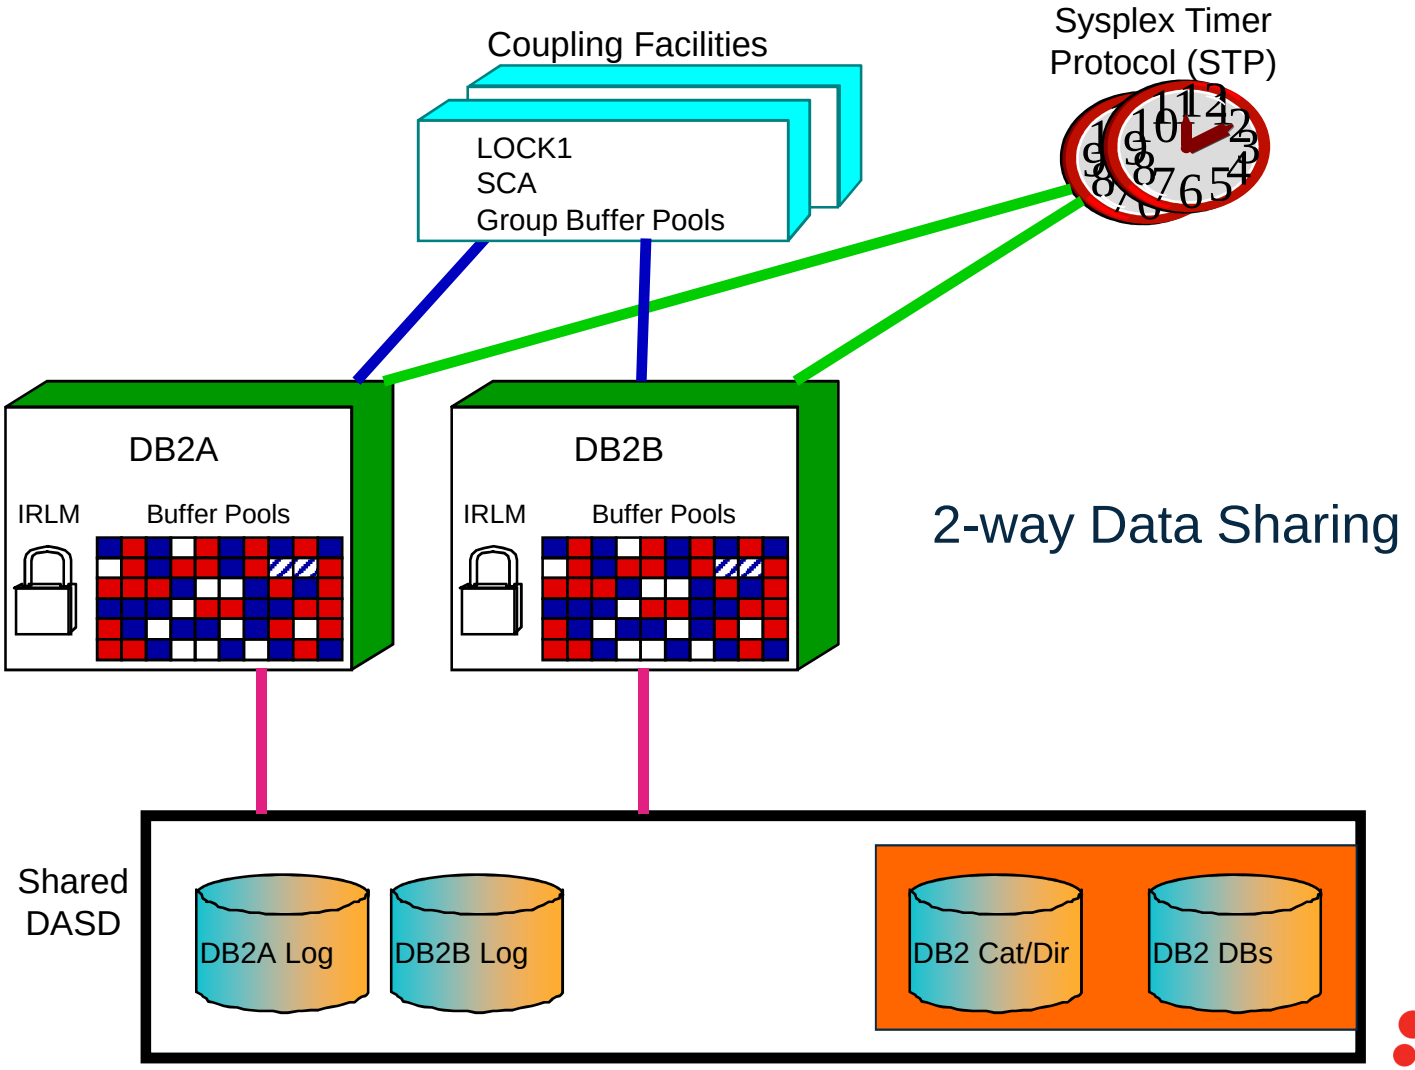

DB2 N-way Data Sharing

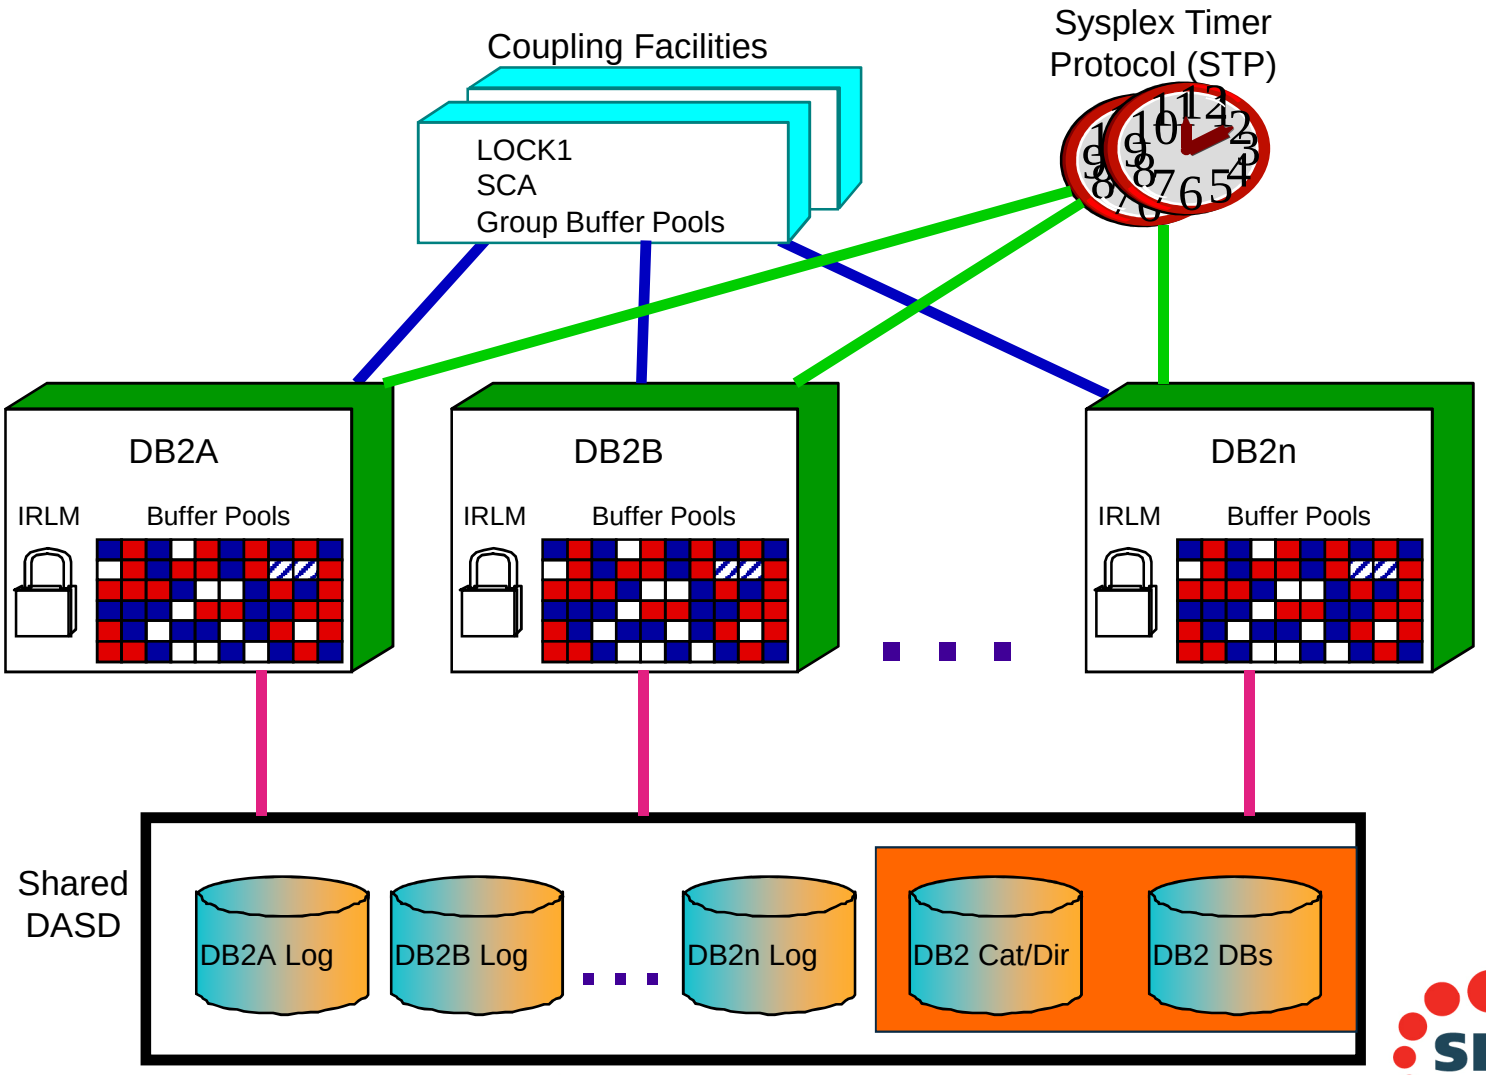

DB2 CF Structures

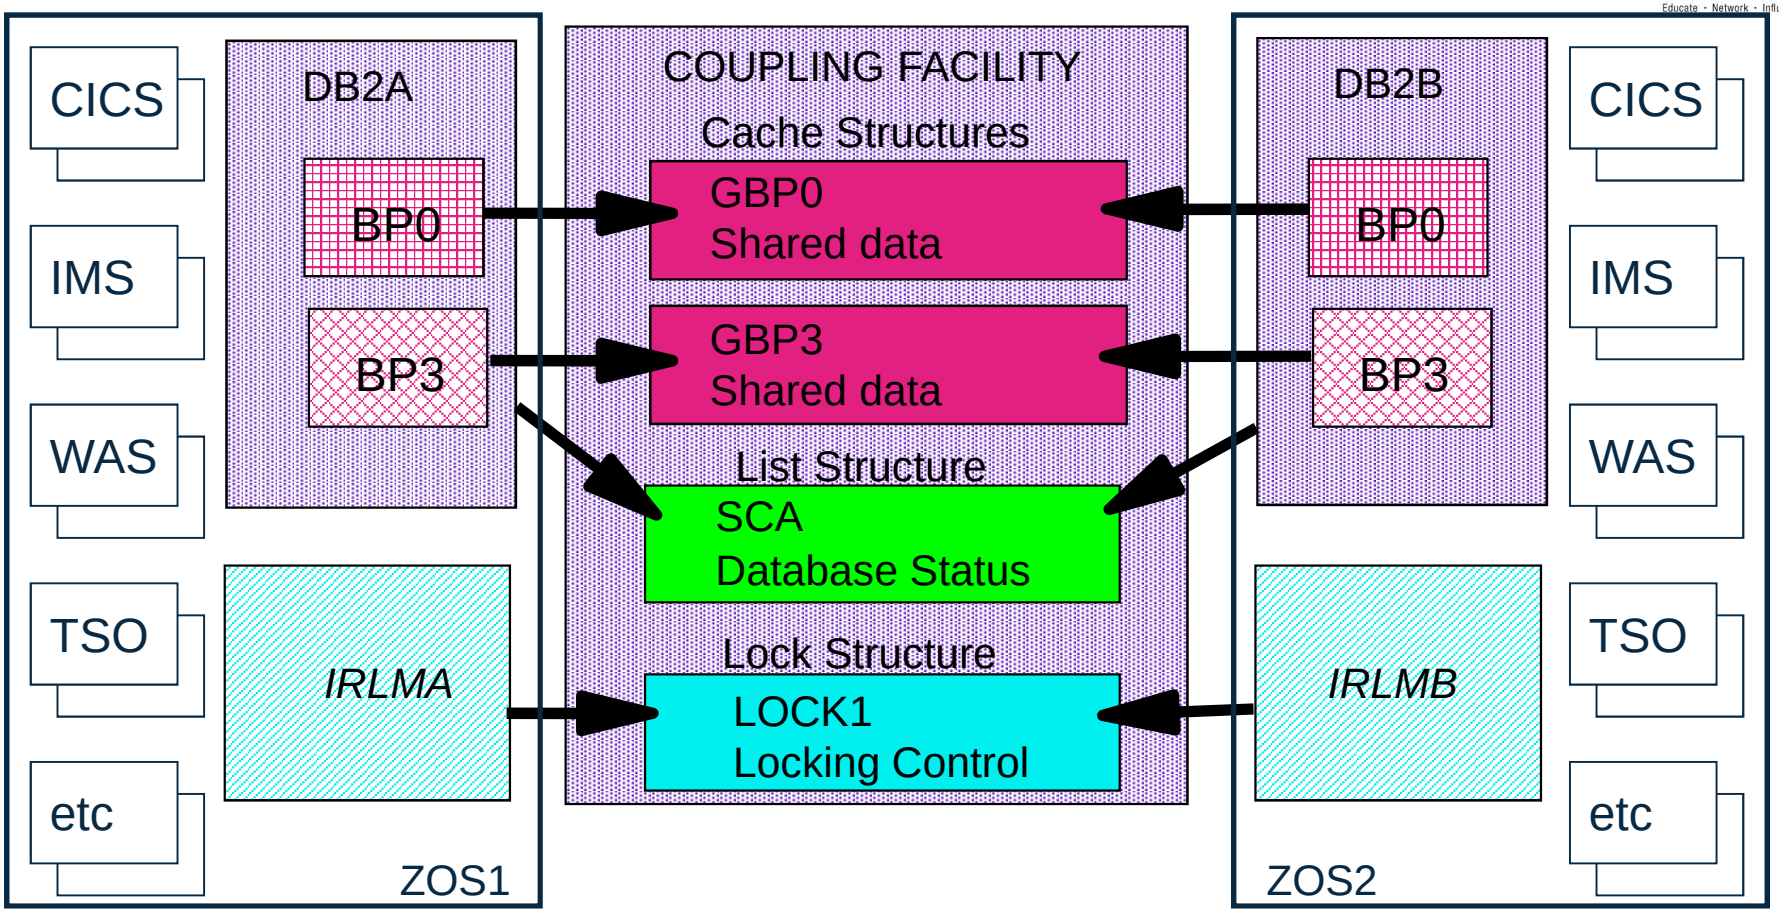

- *ssnm*MSTR allocates SCA (*grpname*_SCA)
- *ssnm*DBM1 allocates GBPs (*grpname*_GBP*n*)
- *ssnm*IRLM allocates LOCK1 (*grpname*_LOCK1

Shared Communications Area (SCA)

- **Used by DB2 to maintain group-wide status information**
  - Recovery Pending
  - Copy Pending
  - Write Error Ranges
  - Logical Page List
  - GRECP Status
  - BSDSs of all DB2s in data sharing group
  - System checkpoint intervals
  - Database Exception Table (DBET)
- **SCA is generally not a performance concern**

Lock Structure (LOCK1)

- **Used by IRLM to manage global locking**

- **Holds L-locks and P-locks**
  - L-locks to track concurrency
  - P-locks to track coherency
    - P-locks new for data sharing
- **Consists of a lock (or hash) table and a modify lock list**
