Db2 for z/OS Data Sharing: Configurations and common issues

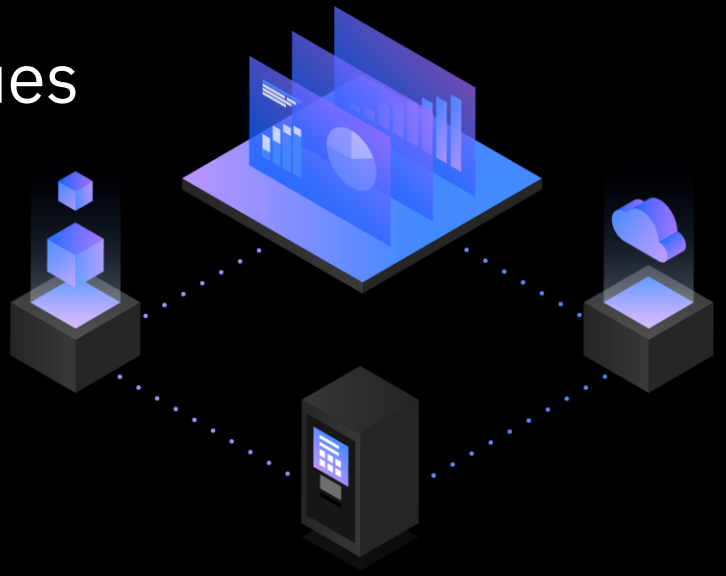

New England Db2 User's Group

December 1, 2022

Mark Rader
Db2 for z/OS Specialist
IBM Z Washington Systems Center

IBM

Agenda

- Db2 data sharing overview
  - What is Db2 data sharing?
  - What are the major components?
- Db2 data sharing configurations
  - General options
- Common issues
  - Performance
  - Availability
  - Dynamic workload balancing
- Recent data sharing enhancements
  - Db2 12 and Db2 13

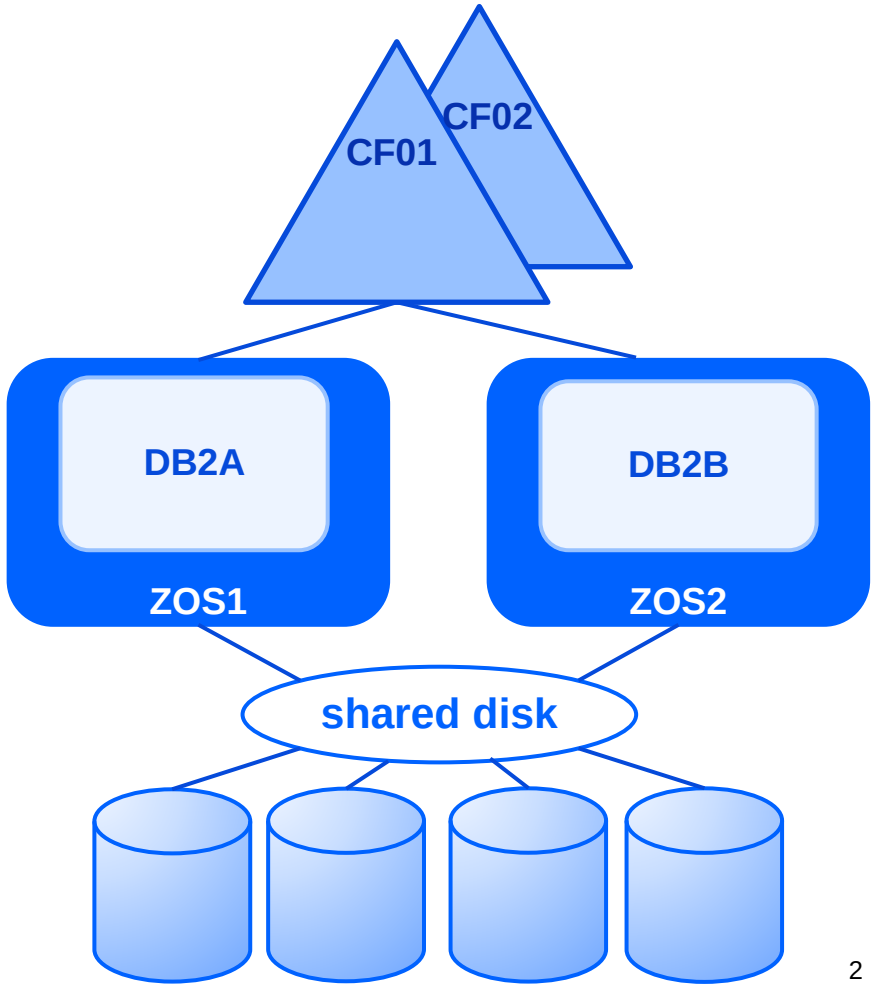

Db2 data sharing overview

- Db2 data sharing and Parallel Sysplex constitute the infrastructure that provides enterprise applications using relational data the highest levels of:
  - Availability – protection against planned and unplanned outages
  - Non-disruptive scalability – growth to handle business, market, or regulatory changes
  - Dynamic workload balancing – using available system capacity, handling spikes

- Application investment protection
  - SQL interface is unchanged for data sharing
  - Excellent scaling
  - Applications do not need to become "cluster aware" as processing nodes (Db2s and LPARs) are added to your configuration

Parallel Sysplex

- Parallel Sysplex components:
  - Sysplex Timer Protocol (STP)
  - Coupling facility (CF) LPARs
    - High-speed shared memory
    - CF control code (CFCC)
    - Structures: lock, cache, list
  - CF links
  - CF resource management (CFRM) policy
  - z/OS components
    - Cross-system extended services (XES)
    - Cross-system coupling facility (XCF)

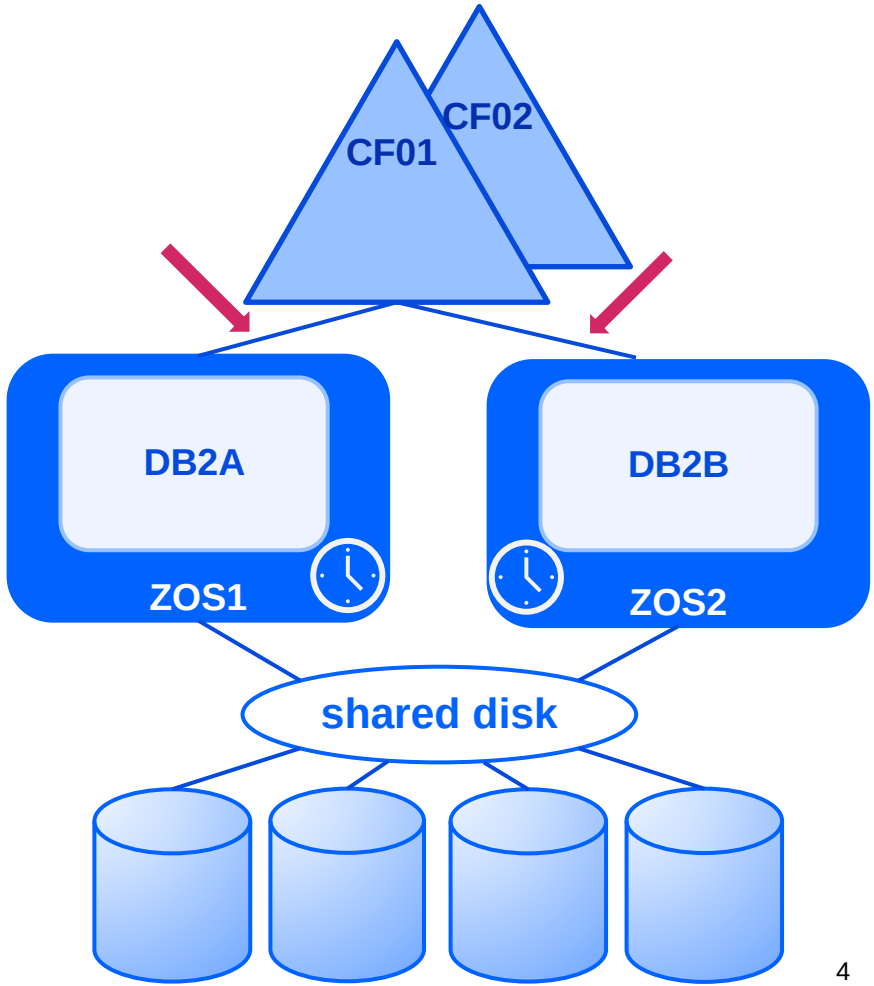

Db2 data sharing concepts

- A Db2 ***data sharing group*** consists of:
  - 1 to 32 Db2 ***members*** (subsystems) with:
    - Single Db2 Catalog and Directory
    - Db2 and user data on shared disk
    - Single LOCATION for distributed access (distributed data facility, or DDF)
  - Each Db2 ***member*** has:
    - Read/write access to its own active log, archive log, and boot strap data sets (BSDS)
    - Read access to every other member's active log, archive log, and BSDS

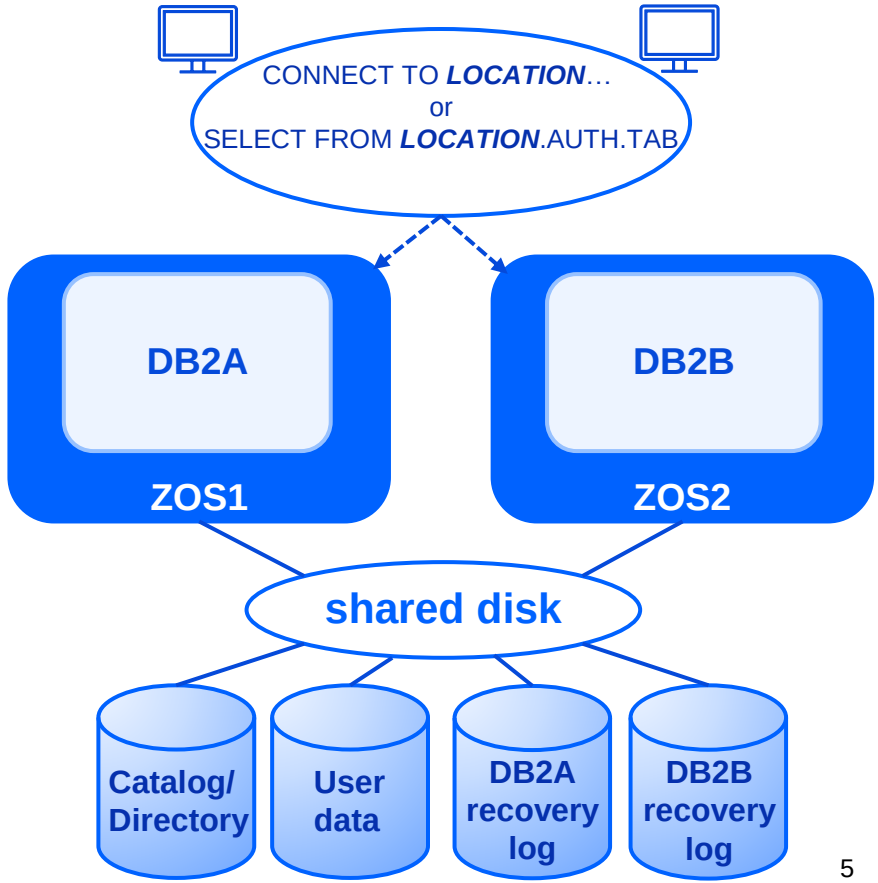

Db2 data sharing concepts

- Coupling facilities (CFs) contain:
  - 1 lock structure (LOCK1) per data sharing group *
  - 1 shared communications area (SCA) per group *
  - Multiple group buffer pools (GBPs) per group
    - 1 GBP per local buffer pool (BP) containing shared data
    - Db2 allocates GBP0, GBP8K0, GBP16K0, GBP32K

* = Required structures
- Db2 will not start if it cannot allocate LOCK1 or SCA
- Db2 will fail if it loses access to LOCK1 or SCA and cannot rebuild it

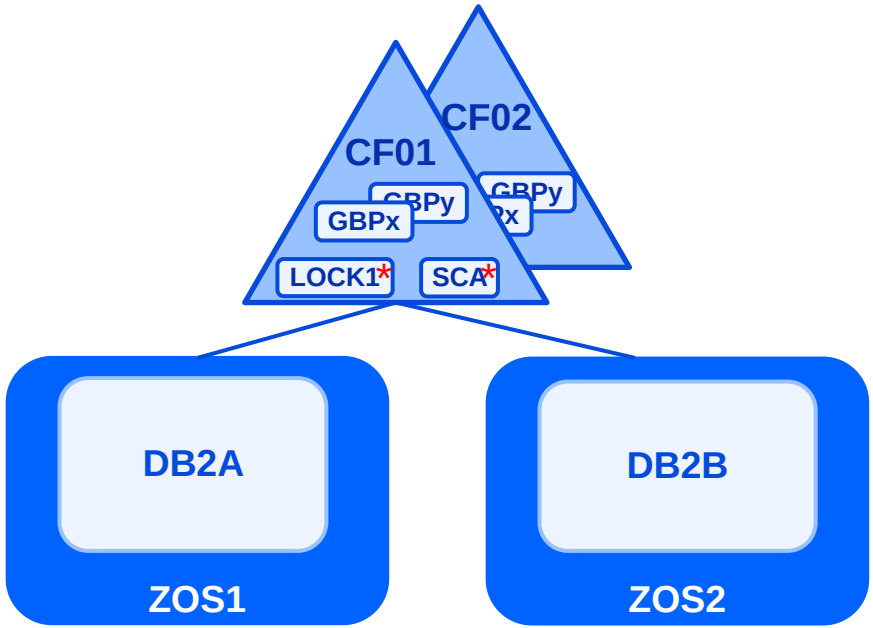

Shared communication area (SCA)

- List structure
- DB2xMSTR
- Required structure *
  - All Db2 members required to rebuild SCA
- Group-wide status information
- Generally not a performance concern

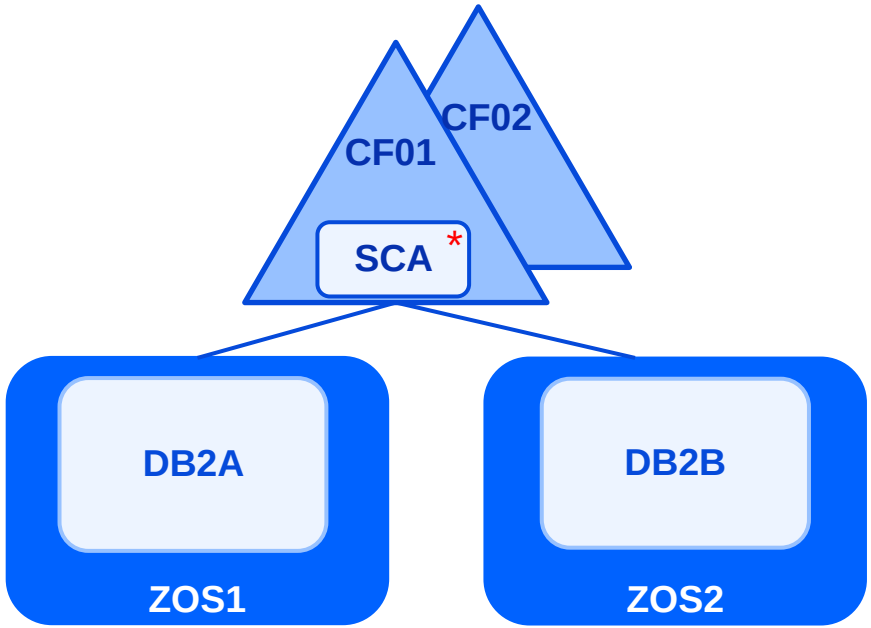

Lock structure (LOCK1)

- DB2xIRLM
- Required structure *
  - All IRLMs required to rebuild LOCK1
- Size influences performance
- Lock table
  - Lock table entries (LTEs) represent hash classes
    - S – shared *or*
    - X – exclusive
- Modify lock list (record list entries or RLEs)
  - Update-type locks
  - Locks retained upon Db2 or z/OS failure
    - Resources protected by these locks not available
    - ***Db2 restart required to release retained locks***

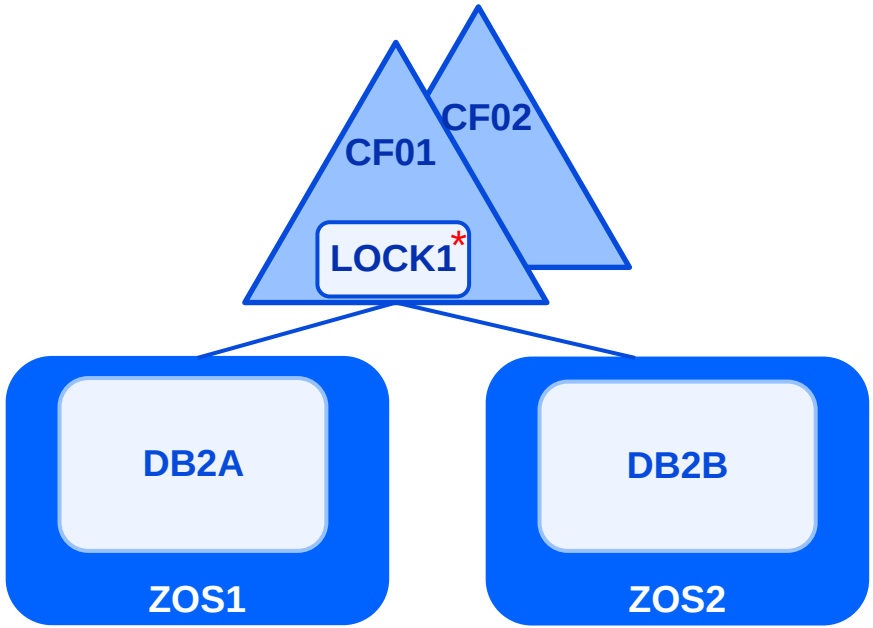

Lock structure (LOCK1)

- Logical locks (L-locks)
  - Track concurrency
    - Many readers *OR*
    - One writer
- Physical locks (P-locks)
  - Track coherency
    - Resource held in 1 or more members
    - Other member wants to change resource
  - Negotiable
    - DB2A and DB2B negotiate final lock state
    - One example: GBP dependency

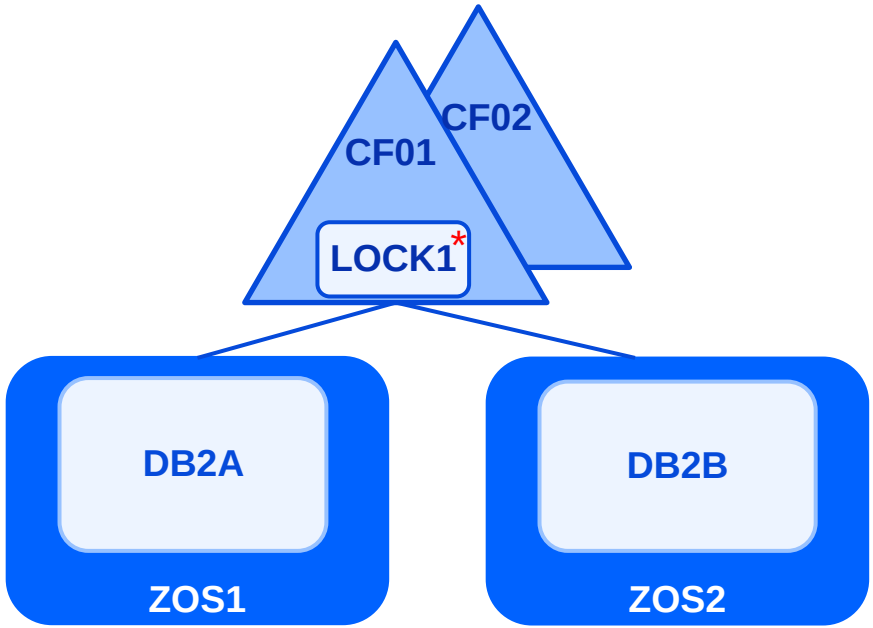

Group buffer pools (GBPn)

- Cache structures
- DB2xDBM1
- Buffer coherency
- Size influences performance
- Directory entries
  - Pointers to pages in local buffer pools, GBPs
- Data elements
  - Table or index pages
  - Changed pages (default)
    - Other options available