  - Lock table controls access to resources
    - Minimal information to track readers (‘S’ locks) and writers (‘X’ locks)
    - Based on hashing algorithm
  - Modify lock list contains detailed information for update-type locks
    - Entries become retained locks in case of an IRLM or DB2 or LPAR or CEC failure
    - Retained locks must be released via restart of the DB2 member that was subject to the failure

Lock Structure (LOCK1)

- **Simplified view**

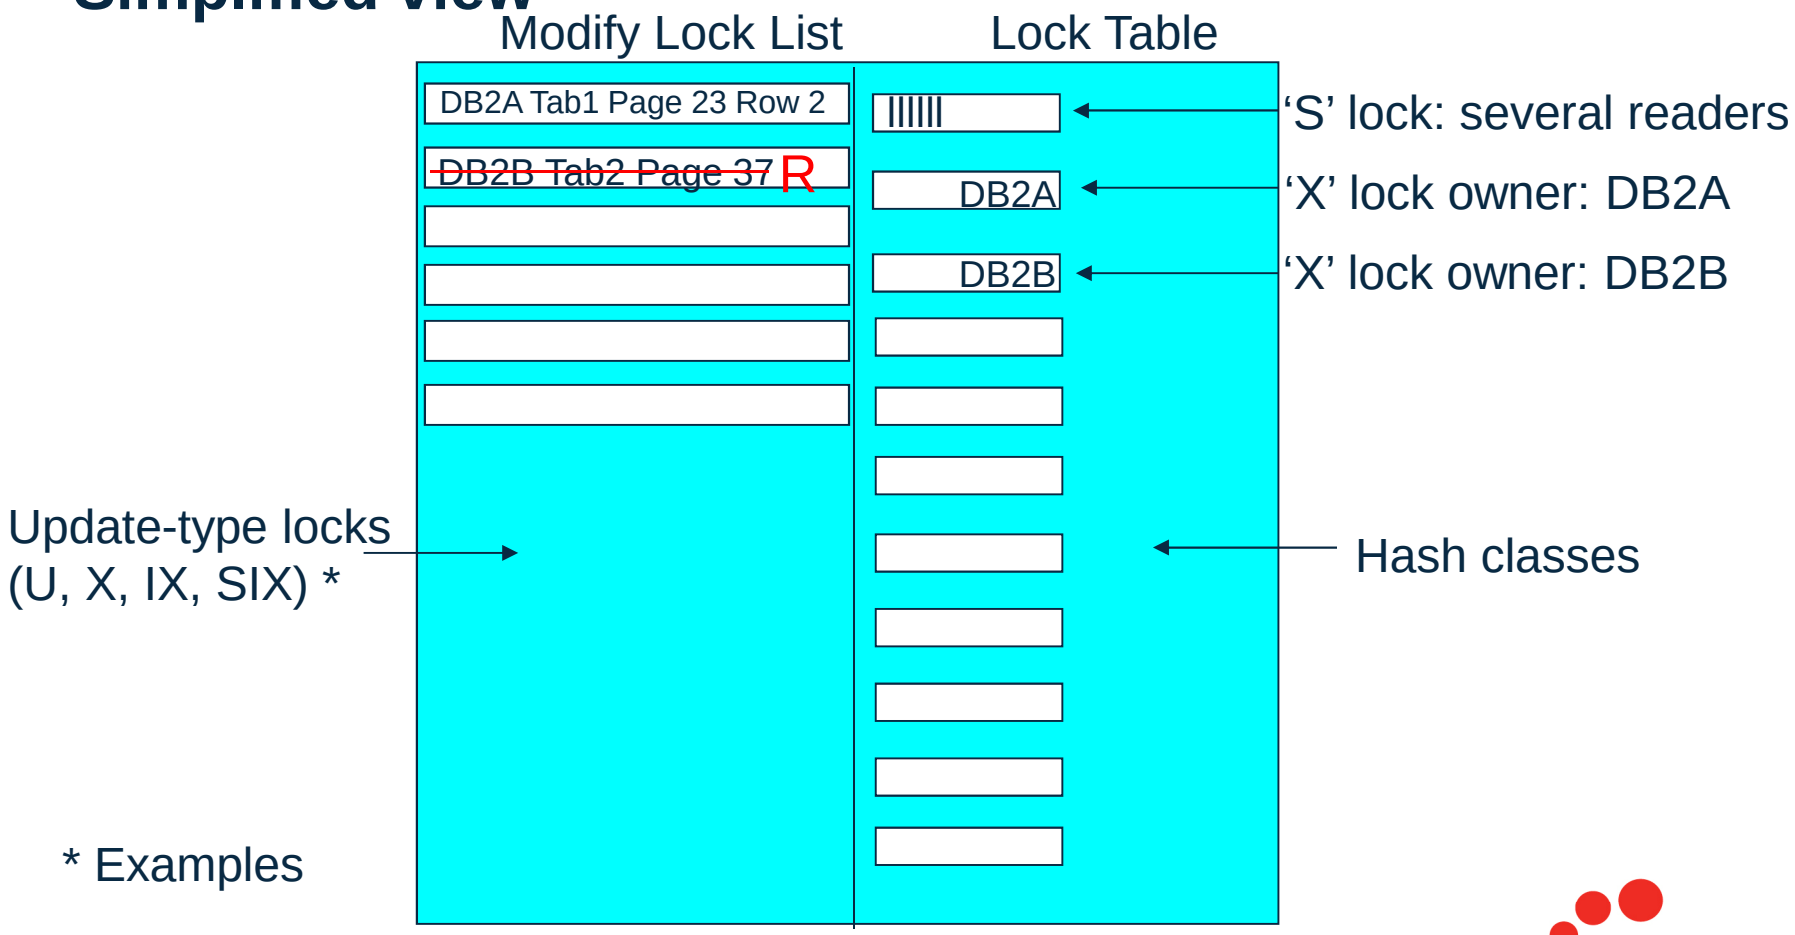

Group Buffer Pools (GBPs)

- **DB2 uses GBPs to**
  - Manage buffer coherency
  - Cache changed pages
    - Optionally cache read-only pages
- **GBP consists of directory entries and data elements**
  - Directory entries manage coherency by tracking interest in a data or index page by any DB2 member in the data sharing group
  - Data elements are the cached pages that a DB2 member changed
  - Default 5:1 ratio – AutoAlter can change dynamically

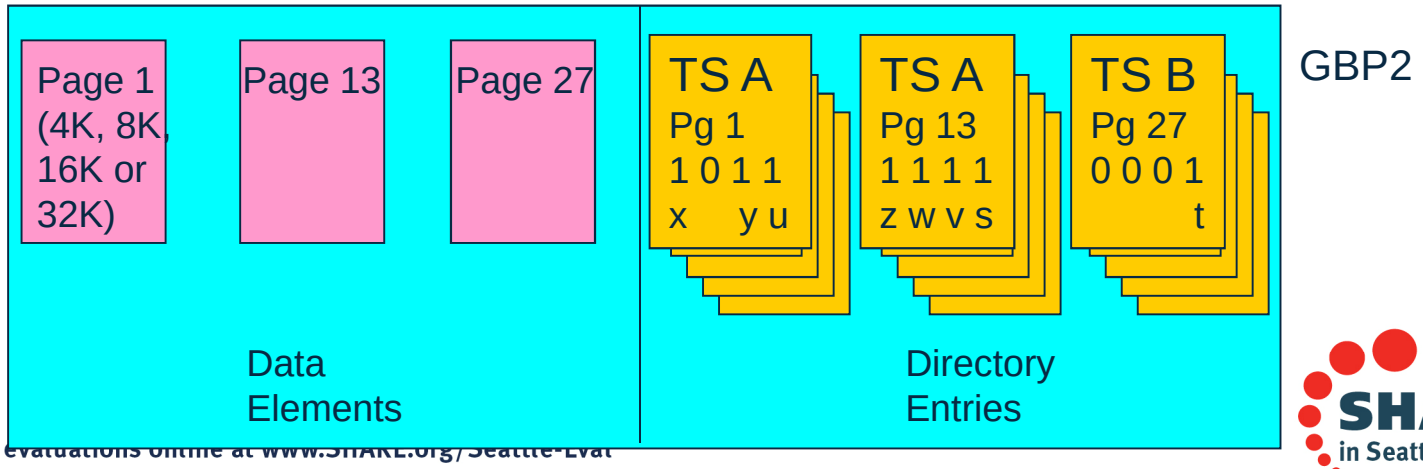

Agenda

- **DB2 data sharing overview**
  - What is it? Why do our clients implement it?
- **DB2 data sharing concepts**
  - Parallel Sysplex ®
  - Coupling Facilities (CFs)
  - CF Structures
- **DB2 data sharing configurations**
- **Common Issues**
  - Performance
  - Availability
  - Dynamic workload balancing
- **Resources and FAQs**

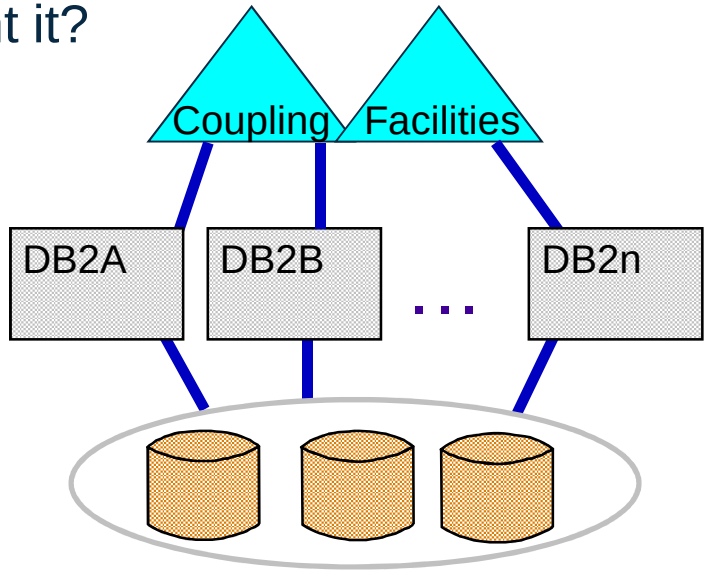

Parallel Sysplex Configurations

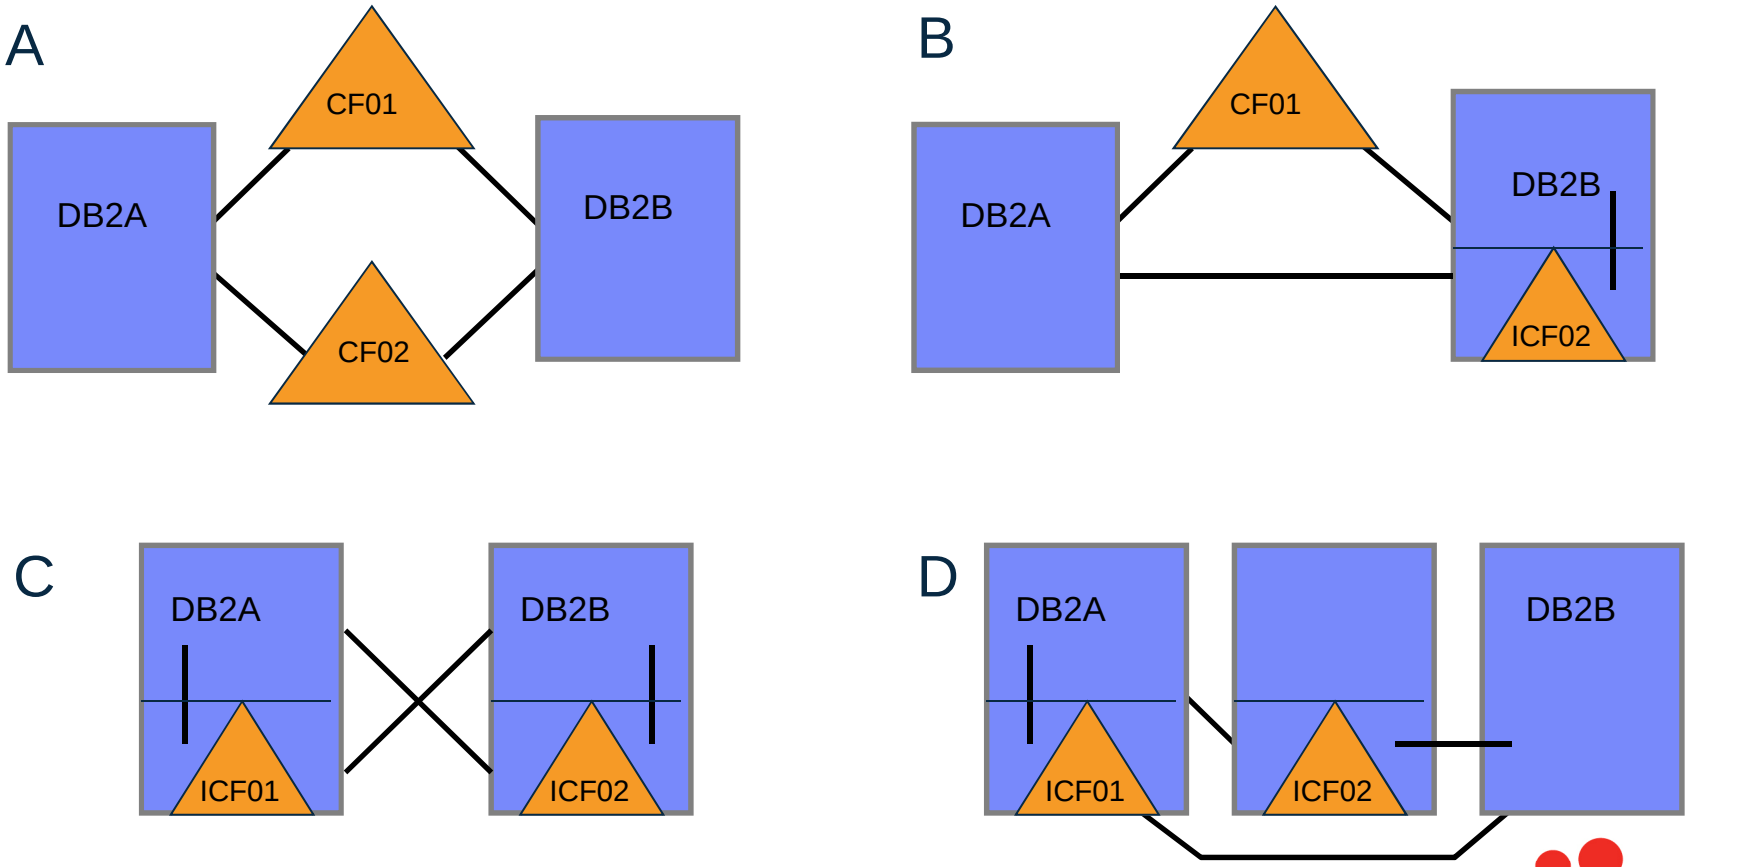

Parallel Sysplex Configurations

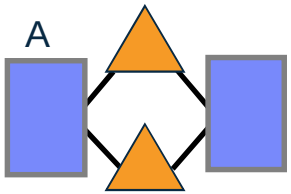

- A: “Traditional” configuration
  - LOCK1 and SCA in one CF
  - Duplexed GBPs spread across both CFs
    - Primary and secondary GBPs balanced based on load

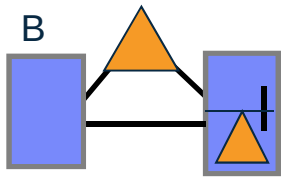

- B: One Integrated CF (ICF), one external CF
  - LOCK1 and SCA in external CF
  - Duplexed GBPs spread across both CFs
    - Primary GBP in ICF has advantages for ‘local’ DB2

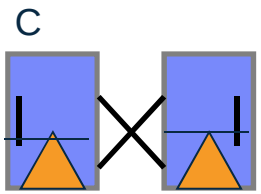

- C: Two ICF configuration
  - Lock1 and SCA duplexed; allocated in both CFs
    - “System-managed” duplexing
    - Performance implication for LOCK1 requests
  - Duplexed GBPs spread across both CFs
    - Primary GBP in ICF has advantages for ‘local’ DB2

- D: Special case of ‘B’

Sysplex in a Box

- **Some customers use this configuration to avoid planned outages**
  - Still at risk of unplanned outages, especially hardware

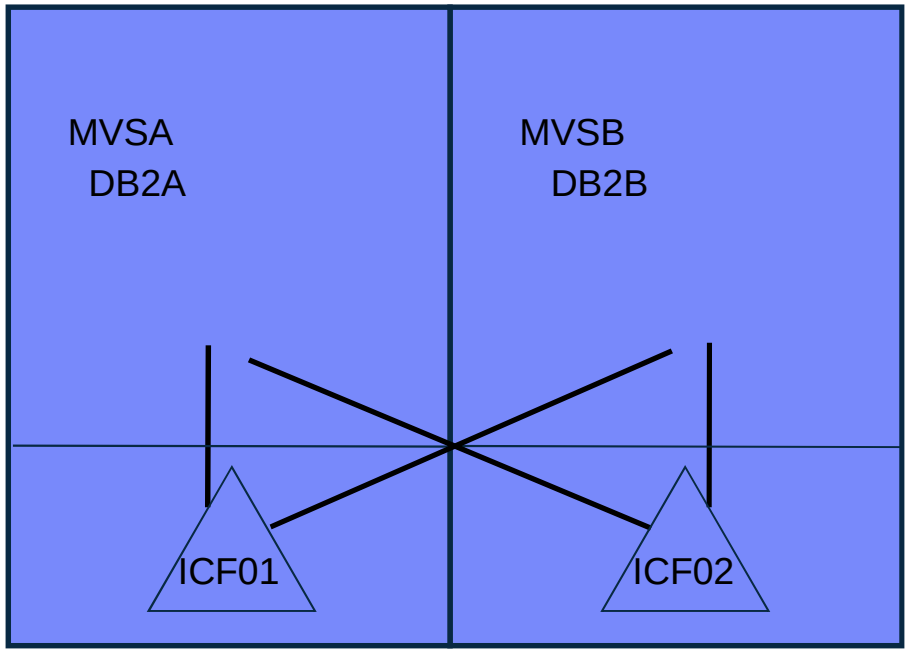

Agenda

- **DB2 data sharing overview**
  - What is it? Why do our clients implement it?
- **DB2 data sharing concepts**
  - Parallel Sysplex
  - Coupling Facilities (CFs)
  - CF Structures
- **DB2 data sharing configurations**
- **Common Issues**
  - Performance
  - Availability
  - Dynamic workload balancing
- **Resources and FAQs**