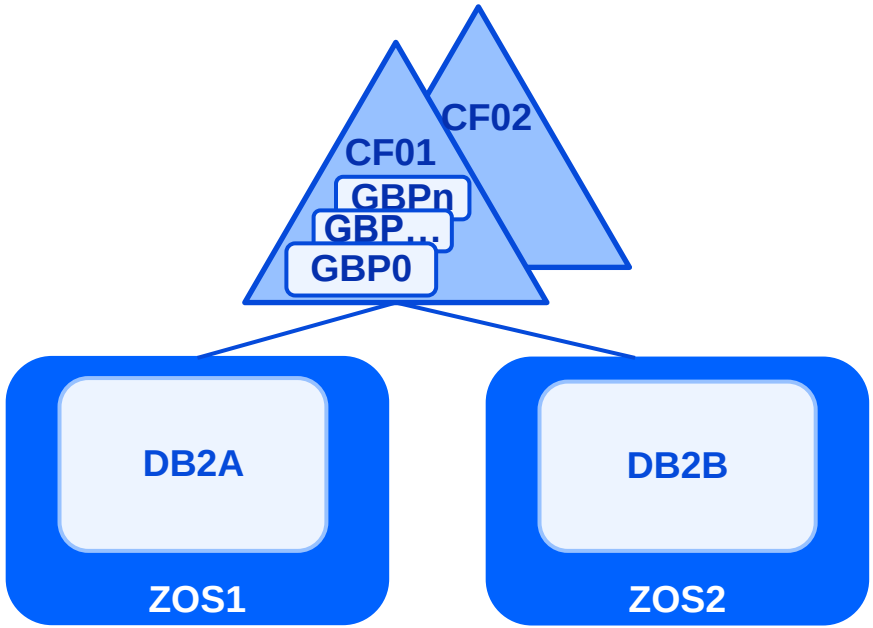

Group buffer pools (GBPn)

- Changed pages
  - ***Castout*** to disk asynchronously
    - Dashed **green** arrow
- Most current state of page
  - Only copy ? – needs to be protected to avoid lengthy recovery
- Duplexed GBPs recommended – ***strongly recommended***
  - Primary GBPs – all access (register page, cache changed pages, castout, read for local buffer miss)
  - Secondary GBPs – copy of changed pages on other CF
  - “User managed” duplexing

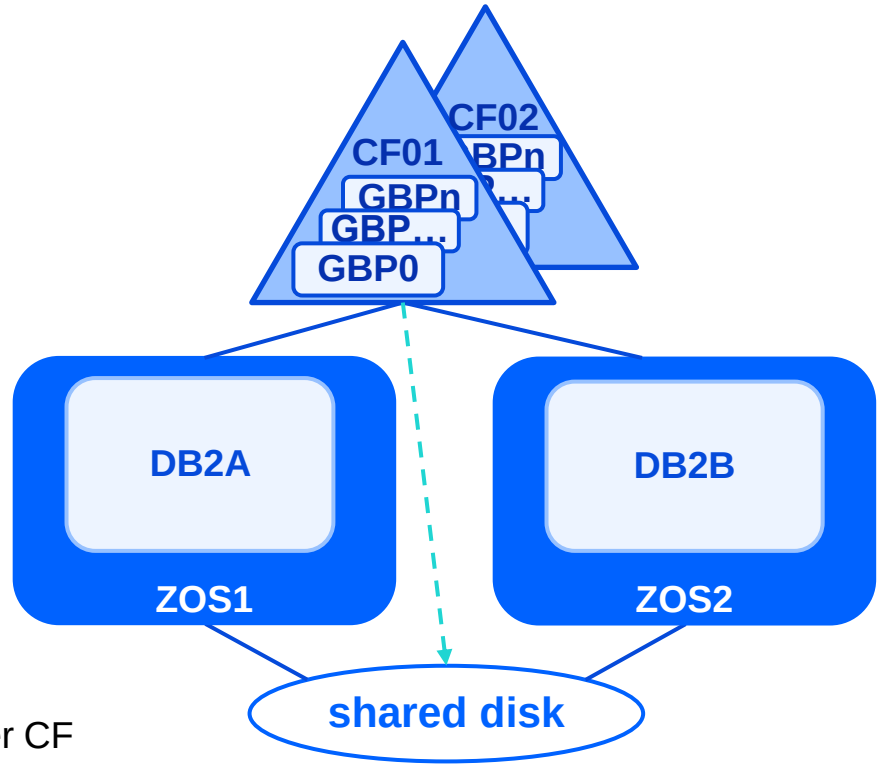

Group buffer pool (GBP)

- Simplified view (primary GBP)

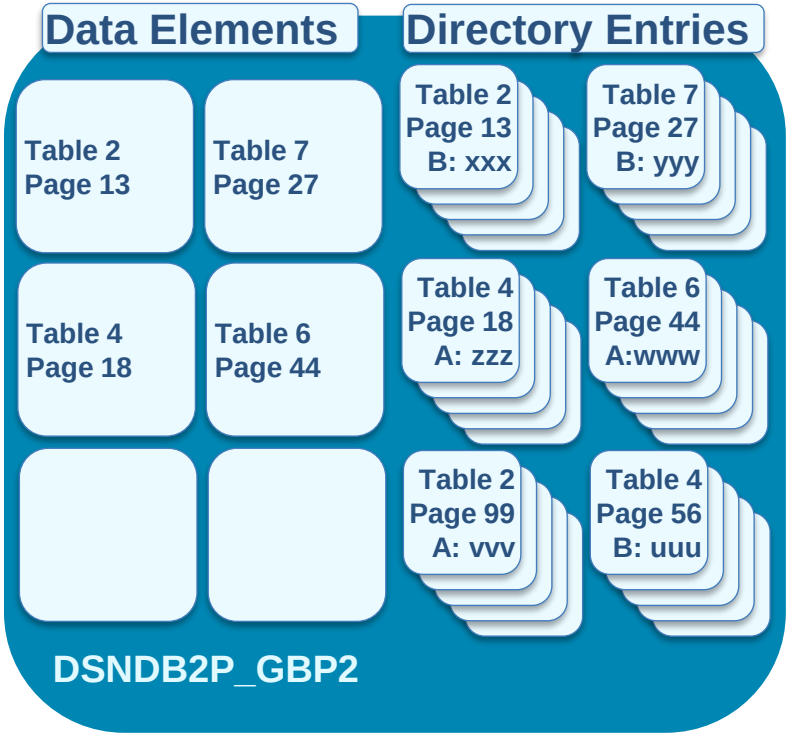

Data sharing configurations

General options

Db2 non-data sharing

- Db2 subsystem DB2P
  - DB2PDBM1
  - DB2PMSTR
  - DB2PIRLM
  - DB2PDIST
  - DB2PWLMx
- Local attach to ‘DB2P’
  - CICS, IMS, TSO, batch, Websphere on Z, Call attach facility…
- Network attach:
  - IP address or domain name *and*
  - Port *and*
  - Location (e.g. LOCDB2P)

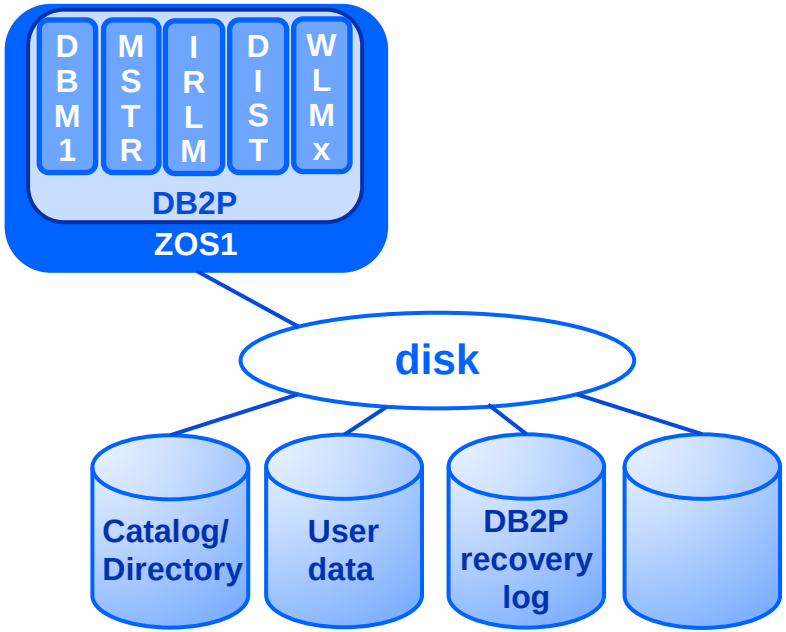

Db2 1-way data sharing

- Db2 ***data sharing group***
  - DSNDB2P – group name
  - Db2 member **DB2A**
    - Address spaces DB2A*nnnn*
    - Group attach name **DB2P**
    - Benefit: no change to existing JCL required
- CF Structures
  - **DSNDB2P_SCA** – required *
  - **DSNDB2P_LOCK1** – required *
  - DSNDB2P_GBPn – optional at this point

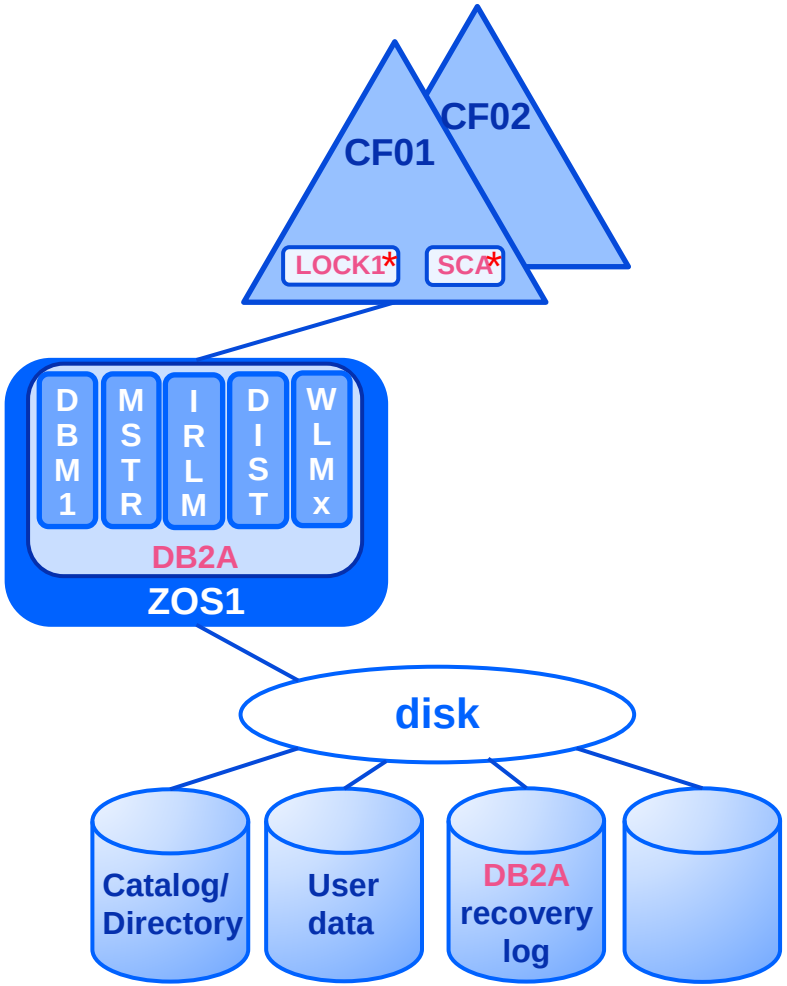

Db2 2-way data sharing

- Db2 ***data sharing group***
  - DSNDB2P
  - Db2 member **DB2A**
    - Group attach name **DB2P**
  - Db2 member **DB2B**
    - Group attach name **DB2P**
  - Up to 32 members
- CF Structures
  - DSNDB2P_SCA
  - DSNDB2P_LOCK1
  - **DSNDB2P_GBPn**
    - One for each buffer pool containing shared data

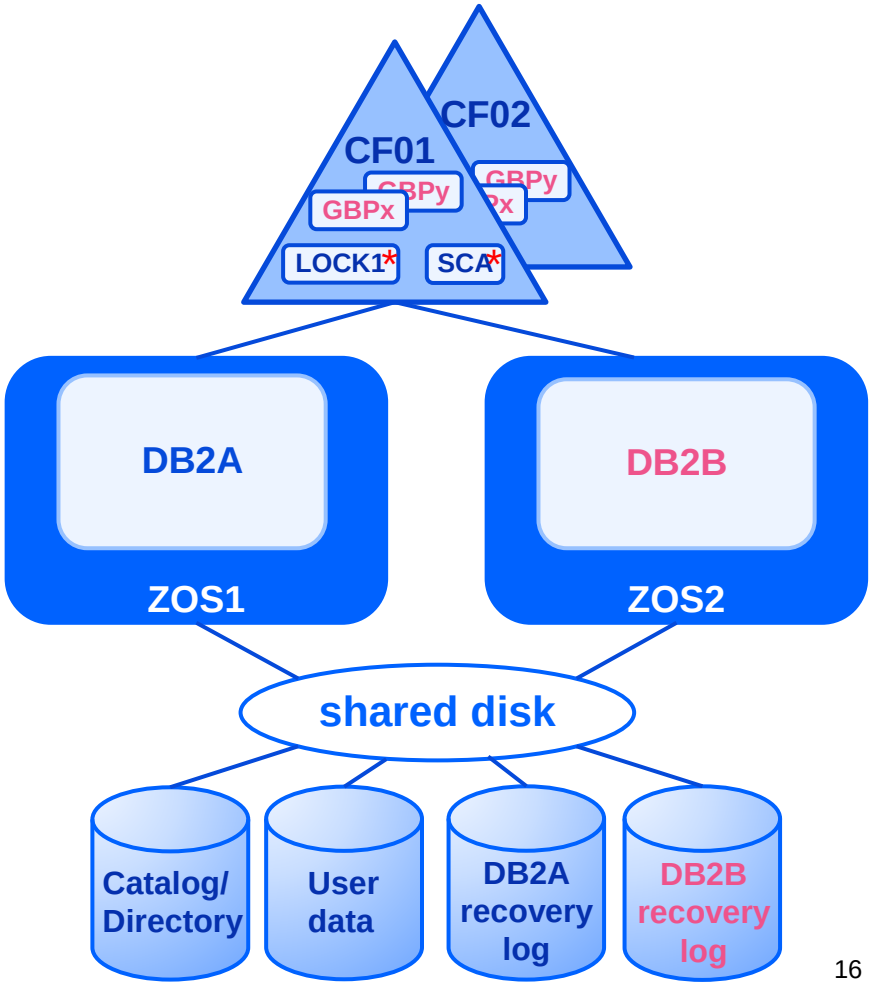

Parallel Sysplex configurations - A

- Original configuration
  - 2 external CFs
  - SCA and LOCK1 isolated from Db2, IRLM
  - GBPs spread across CFs
  - 4 systems (CECs)
  - Tolerate the failure of any one system
    - Applications continue to access data

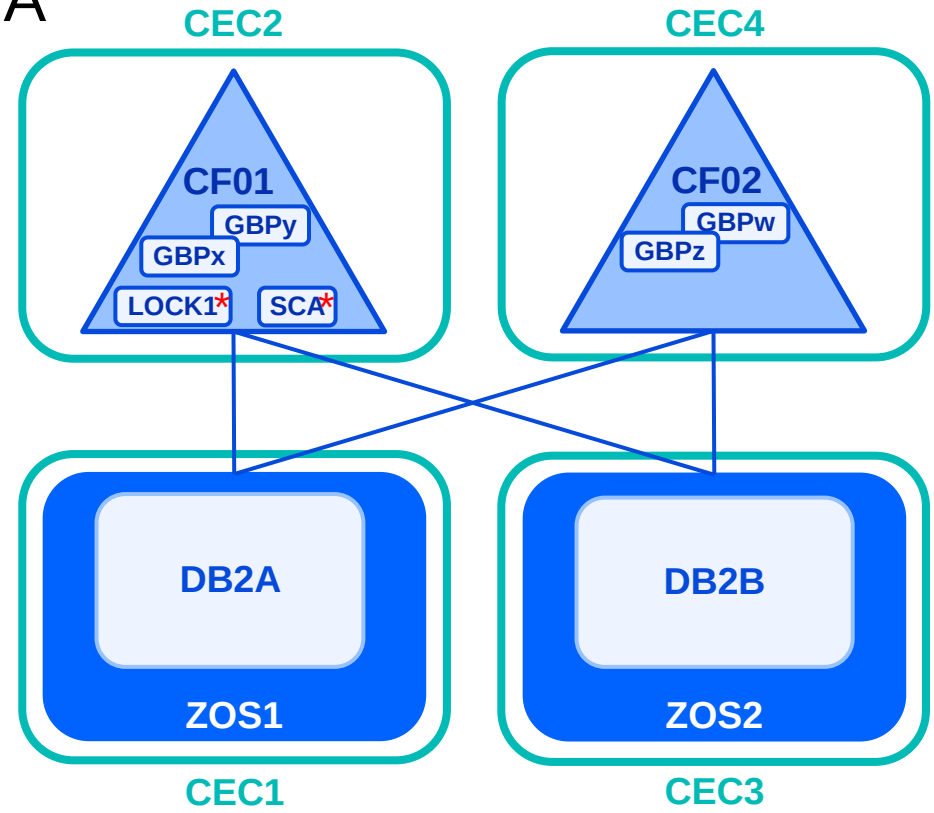

Parallel Sysplex configurations - B

- Second configuration
  - 1 external CF – CF01 on **CEC2**
  - 1 integrated CF (ICF) – ICF02 on **CEC3**
  - SCA and LOCK1 isolated from Db2, IRLM
  - GBPs spread across CFs
    - Duplexed GBPs
    - Primary GBPs spread across CFs
      - Secondary GBP on other CF
  - Tolerate the failure of any one system
    - Applications continue to access data

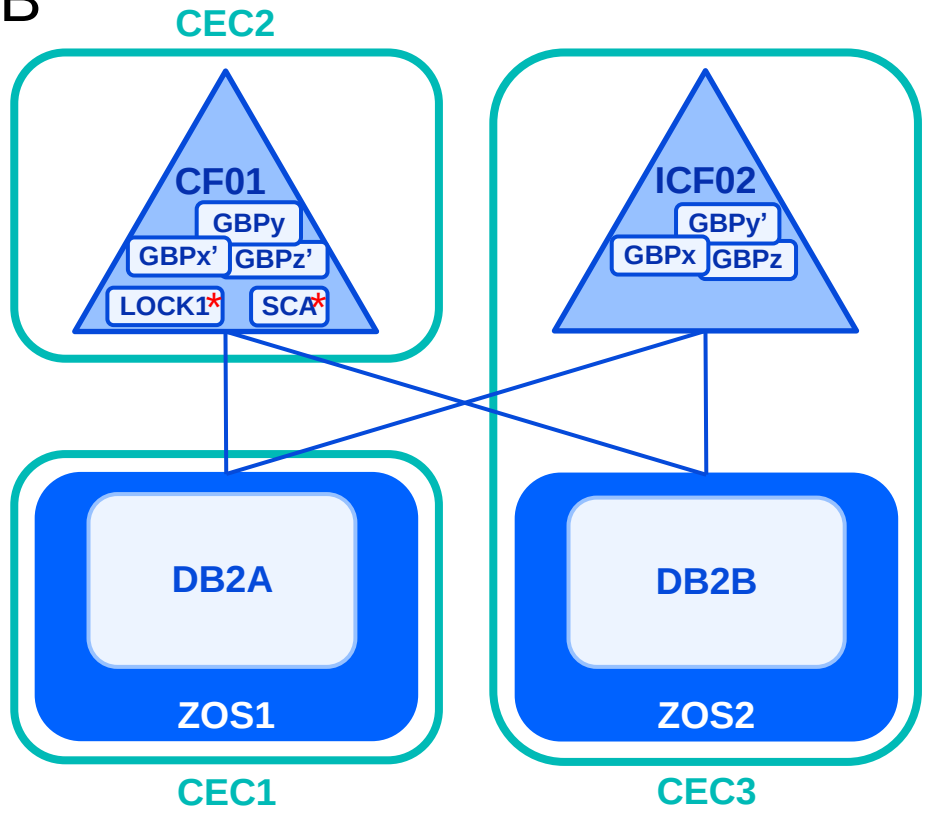

Parallel Sysplex configurations - C

- Third configuration
  - 2 integrated CFs, 2 CECs
  - SCA and LOCK1 **not** isolated from Db2, IRLM
  - Duplexed GBPs spread across CFs
- What if one system is lost?
  - It depends…
  - If CEC3, would be loss of 1 Db2: DB2B
    - DB2A would continue running
  - If CEC1, would be loss of entire data sharing group
    - DB2B alone would not be able to rebuild LOCK1 and SCA, so DB2B would fail, too.