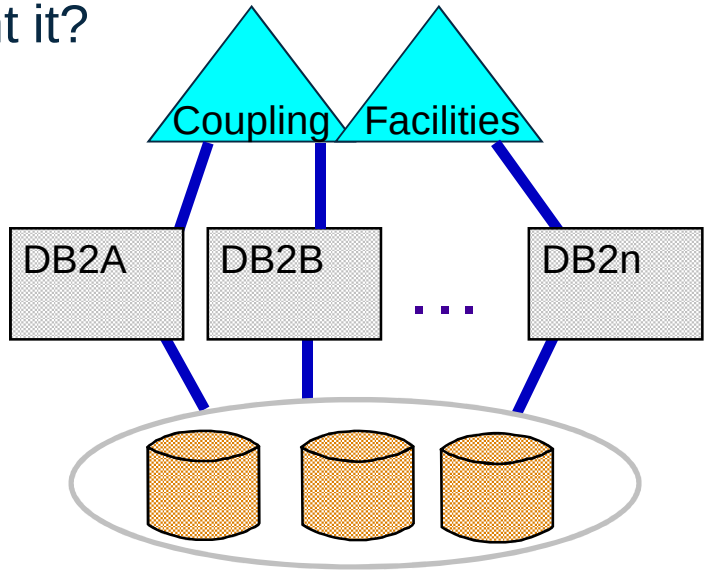

Design Goals - Data Sharing Performance

- **Little or no performance impact if data not actually shared, i.e. if no inter-DB2 R/W interest**
  - Dynamic recognition of sharing
- **Minimal and acceptable CPU overhead if inter-DB2 R/W interest exists**
  - Overhead will vary based on individual workload characteristics
- **Near-linear scalability when adding 3rd through nth nodes**
  - Incur majority of data sharing overhead going from 1-way to 2-way
  - Beyond 2-way, additional overhead is minimal

Critical Performance Factors

- **Two factors to preserving data integrity in a data sharing environment**
  - Inter-system ***concurrency*** control - global locking
    - Multiple readers OR
    - One writer
  - Inter-system buffer ***coherency*** control – managing changed data
    - When one system changes data rows that also reside in other system(s)
- **"Data sharing overhead" is attributable to the extra CPU cost needed to manage these two factors**
  - Thousands to tens of thousands of messages per second
    - Extreme example: 166,114 synchronous lock requests per second
      - average in one RMF interval
- **Most CF messages for DB2 / IRLM are synchronous**
  - Incur host CPU overhead for duration of round trip to CF

Data Sharing: Locking

- Global locking using Parallel Sysplex® coupling technology
  - ***Inter-system concurrency control***

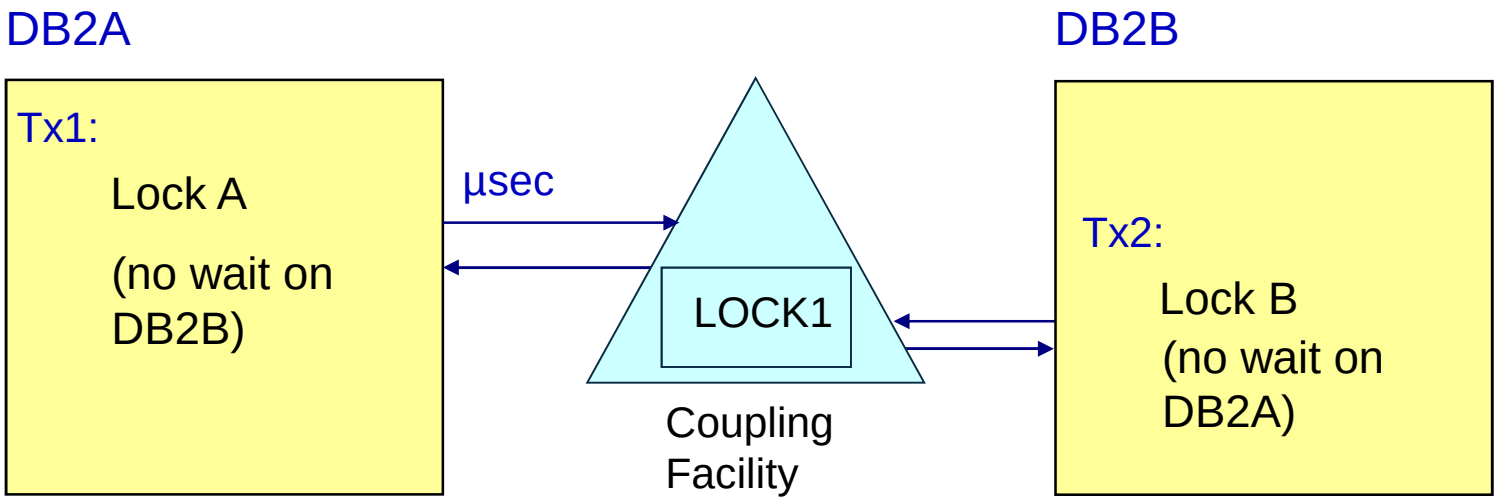

  - Cost of obtaining lock does not increase when adding 3rd through nth members
  - This example assumes no contention. Contention results in additional messaging

Data Sharing: Managing changed data

- ***Inter-system buffer coherency control***
  - Example: DB2A has write interest in the table space, and page P1 is in DB2A's buffer pool

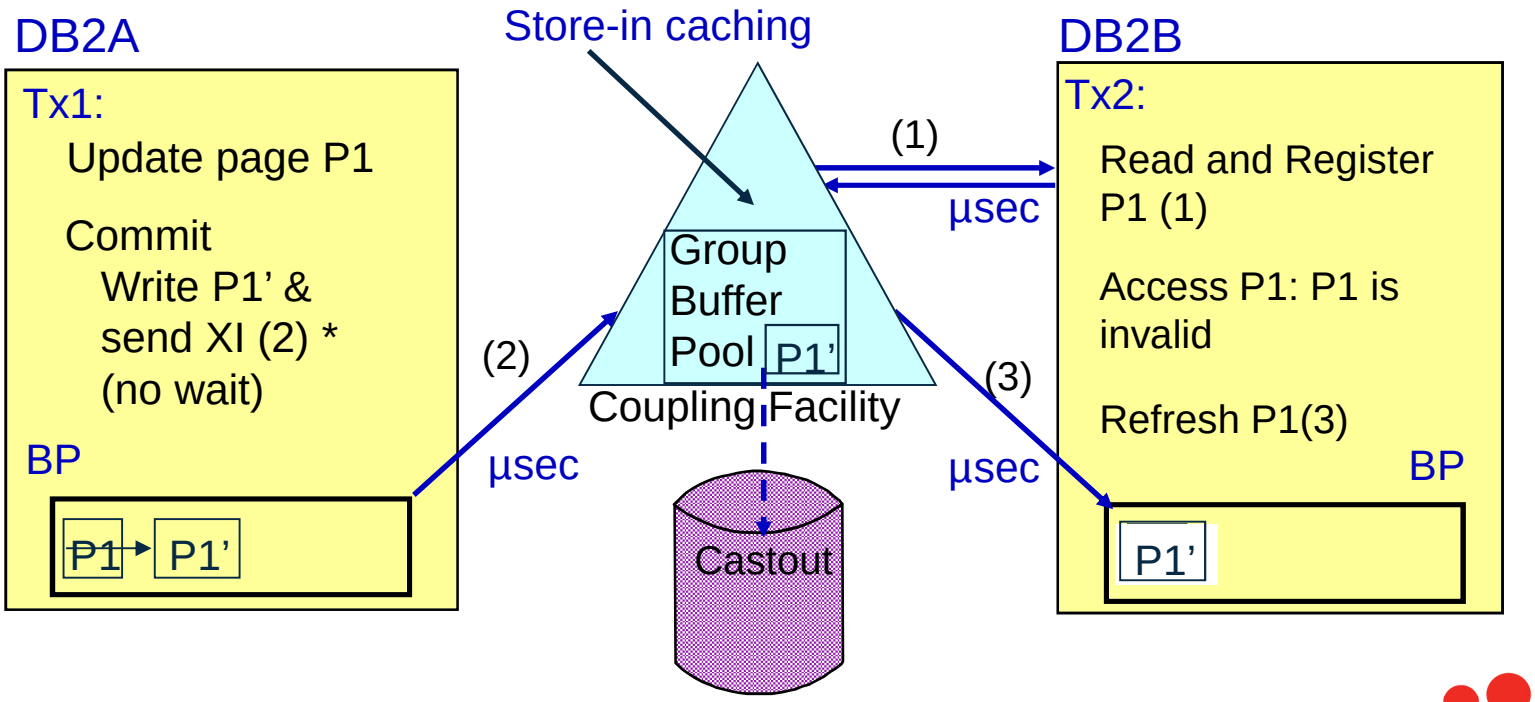

- * Cross-invalidate (XI) to other member without interrupt

Data Sharing Performance Summary

- **CPU cost of data sharing varies based on:**
  - CF access intensity for locking and caching. This varies based on:
    - Percentage of CPU time in DB2
    - Degree of read/write sharing
    - Number of locks obtained
    - Access rate to shared data
    - Insert/delete intensity
    - Release of DB2
  - Hardware configuration
    - Slow CF increases CPU overhead on host LPAR
  - Lock contention rates
- **Data sharing cost varies from one workload to another**
  - 'Typical' 2-way data sharing overhead about 10%
    - Important: this is relative to the workload, not an address space
  - Individual jobs/transactions may have higher overhead
  - < 0.5% added cost per member past 2-way
- **Balanced configuration important for best performance**
  - CF technology = Host LPAR CEC technology
  - Two CF links per Host/CF, two dedicated CF engines per CF LPAR

Data Sharing Performance in Production

- **Host CPU effect with primary application involved in data sharing**
  - 10% is a typical average
  - Scalability and performance for real life customer workloads

Industry	Trx Mgr / DB Mgr	z/OS Images	CF access per Mi	% of used capacity
Pharmacy	CICS/DB2	3	8	10%
Insurance	CICS/IMS+DB2	9	9	10%
Banking	IMS/IMS+DB2	4	8	11%
Transportation	CICS/DB2	3	6	8%
Banking	IMS/IMS+DB2	2	7	9%
Retail	CICS/DB2+IMS	3	4	5%
Shipping	CICS/DB2+IMS	2	8	9%

Note: “Mi” stands for ‘million instructions’

CASTOUT processing

- **CASTOUT is the process that 'hardens' changed pages in the GBP to disk.**
- **CASTOUT will occur when:**
  - CLASST exceeded
    - Analogous to VDWQT
  - GBPOOLT exceeded
    - Analogous to DWQT
  - GBP checkpoint
  - No more inter-DB2 interest in the page set
  - GBP being rebuilt, but alternate GBP is not big enough to contain cached pages
- **Keep CASTOUT steady**
  - Use BP VDWQT for steady writes to GBP
  - Use GBP CLASST for steady castout

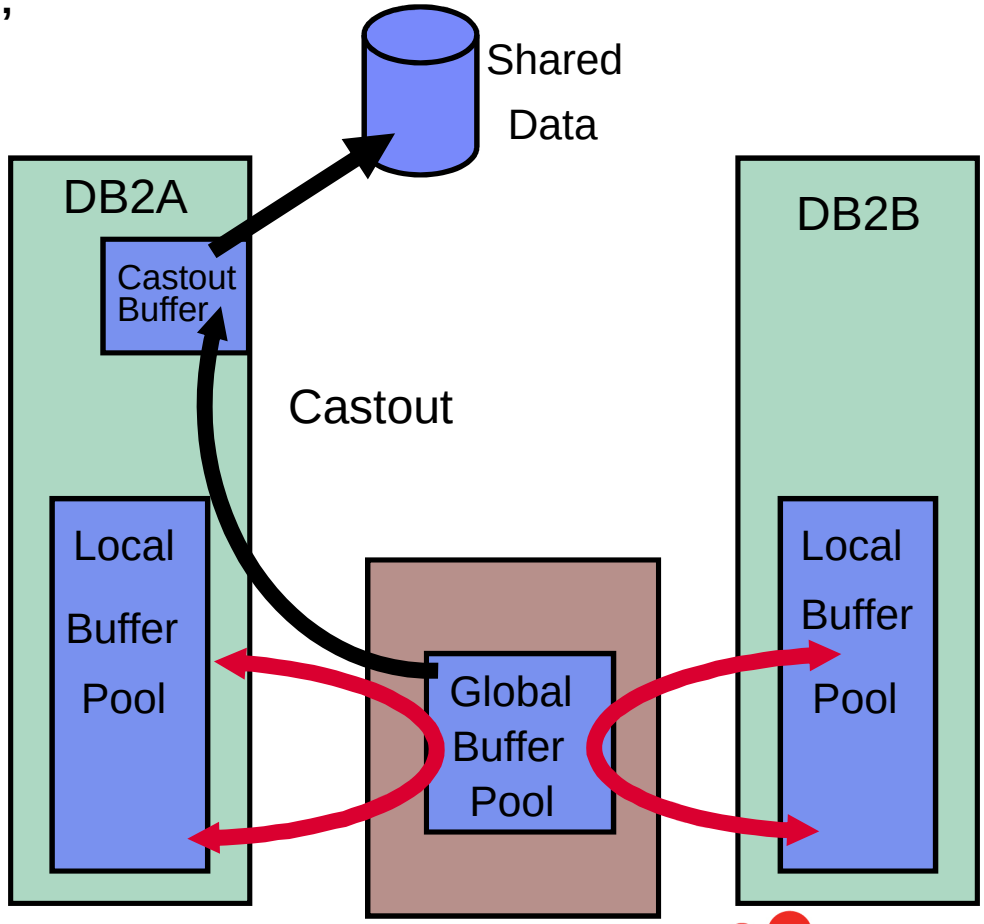

DB2 Data Sharing Performance: Common Issues

- **Under-configured CF**
  - High CF CP busy can increase host CPU overhead
    - Elongated sync service time
    - Prefer 2 dedicated CF engines
  - CF Structures too small – could lead to:
    - LOCK1 – increased false contention: overhead
    - LOCK1 – insufficient modify lock list entries: IRLM slow downs
    - GBPs – directory entry reclaims and XIs due to reclaims: overhead, extra I/O
    - GBPs – write failures due to lack of storage: slow downs, pages on LPL
- **Unbalanced Host-CF technology**
  - E.g. Host LPAR is on zEC12, but CF LPAR is on z10 or z196
    - High host CPU for synchronous CF requests
    - Sync to async conversion: increased lock service time, increased elapsed time, decreased throughput
- **Heavy update plus slow castout can result in structure full warnings**
  - Use VDWQT to write to GBP steadily, rather than many pages at commit
- **High rate of data set physical close and open**
  - Tablespace or Indexspace CLOSE (YES|NO)
  - DSNZPARM: PCLOSEN and PCLOSET
- **Long running units of recovery (UoRs) reduce lock avoidance**

Design for High Availability

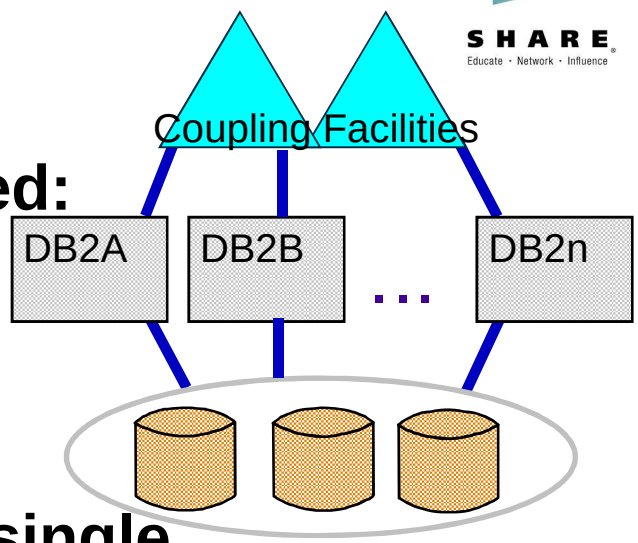

- **Most single points of failure eliminated:**
  - DB2 subsystem or z/OS system
  - CPC (or CEC)
  - I/O path
- **Goal: Continuous availability across planned or unplanned outage of any single hardware or software element**
- **Strategy:**
  - Remove all causes for planned outages
  - Build on legacy of robust, fault tolerant MVS components
  - On a failure:
    - Isolate failure to lowest granularity possible
    - Automate recovery and recover fast

DB2 Member Outage – Planned

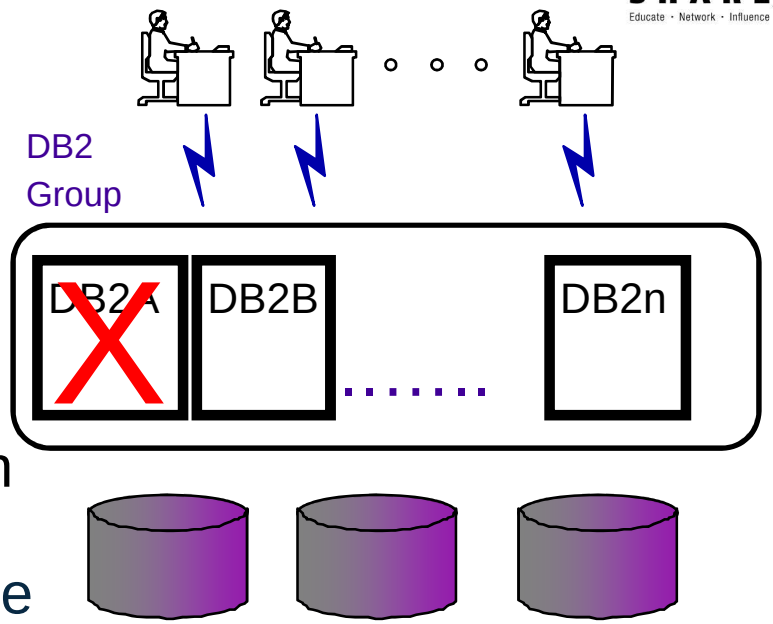

- **"Rolling" maintenance**
  - One DB2 stopped at a time
  - DB2 data continuously available via the N-1 members
  - Other members temporarily pick up work of the member that is down
  - Batch work can be offloaded to another member with more available capacity to reduce the batch window
- **Applies to hardware and operating system changes, too**
  - Rolling IPLs
- **KEY TO SUCCESS: Applications must be able to run on more than one DB2 member!**

DB2 Member Outage – Planned

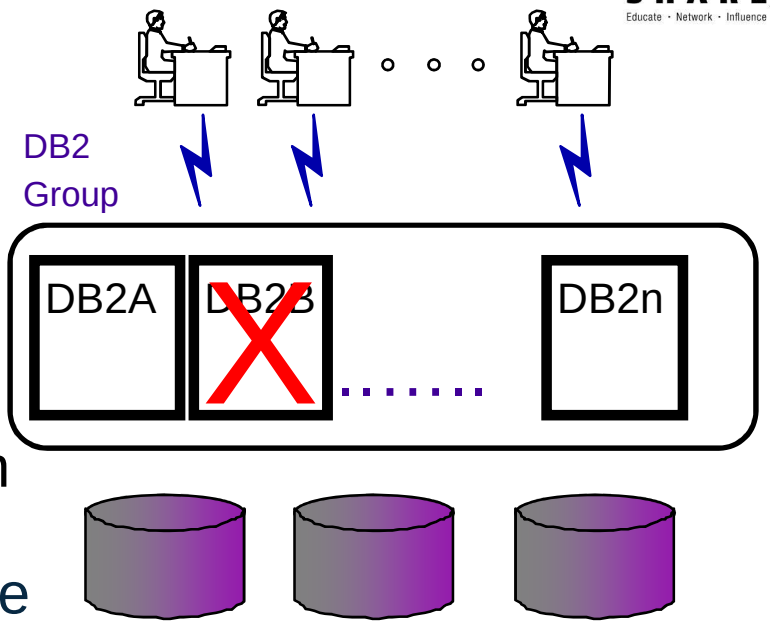

- **"Rolling" maintenance**
  - One DB2 stopped at a time
  - DB2 data continuously available via the N-1 members
  - Other members temporarily pick up work of the member that is down
  - Batch work can be offloaded to another member with more available capacity to reduce the batch window
- **Applies to hardware and operating system changes, too**
  - Rolling IPLs
- **KEY TO SUCCESS: Applications must be able to run on more than one DB2 member!**

DB2 Member Outage – Planned

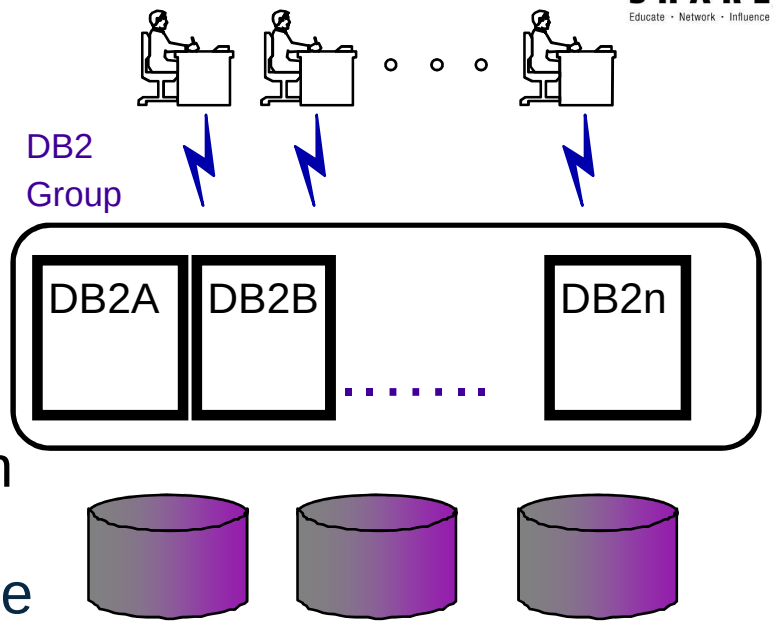

- **"Rolling" maintenance**
  - One DB2 stopped at a time
  - DB2 data continuously available via the N-1 members
  - Other members temporarily pick up work of the member that is down
  - Batch work can be offloaded to another member with more available capacity to reduce the batch window
- **Applies to hardware and operating system changes, too**
  - Rolling IPLs
- **KEY TO SUCCESS: Applications must be able to run on more than one DB2 member!**

DB2 Release Migration

- **DB2 Data Sharing Group can be available – and applications executing - across release migrations**
- **N/N+1 release levels can coexist**
- **Coexistence of mixed releases in a data sharing group can add complexity**
  - Consult DB2 manual, *Data Sharing: Planning and Administration* for details
- **Process:**
  - Apply SPE to each DB2 member
  - Restart each DB2 member
  - Put in new release
  - Restart each DB2 member

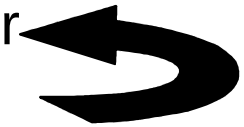

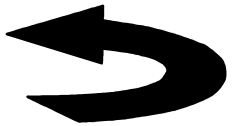

- **Catalog migration done once per data sharing group**

DB2 Member Outage – Unplanned

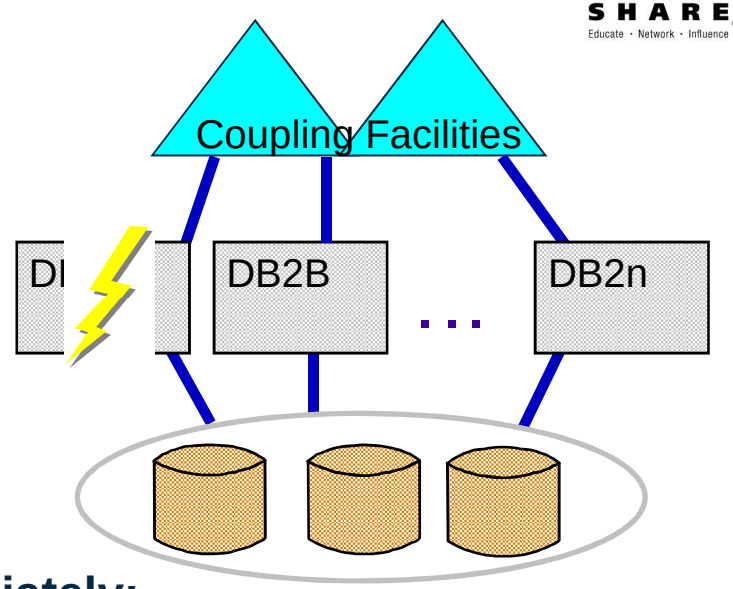

- **The other "surviving" members remain up and running**
- **The architecture allows all members to access all portions of the data bases**
- **Work can be dynamically routed away from the failed DB2 member – assuming applications can run on >1 DB2**
- **Failed member must be restarted immediately:**
  - The failed member holds "retained locks" to protect inconsistent data from being accessed by other members
  - If LPAR failure, automating Cross System Restart is best practice
- **z/OS Automatic Restart Manager (ARM) can automatically restart failed DB2 members**
- **Restart 'Light' minimizes storage impact of restarting DB2 after an LPAR failure**