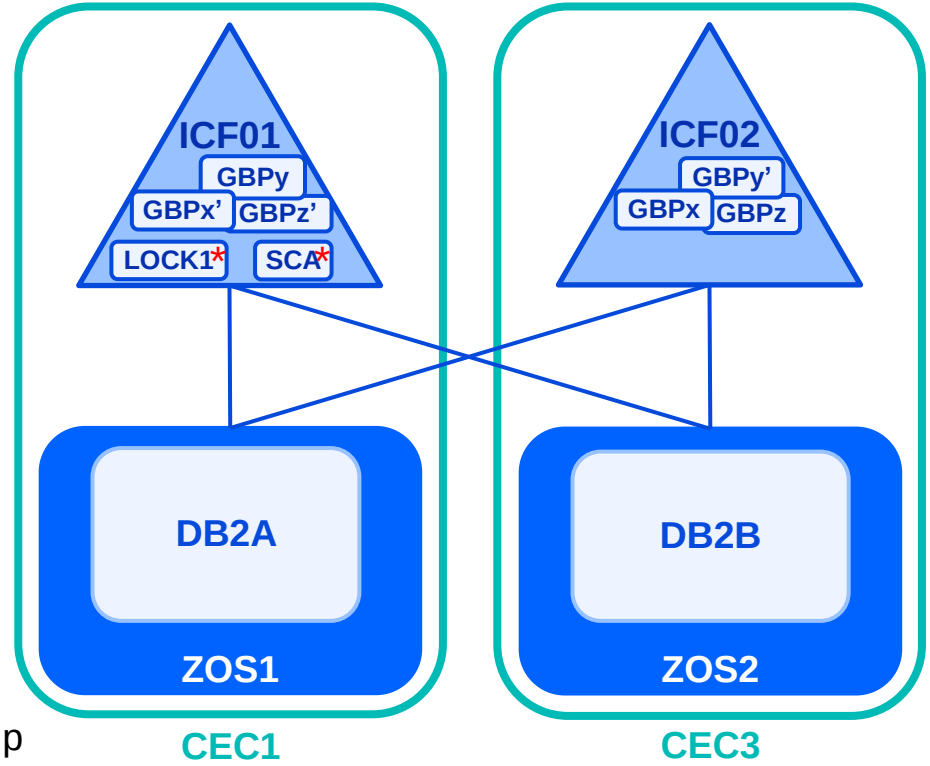

Parallel Sysplex configurations – C: recommended

- **System** managed duplexing
  - 2 integrated CFs, 2 CECs
  - SCA and LOCK1 duplexed
    - Synchronous duplexing
    - GBPs deploy **user** managed duplexing
- If either CEC fails
  - Db2 on other CEC available
  - Restart failed Db2 on other LPAR
    - Release retained locks
- Increased cost of lock requests
  - Performance and CPU costs
  - ***Db2 12: Asynchronous CF duplexing for LOCK1 mitigates performance and CPU costs***

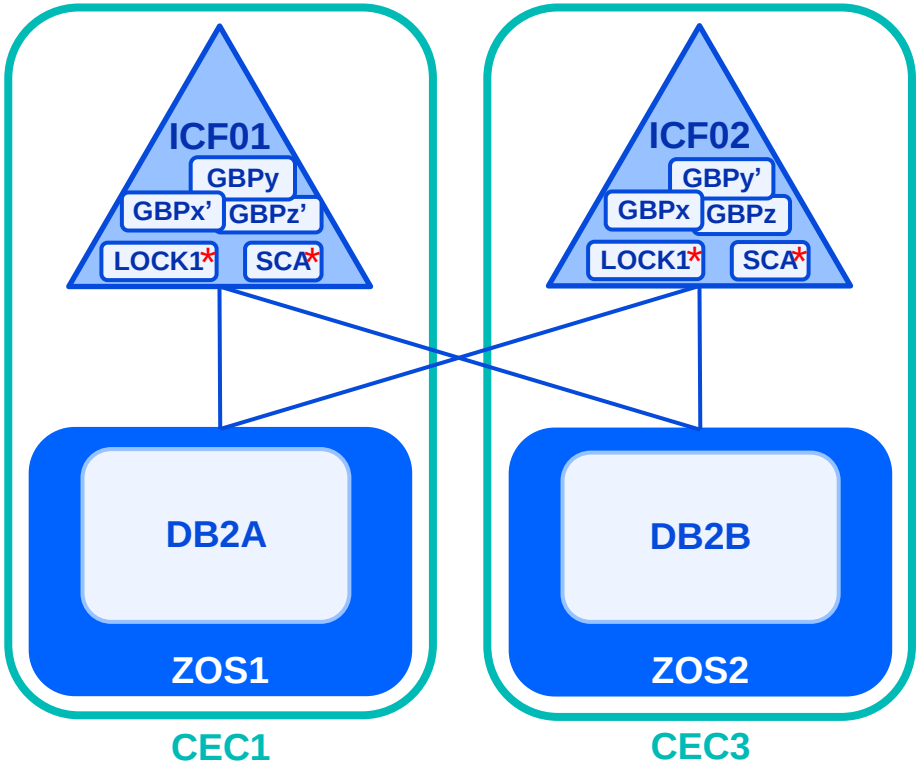

Parallel Sysplex configurations - D

- Special case of configuration B
  - Still all-ICF configuration
  - SCA and LOCK1 on ICF01
  - SCA and LOCK1 isolated from Db2 and IRLM
  - Duplexed GBPs across 2 ICFs
- If any single CEC fails
  - Data sharing group remains available

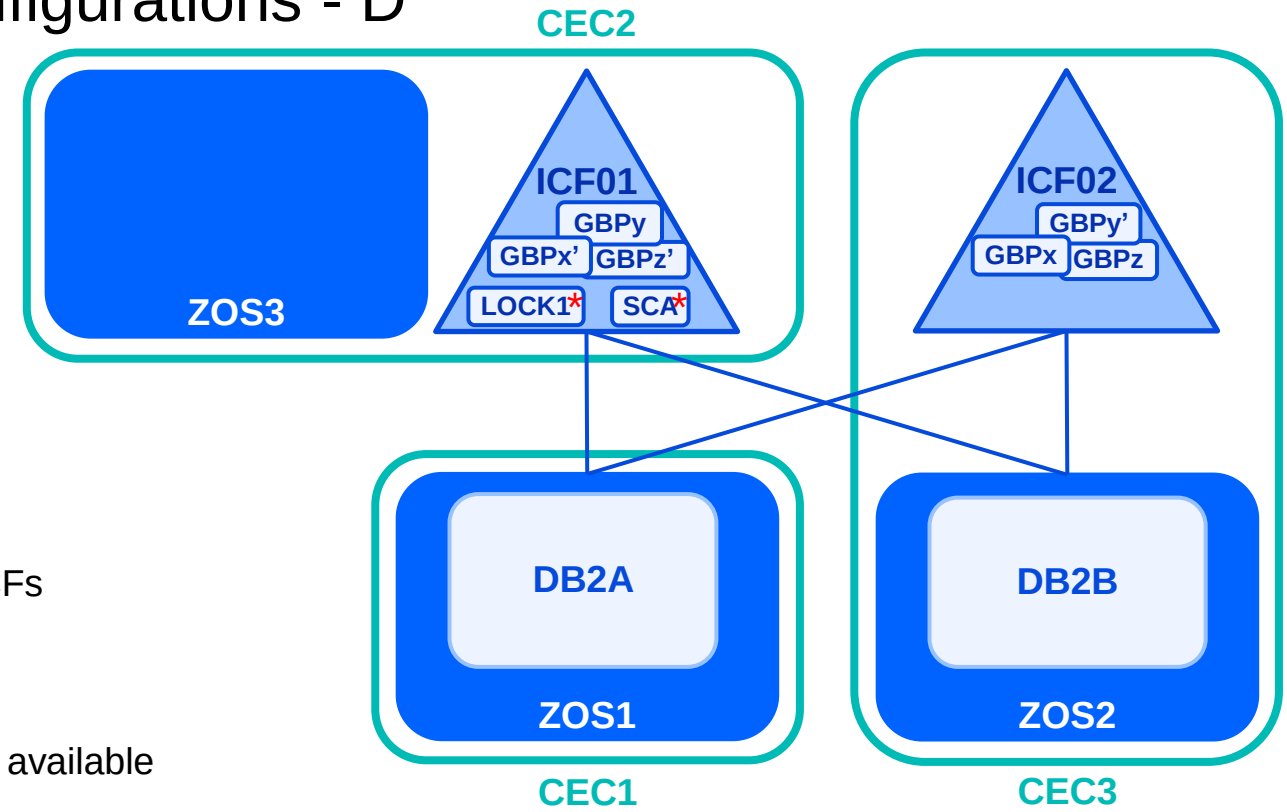

Parallel Sysplex in a box

- All members and LPARs on one CEC
  - No hardware failure isolation at system level
  - Avoid planned outages for software changes
  - Avoid unplanned software outages if applications are cloned
- Risk of hardware outage
  - Unplanned outages
  - Some planned outages

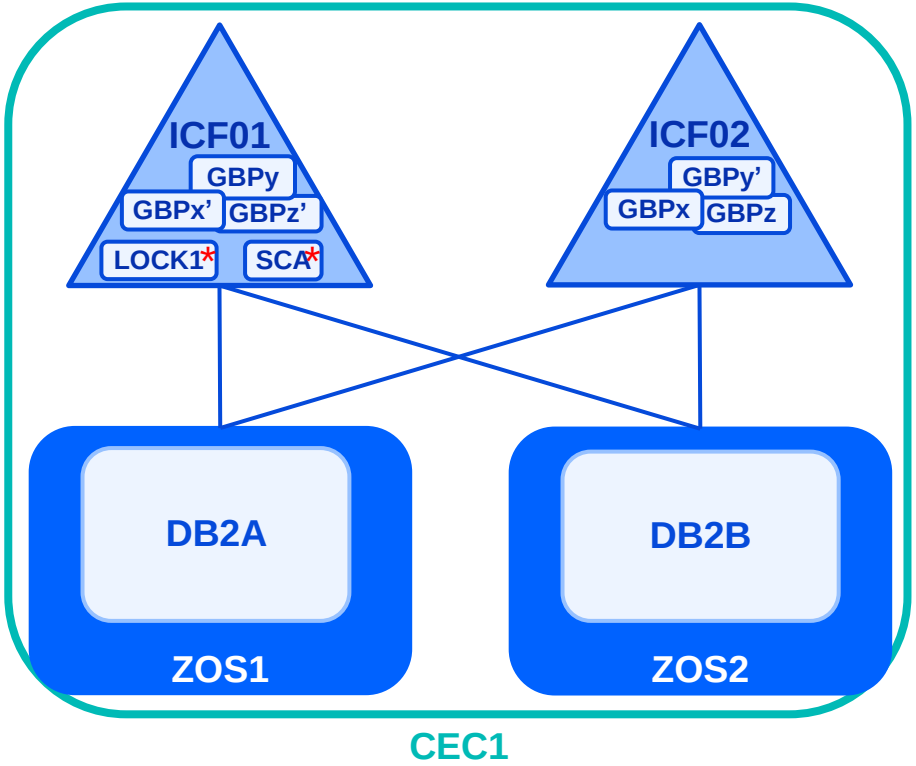

Common issues

Performance

Design goals: data sharing performance

- Little or no performance impact if data not actually shared
  - If no inter-Db2 read/write interest
  - Dynamic recognition of sharing
- Minimal and acceptable CPU overhead if inter-Db2 read/write interest exists
  - Overhead varies based on workload characteristics
- Near-linear scalability when adding 3rd through nth members
  - Majority of data sharing overhead appears going from 1-way to 2-way data sharing
  - Beyond 2-way, additional overhead is minimal