Coupling Facility Outages

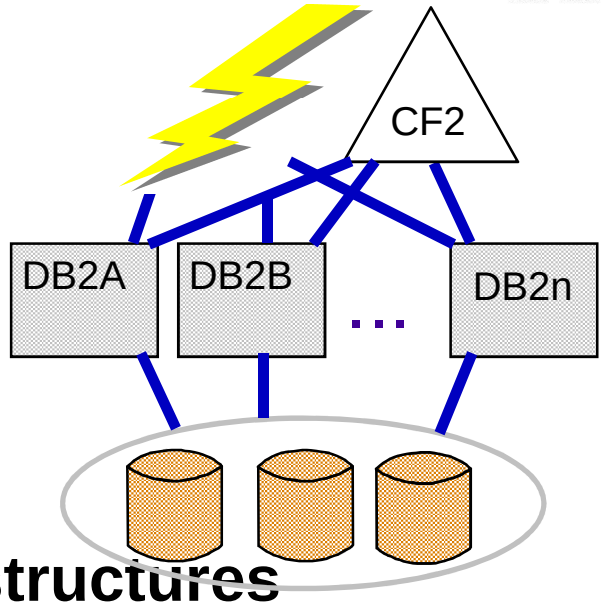

- **Planned outages: use the z/OS operator command to "rebuild"**
  - REBUILD moves the structures to another CF LPAR
  - No outage to data sharing group
  - REALLOCATE to ‘move’ back
- **Unplanned outages: the system automatically recovers the lost structures**
  - Lock & SCA are dynamically rebuilt into alternate CF
    - Spare CF capacity required to house the structures
    - ‘White space’ part of CF capacity planning
  - GBPs must be duplexed for high availability
  - DB2 for z/OS allows duplexing of Lock and SCA but duplexing not necessary for high availability
    - Unless in a 2-ICF configuration
      - Refer to slides 18-19 and Configuration C

Duplexing – 2 Kinds

- **“User-managed” duplexing**
  - Applies to GBPs
  - “User” = DB2; DB2 is responsible for managing two structures in different CF LPARs
- **“System-managed” duplexing**
  - Applies to LOCK1 and SCA
  - XES (z/OS) is responsible for managing two structures in different CF LPARS
    - DB2 is not aware of second structure
    - Everything happens on the primary, then on the secondary
  - Recommended if CFs run on Integrated Coupling Facilities (ICFs) and co-reside on CEC with DB2 members
    - Avoid ‘double failure scenario’
    - But up to 4x cost for synchronous lock requests

GBP Duplexing

- **Allocate secondary GBP on alternate CF**
- **Write changed pages to both primary and secondary**
- **If loss of connectivity or loss of structure -**
  - Switch to secondary (seconds)
  - No rebuild required; changed pages already in GBP
  - Cross-invalidate buffers and gradually repopulate directory entries
- **No application outage unless both primary and secondary GBPs are lost**
- **No excuse not to implement GBP duplexing and AutoAlter**

CF Structure Duplexing

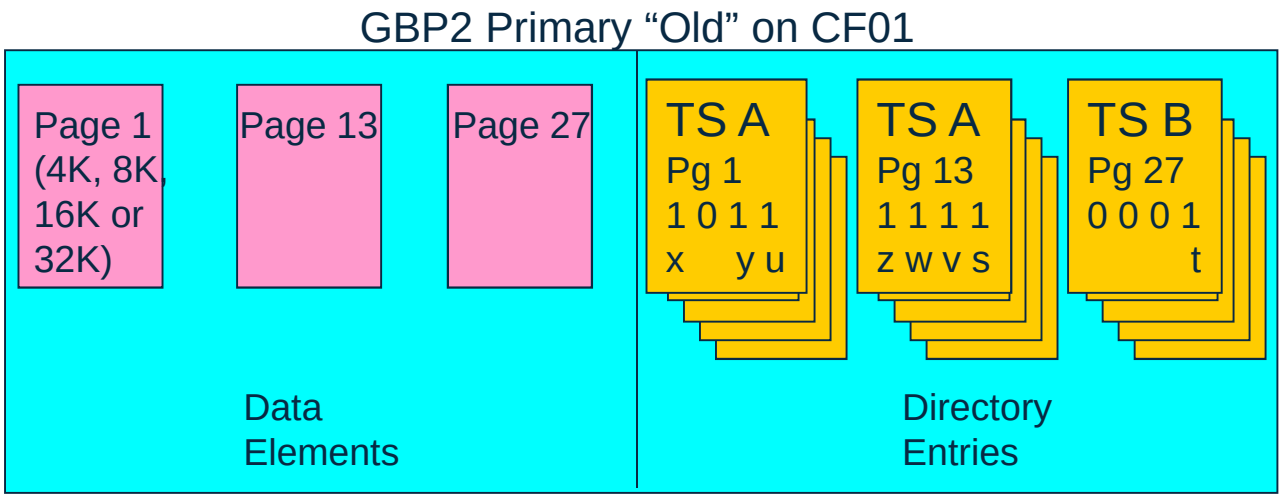

-Cache changed pages
-Register interest
-Castout
-Search if local buffer miss

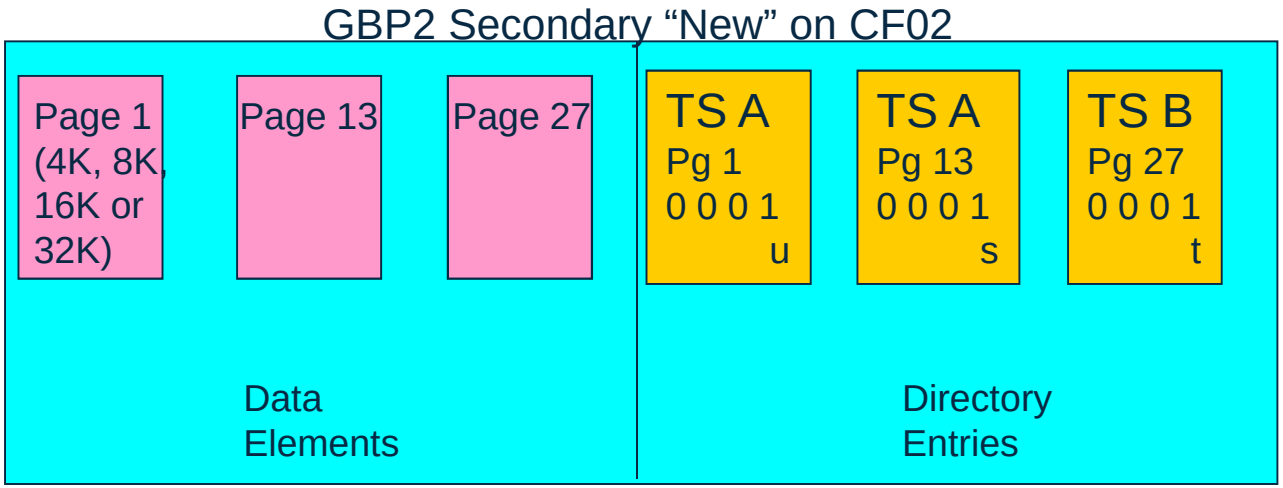

-Cache changed pages

V5 GBP Duplexing Recovery

IRWW workload
Measured at 08/22/98

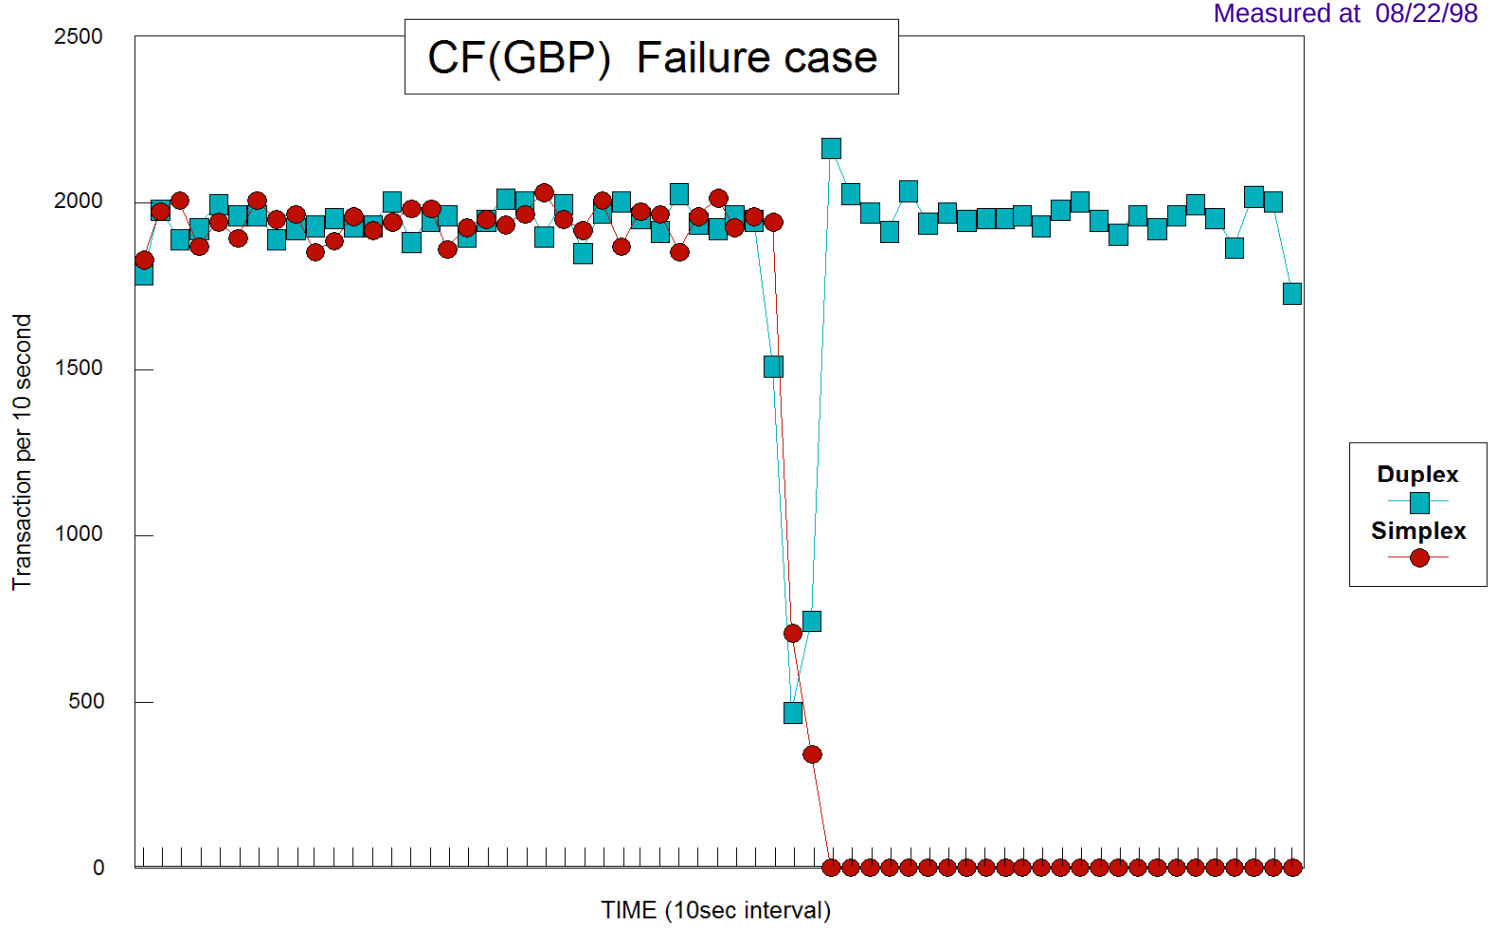

DB2 Data Sharing Availability: Common Issues

- **2-CEC / 2-ICF configuration without System Managed Duplexing**
  - Risk of 'double failure' – LOCK1 and an IRLM, or SCA and a MSTR
    - If all IRLMs survive, they can rebuild LOCK1; same with MSTRs and SCA
    - If LOCK1 and IRLM both affected by outage, entire DB2 data sharing group fails (same for SCA and MSTR)
- **Reliance on manual restart processes** ('emergency bridge call')
  - Applies to same system restart or cross system restart
    - Most clients have some automation to restart in place in case of DB2 failure
- **Delay of cross-system restart of DB2 in case of LPAR outage**
  - z/OS has capabilities to isolate failing or 'sick' LPAR from PSX without interaction
  - Manual processes or operator prompts delay completion of DB2 restart
    - Retained DB2 locks not released until DB2 restart; potential application impact
- **Not implementing GBP duplexing**
  - Fear of very small performance cost has led to application outages if GBP, GBPs, or CF lost
    - GRECP recovery automatic by DB2 in most scenarios in DB2 10
    - Outage can be minutes to over an hour. No outage with GBP duplexing
- **Not cloning applications and routing around failures**
  - CICS, MQ, batch initiators, etc., can all be single points of failure (SPoF)
  - Numerous examples where DB2 data available, but application cannot access DB2 since application cannot run on other LPAR where DB2 is available
- **Long running UoRs may impact service restoration**
- **Disaster Recovery**
  - Not covered in this presentation

Systems Management

- **Goal: single-system image**
- **DB2 data sharing "ease of use" features:**
  - Command prefix support
  - Group attachment name
    - Subgroup attach in DB2 10
  - DDL, Bind, utilities, authorization are all "group scope"
  - Single DDF location name for the DB2 data sharing group
    - Subset: LOCATION ALIAS
    - DB2 10: dynamic location alias
  - "Group scope" on display output
  - "Group scope" for online performance monitors
  - Log merging for replication (IFCID 306)
  - Symbolic &IWMSSNM for WLM Application Environments

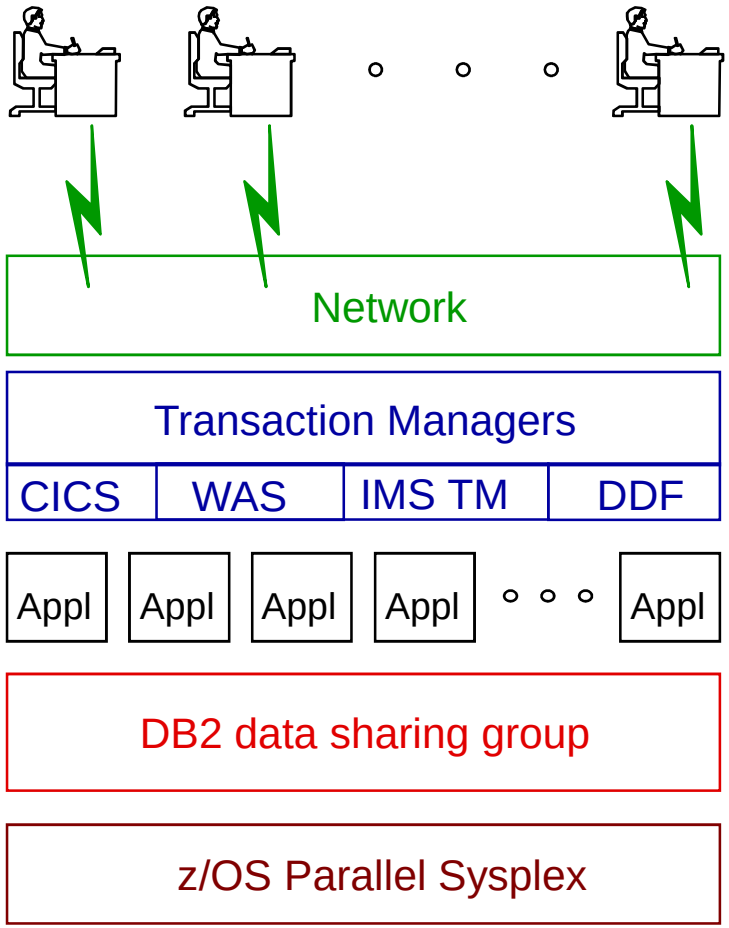

Dynamic Workload Balancing

- **Workload Manager (WLM)**
- **CICSPlex System Manager (CPSM)**
  - Route workload between CICS TORs and AORs
- **IMS Transaction Manager (TM)**
  - Shared message queues
  - BMPs
- **Websphere (incl. MQ)**
  - WAS on z – separate instances
  - WAS on distributed
  - MQ shared queues
- **Distributed access (DDF)**
  - Dynamic Virtual I/P Addressing (DVIPA)
  - Weighted server list sent to clients by DB2
- **Sysplex Distributor (z/OS Commserver)**

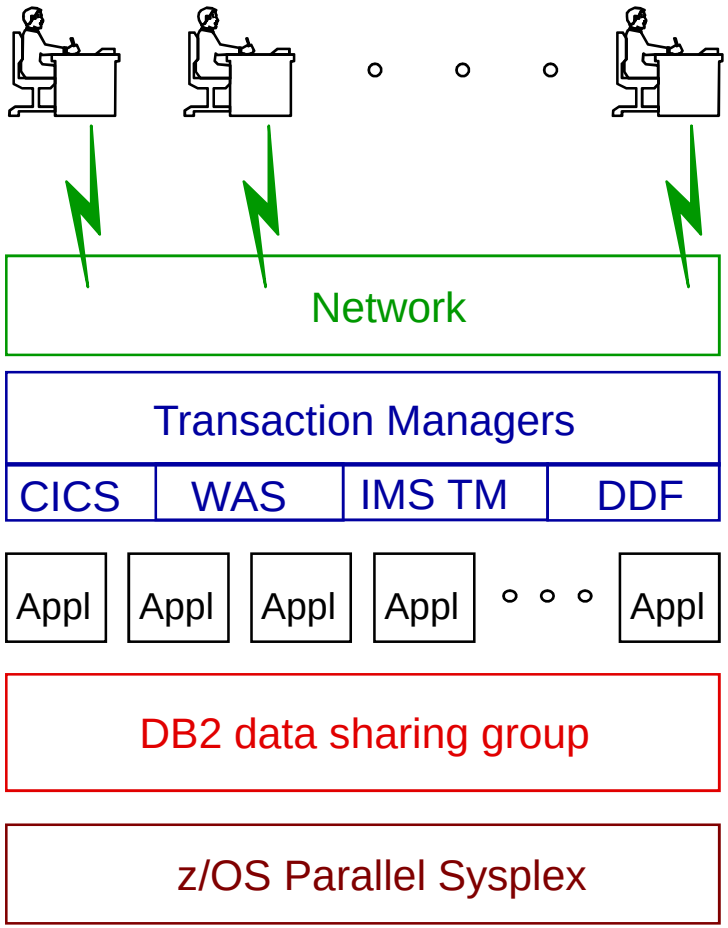

Websphere

To exploit DB2 Data Sharing workload balancing and transparent failover, application server connection pool allows access across transports to each DB2 member