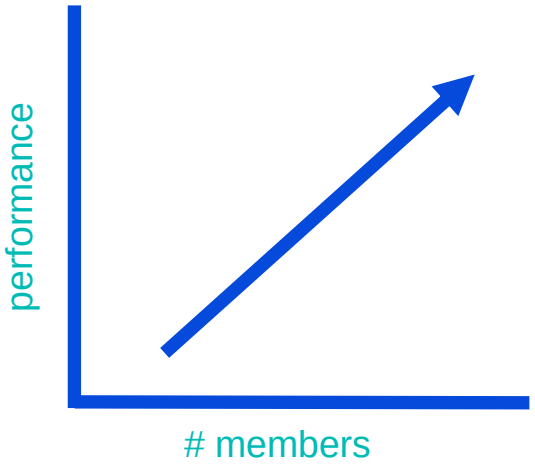

Data sharing performance: critical factors

- Two factors to preserving data integrity in data sharing
  - Inter-system ***concurrency*** control – global locking
    - Multiple readers ***OR*** one writer
  - Inter-system ***coherency*** control – managing changed data
    - One system changes data pages that also reside in other system(s)

LOCK1

LOCK1 + GBPs

- Data sharing overhead based on CPU to manage these factors
  - Thousands to tens of thousands of messages per second
    - Early extreme example: 166,114 synchronous lock requests per second (2008)
    - Several sites have exceeded 200,000 synchronous lock requests per second
      - Latest report peaked at over 387,000 synchronous lock requests per second!
    - One report: 643,300 synchronous lock requests per second!

- Most CF messages for Db2 or IRLM are synchronous
  - Host CPU cost for duration of round trip to CF → CPU busy

Data sharing: managing changed data

- Update and cross-invalidate (XI)

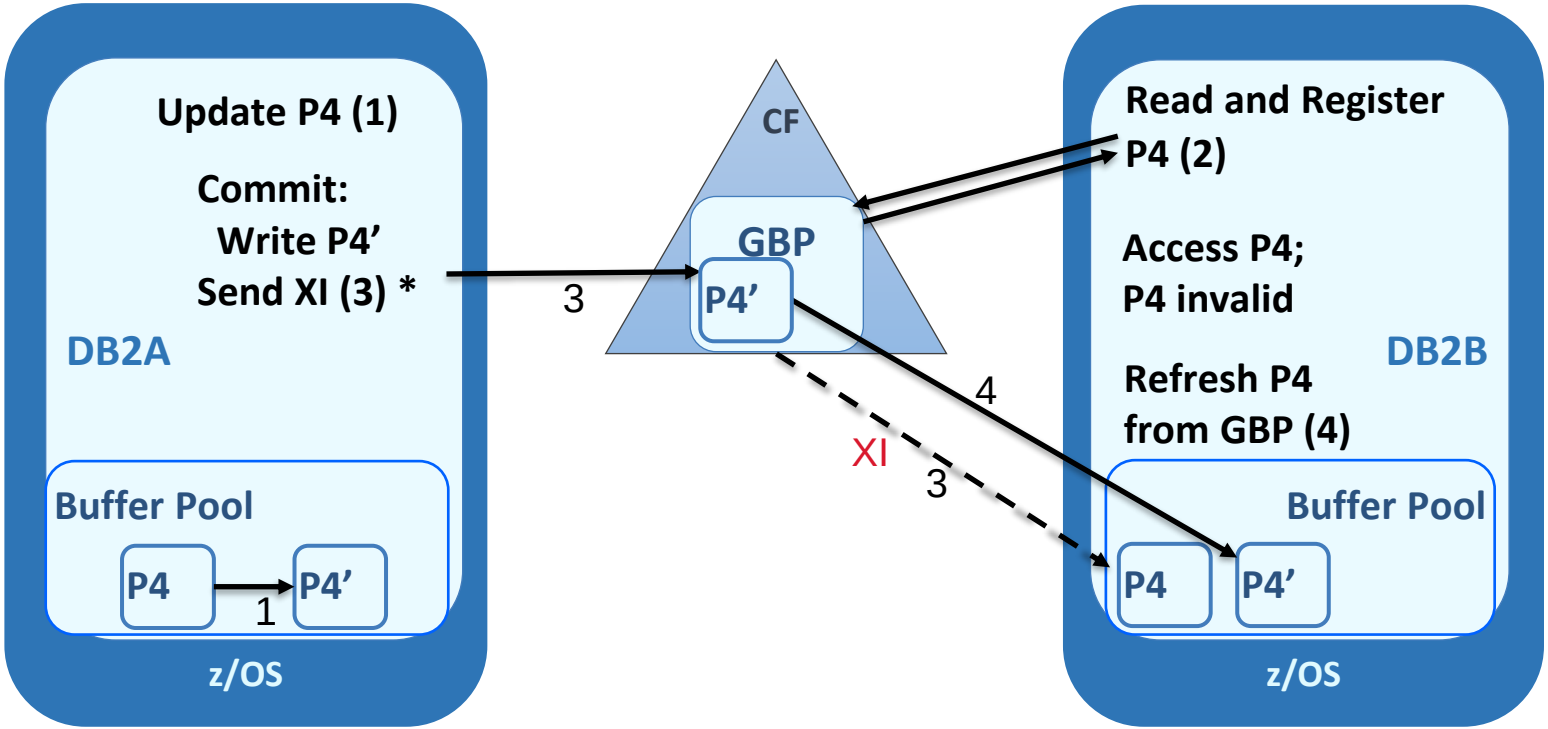

Data sharing: managing changed data

- Castout to disk

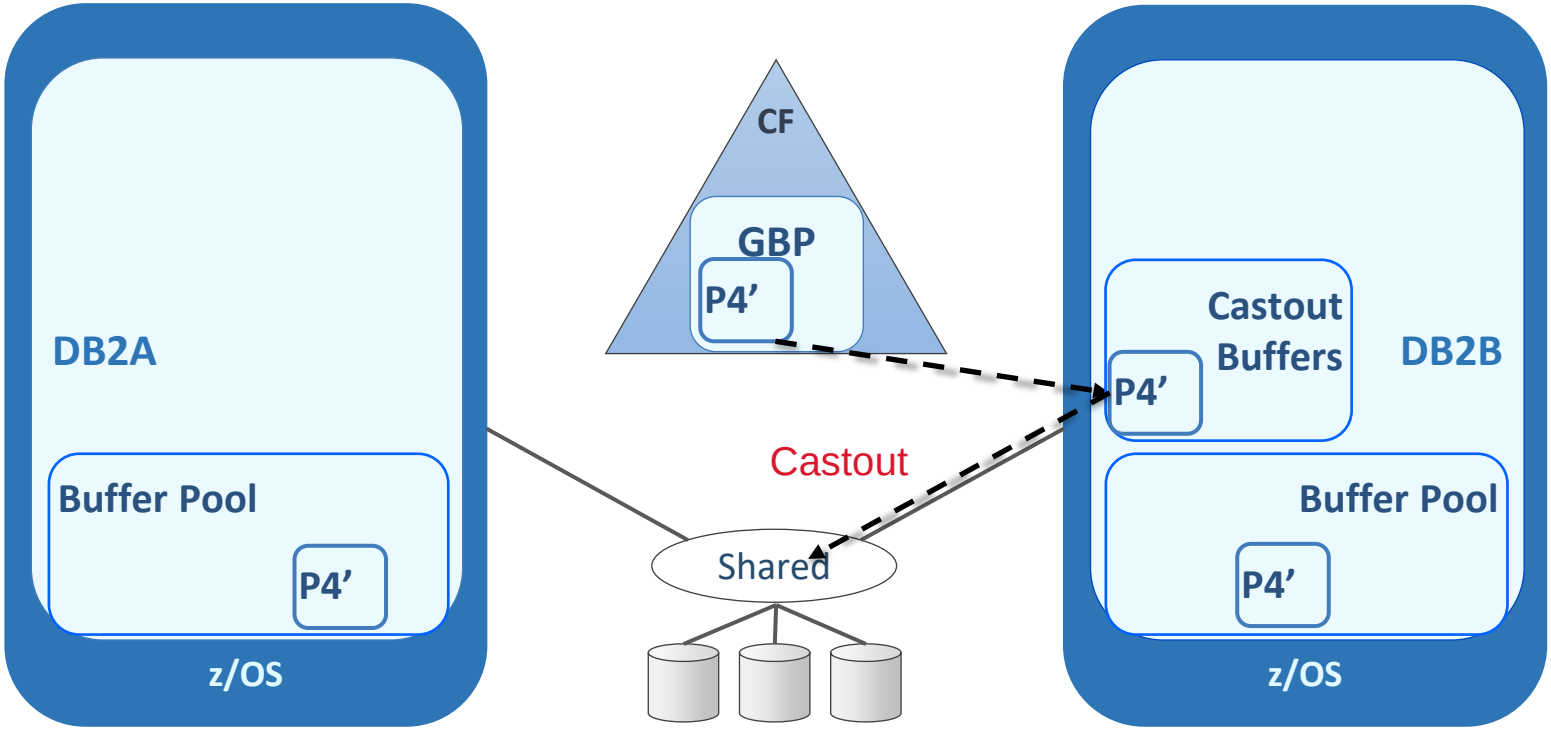

Data sharing performance variables

- CPU cost of data sharing varies based on:
  - Hardware configuration
    - # of CF processors, # of CF links, speed of CF processors, etc.
    - E.g. slow CF increases CPU overhead on host LPAR
  - Lock contention rates
  - CF access intensity for locking and caching
    - Percentage of CPU time in Db2
    - Degree of read/write sharing
    - # of locks obtained
    - Access rate to shared data
    - Insert/delete intensity
    - Release of Db2

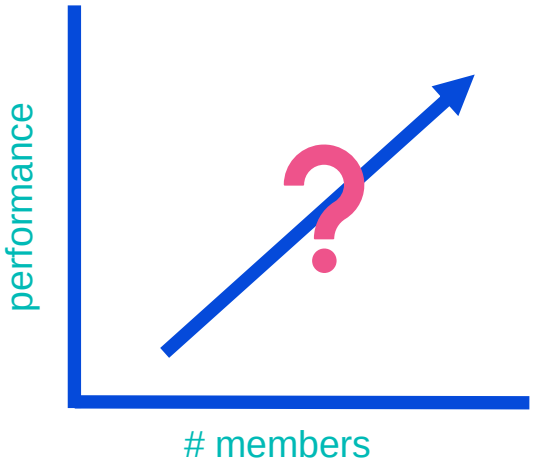

Data sharing performance

Data sharing cost varies by workload

- 'Typical' 2-way data sharing overhead about 10%
  - Important: this is relative to the overall workload, not to an address space
- Individual jobs/transactions may have higher overhead
- < 0.5% added cost per member for 3rd through nth member

Balanced configuration important for best performance

- CF technology = host LPAR CEC technology
- At least 2 CF links per host/CF
- At least 2 dedicated CF engines per CF LPAR for production

Data sharing performance: common issues

Which team is responsible?
Performance?
z/OS?
Db2?
Development?