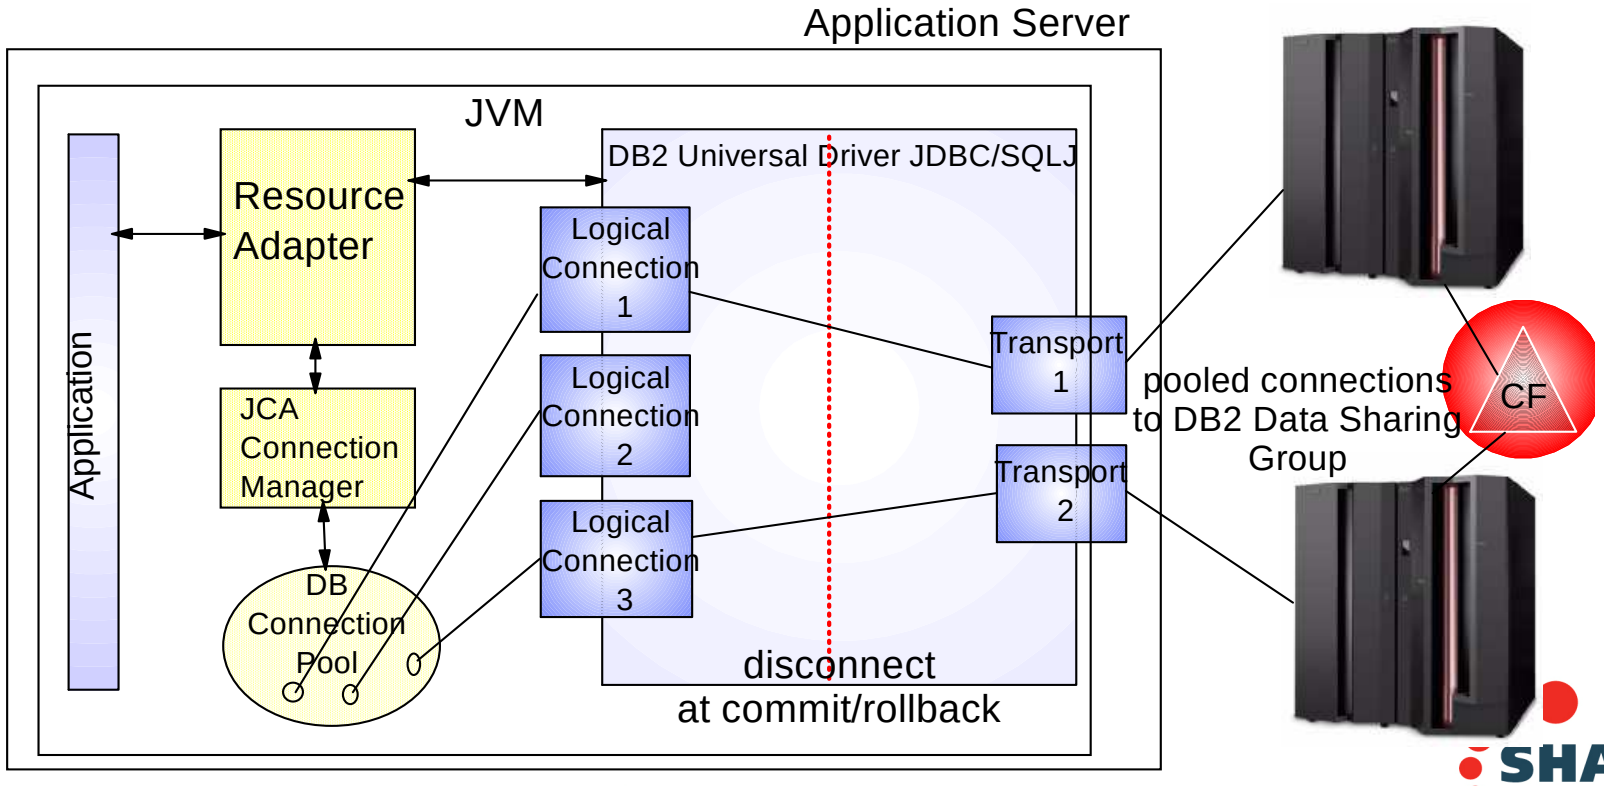

Sysplex Distributor and DVIPA

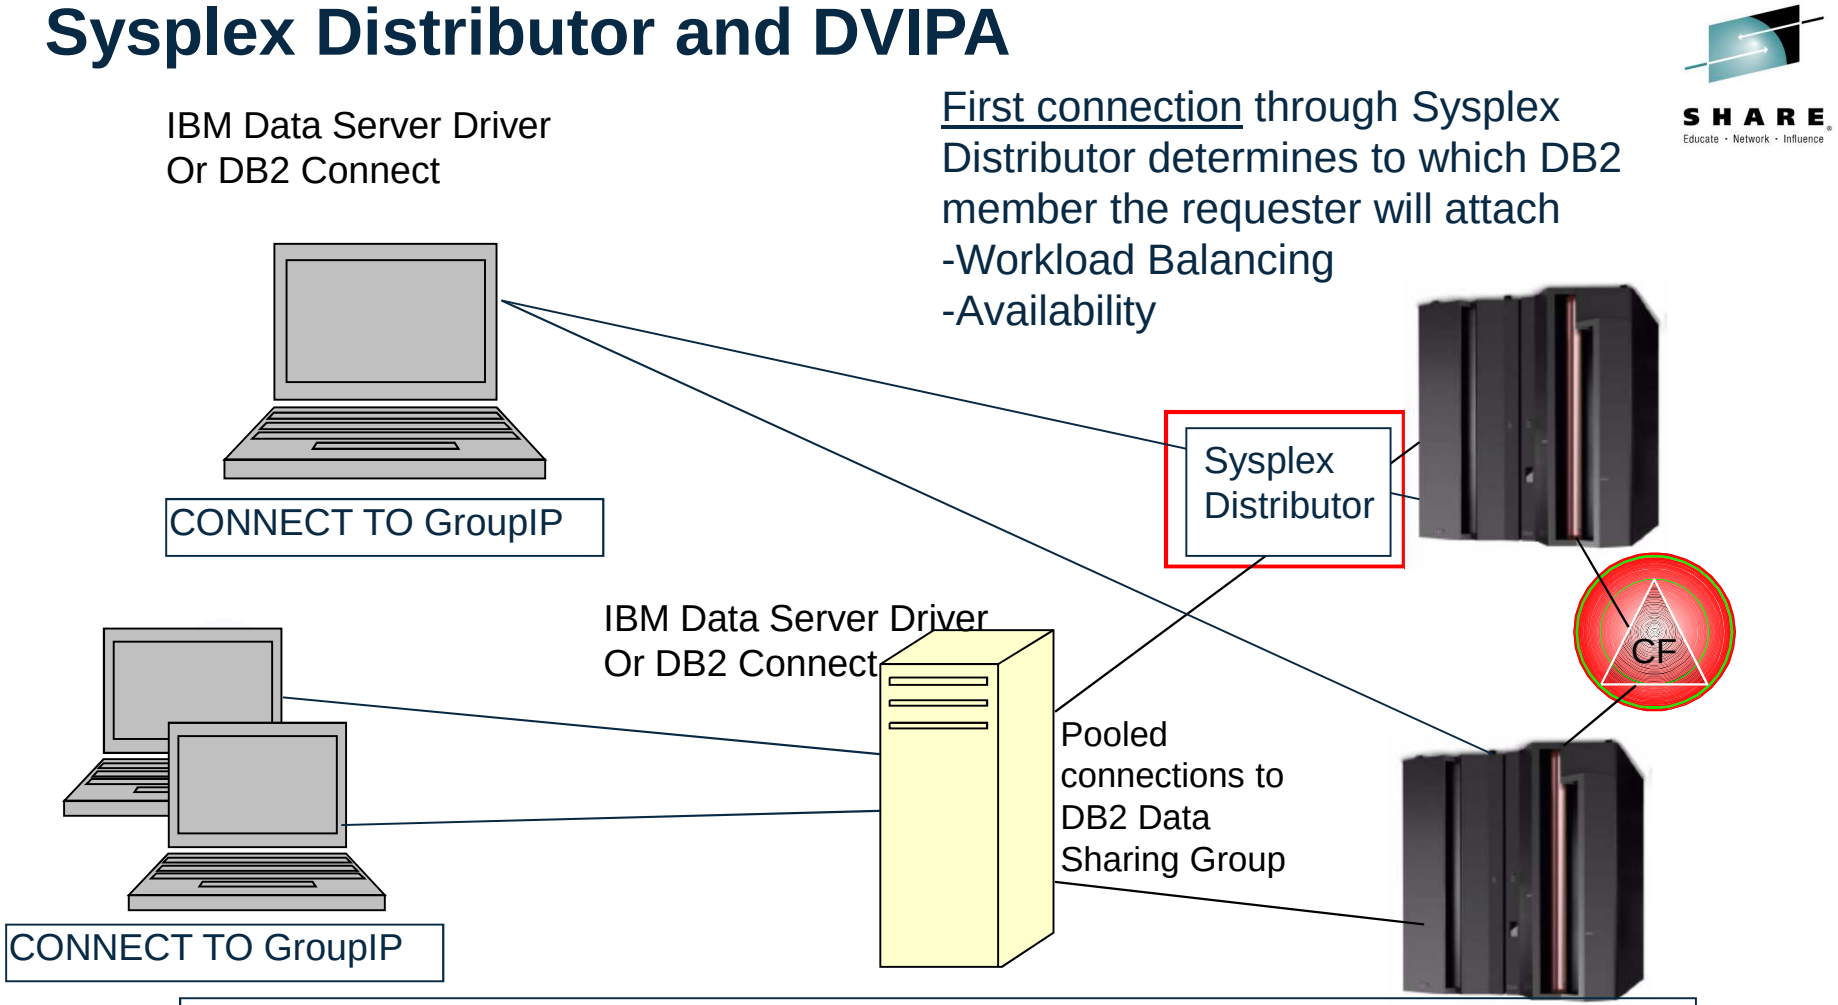

- Direct client connection is preferred (upper part of diagram)
- Gateway configurations still exist but not strategic
  - Path length, SPoF, PD/PSI complexity
  - DB2 10 NFM has Profile Support - reduces need for gateway configuration

IBM Data Server Client Packages

- **IBM Data Server Driver for JDBC and SQLJ**
  - For Java-only environments
- **IBM Data Server Driver**
  - Includes all runtime drivers
- **Also available:**
  - IBM Data Server Runtime Client
    - Includes Command Line Processor
    - Larger footprint
  - IBM Data Server Client
    - Full function development and DBA client
    - Largest footprint
- **All downloadable from www.ibm.com**
  - All require DB2 Connect license

DB2 Data Sharing Workload Balancing: Common Issues

- **Managing DDF traffic across multiple members**
  - DDF traffic ‘stuck’ on busy DB2
    - DB2 10 introduced new parameters:
      - MAXCONQW – time waiting for active DBAT
      - MAXCONQN – depth of queue waiting for active DBAT
    - DB2 will close connections if either value exceeded, allowing client to connect to other member of data sharing group
  - Runaway requesters consume all DBATs (MAXDBAT or CONDBAT)
    - Other DDF requesters shut out, with a variety of symptoms and issues
      - ‘Denial of service attack’
    - DB2 10 Profile support allows limit on connections or threads based on variety of qualifiers
      - Qualifiers and behavior can vary on different data sharing members
- **Clients may think ‘workload balancing’ means 50-50 or 25-25-25-25**
  - Workload balancing based on numerous inputs to WLM and participant subsystems, including ‘health’ and resources
- **SW charge ‘penalty box’**
  - Client may want certain work isolated to one environment to limit ISV SW charges

DB2 Data Sharing: Benefit Summary

- Data sharing technology provides the base to allow DB2 to deliver continuous availability and nearly unlimited scalability into the future
- DB2 Data Sharing is a proven technology
  - Many customers have implemented DB2 data sharing production
- DB2 development continues to deliver data sharing enhancements in each new release

Agenda

- **DB2 data sharing overview**
  - What is it? Why do our clients implement it?
- **DB2 data sharing concepts**
  - Parallel Sysplex
  - Coupling Facilities (CFs)
  - CF Structures
- **DB2 data sharing configurations**
- **Common Issues**
  - Performance
  - Availability
  - Dynamic workload balancing
- **Resources and FAQs**

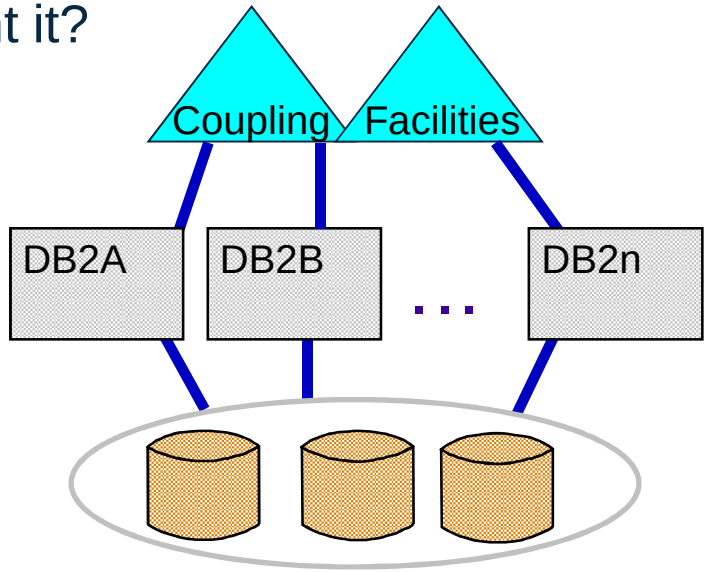

Bibliography

- ***Data Sharing: Planning and Administration***
  - DB2 10 for z/OS: SC19-2973
  - DB2 11 for z/OS: SC19-4055
  - Knowledge Center for DB2 10 and DB2 11
- **Redbooks**
  - *DB2 9 for z/OS: Distributed Functions:* SG24-6952-01
  - *Achieving the Highest Levels of Parallel Sysplex Availability*: SG24-6061
  - *Parallel Sysplex Application Considerations*: SG24-6523
  - *TCP/IP in a Parallel Sysplex*: SG24-5235-02
  - *DB2 for z/OS: Data Sharing in a Nutshell:* SG24-7322
    - Currently being updated
  - *DB2 9 for z/OS: Resource Serialization and Concurrency Control:* SG24-4725-01

Common Questions

- Overhead and capacity
- Application changes

Overhead and capacity

- "I've heard as much as 50% overhead"
  - Not in well-defined PSX and DB2 data sharing environment
  - 5 to 10 % more reasonable estimate
- "What can I expect?"
  - "It depends…"
- "How should I size my CF structures?"
  - CFSizer
    - http://www.ibm.com/systems/support/z/cfsizer/
    - Installation and Migration Guide
      - DB2 10 for z/OS: Installation and Migration Guide: GC19-2974
      - DB2 11 for z/OS: Installation and Migration Guide: GC19-4056
      - Knowledge Center for DB2 10 and DB2 11

Application changes

- "Don't my applications have to change?"
  - SQL interface does not change
  - However, locking and commit frequency may impact data sharing performance
    - Commit frequently – long-time recommendation
    - Take advantage of lock avoidance
      - ISO(CS) or ISO(UR)
      - CURRENTDATA NO
  - New messages and return codes
  - Applications must be able to run on more than one DB2 member for high availability

Thank you

- Thank you for your attention.
- Questions?
  - For questions that come to mind later, please send to:
    - Mark Rader: mrader@us.ibm.com