Under-configured CF

- Too few CF engines
- Not enough memory
- CF structures too small
- CF link contention

Unbalanced host and CF technology

- CF processors slower than host processors

Heavy update plus slow GBP castout

- Applications wait if GBP writes cannot complete

High rate of physical close and open for data set

- In and out of GBP-dependency
- Page-set P-lock negotiation

Data sharing performance: under configured CF typical concerns

Too few CF engines

- High CF CPU busy can increase host CPU overhead
  - Monitor your synchronous lock request service times and conversion from sync to async

LOCK1 too small

- Possible increased false contention
  - Increased IRLM overhead
- Possible insufficient record list entries (RLEs)
  - IRLM slow downs

GBPs too small

- Possible write failures for changed pages
  - Application delays
  - Pages put on LPL
- Possible cross-invalidations due to directory reclaims
  - Overhead
  - Unnecessary I/Os

Monitoring tools

- RMF CF activity report
- Db2 statistics reports
- Db2 –DIS GBPOOL

Data sharing performance: common issues (1)

- Under-configured CF
  - High CF CP busy can increase host CPU overhead
    - Elongated synchronous service time
    - 2 dedicated CF engines, at least, preferred for production
  - 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 cross invalidations (XIs) due to directory reclaims: overhead, extra I/O
    - GBPs – write failures due to lack of storage: slow downs, pages on logical page list (LPL)

Data sharing performance: common issues (2)

- Unbalanced Host-CF technology
  - E.g. Host LPAR is on z14, but CF LPAR is on zEC12 or z13
  - High host CPU for synchronous CF requests
  - Synchronous to asynchronous conversion possible:
    - Increased lock service time
    - Increased elapsed time
    - Decreased throughput
  - Note: if CF is newer technology than host CPU, no problem

Data sharing performance: common issues (3)

- Heavy update plus slow castout can result in structure full warnings and delays
  - Set VDWQT to write to GBP steadily, rather than many pages at commit
  - Set GBPOOLT, CLASST to keep castout steady
- High rate of data set physical close and open
  - Cost of P-lock negotiation from non-shared to shared
  - Table space or index space CLOSE setting
    - CLOSE YES – physical close after pseudo-close intervals, P-lock released
    - CLOSE NO – P-lock not released after pseudo-close intervals
  - DSNZPARM: PCLOSET sets pseudo-close interval
- Long running units of recovery (UoRs) reduce lock avoidance
  - Less impact with Db2 12 lock avoidance algorithms

Common issues

Availability

Design goals: availability

- Parallel components
  - No single points of failure (SPOF)
- Continuous availability
  - Tolerate outage of any single element
  - z/OS fault tolerance
  - Failure isolation
  - Recovery automation
    - ***Responsibilities for planning, testing, updating***

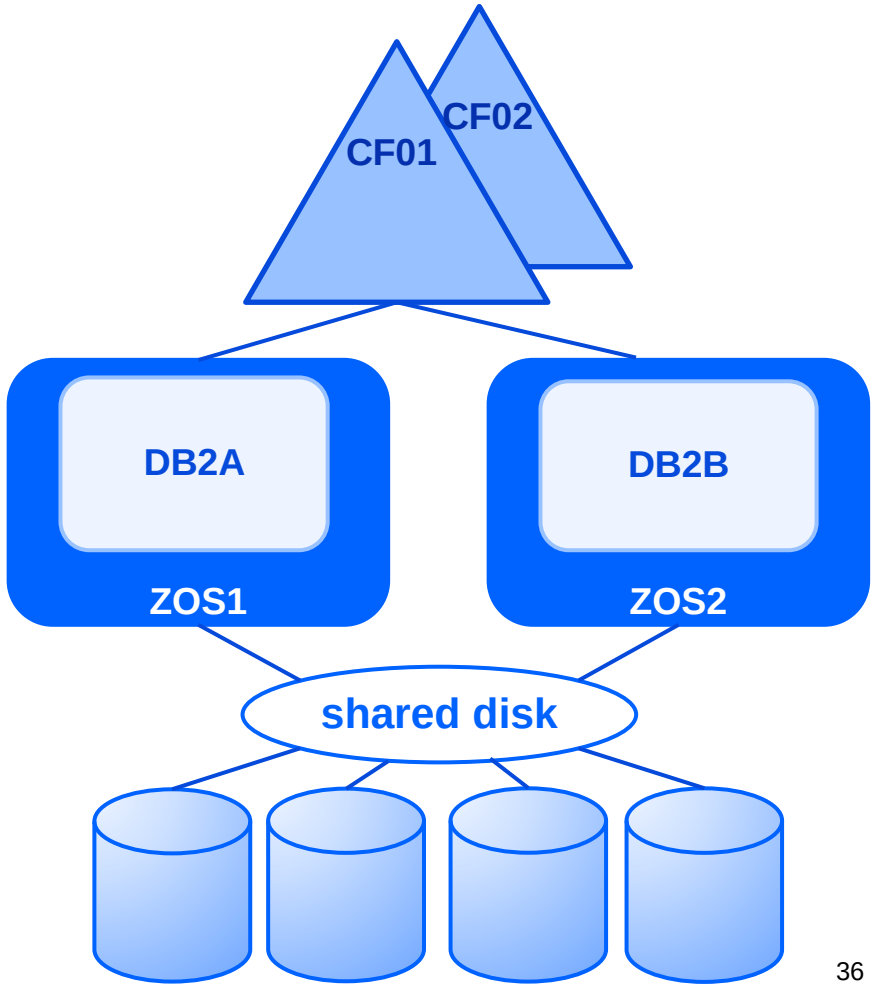

Parallel Sysplex resources

- Network connects clients, employees, partners, and suppliers to a pool of resources

Clients, employees, partners, suppliers

Applications cloned, route around failures

Transaction managers, queue managers, DDF, application servers, batch, TSO, … all on more than one LPAR

Db2 data sharing group

Parallel Sysplex
z/OS LPARs, coupling facilities, shared disk

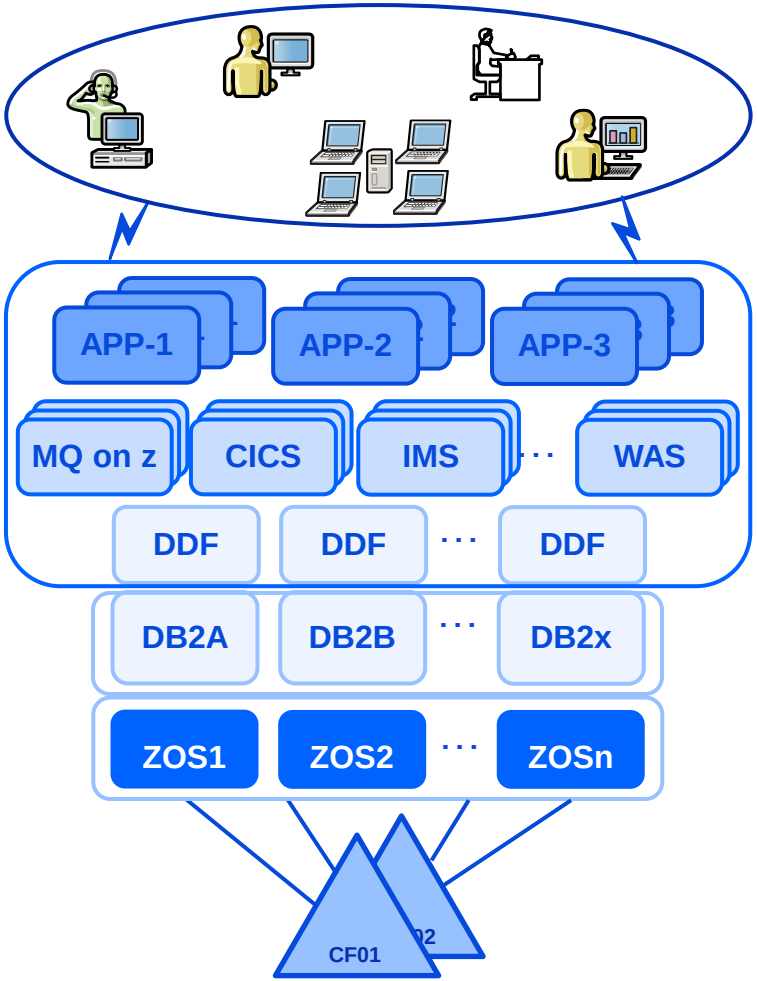

Db2 member: planned outage

- -DB2A STOP DB2 MODE(QUIESCE)
- Sysplex-enabled applications route around outages to resources of other Db2 members
- Rolling maintenance or release migration
  - Point DB2A to new libraries, then start DB2A
  - Then stop DB2B
  - Point DB2B to new libraries
  - Start DB2B
  - Repeat for other Db2 members
- Also use for rolling IPLs, hardware changes

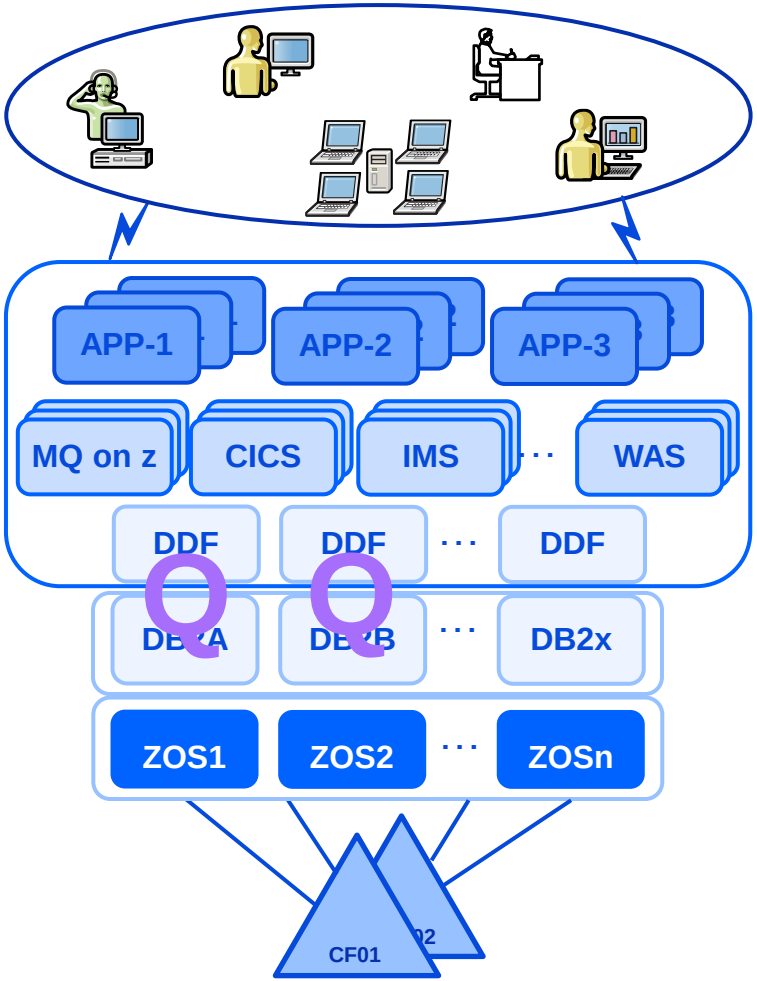

Db2 member: unplanned outage

- In the case of Db2 or IRLM failure
- Sysplex-enabled applications route around outages to resources of other Db2 members
- Retained locks possible
  - Tracked in LOCK1 RLEs
  - May affect multiple users or applications
- Restart Db2
  - DB2A restart required to release DB2A's retained locks
  - Successful restart clears retained locks
- Automation strongly encouraged

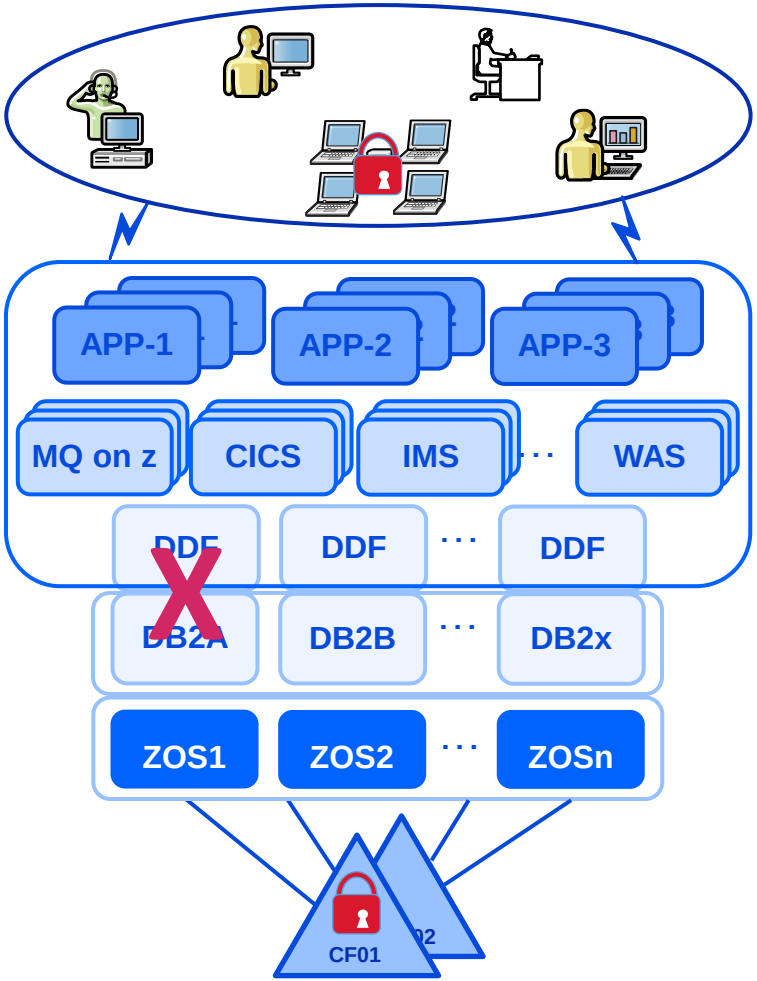

z/OS LPAR: unplanned outage

- In the case of an LPAR failure
  - Db2 impact (and other subsystems)
    - Retained locks possible
- Sysplex-enabled applications route around outages to Db2 and other subsystems
- Restart Db2 on another LPAR
  - DB2A restart required to release DB2A's retained locks
  - Successful restart clears retained locks
- Automate restart of DB2A
  - Possible to use –START LIGHT
  - ***Do not wait for IPL***

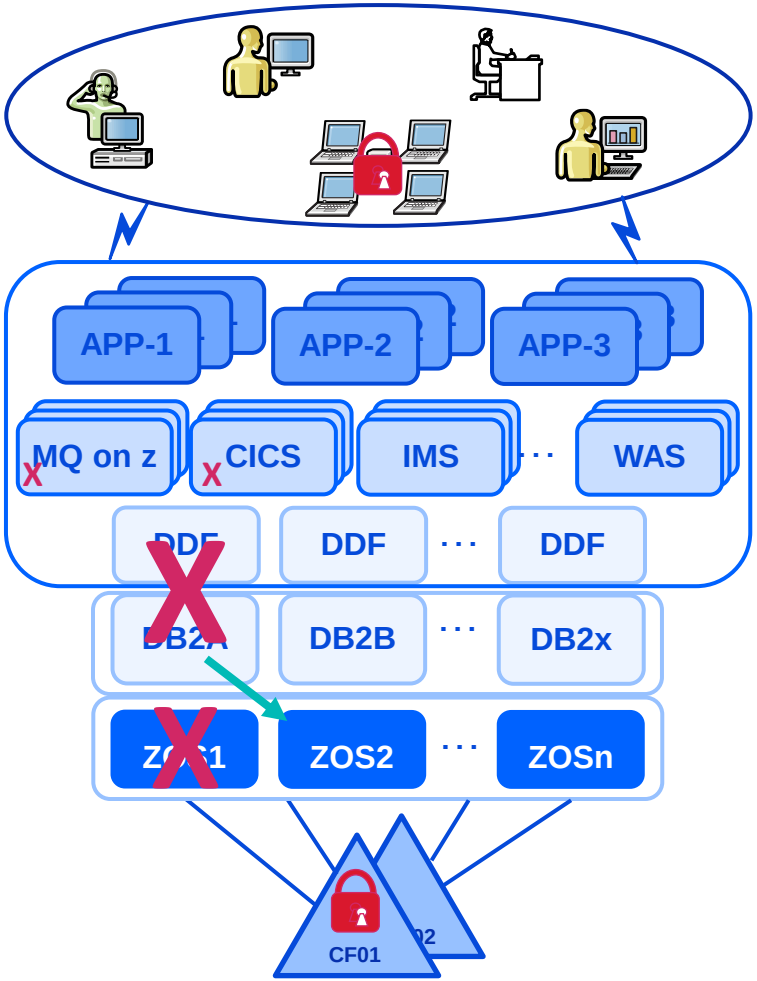

Waiting for the IPL

Coupling facility outages

- Planned CF outages
  - SETXCF START,MAINTMODE, CFNM=*cf_name*
  - SETXCF START,REALLOCATE
  - SETXCF STOP,MAINTMODE, CFNM=*cf_name*
  - SETXCF START,REALLOCATE
  - No outage to data sharing group
  - Note: simplified syntax
    - check IBM Docs for current syntax
- Unplanned CF outages
  - Automatic REBUILD of SCA, LOCK1 in other CF
    - Db2s, IRLMs required
  - Duplexed GBPs revert to simplex mode
  - No outage to data sharing group

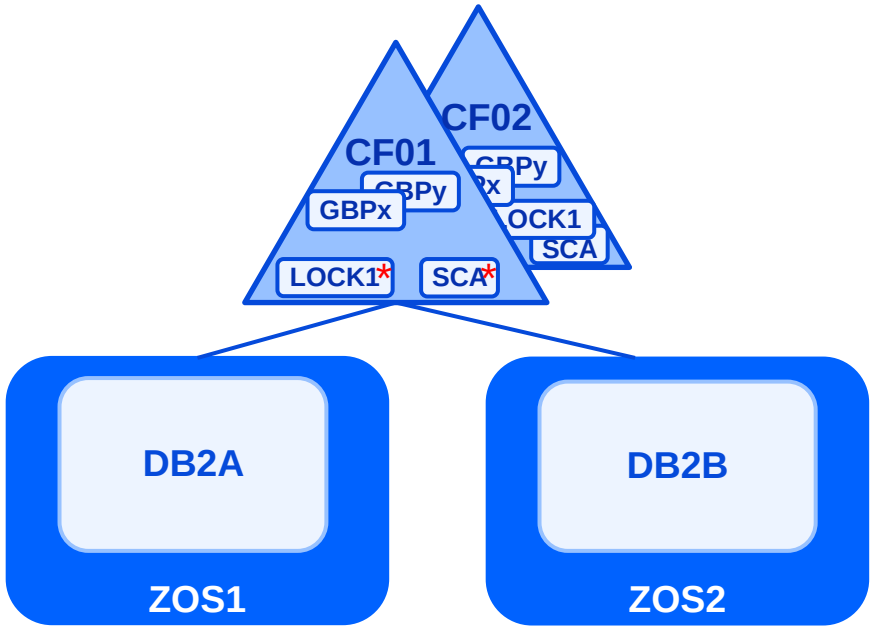

Data sharing availability: common issues

Which team is responsible?
Performance?
z/OS?
Db2?
Development?

Restart processes

- Reliance on manual restart decision
  - (e.g. ‘emergency bridge call’)
- Cross-system Db2 restart not automated in case of LPAR failure

Single points of failure

- 2-CEC, 2-ICF configuration without system managed duplexing (SMD)
  - Synchronous or asynchronous
- Applications not cloned, cannot route around failures

Long-running units of recovery (UoRs)

- Long UoRs delay service restoration for affected applications during outage recovery

Invest in:

- Automation
- Deploying SMD
- Sysplex-enabling applications
- Identifying and shortening units of recovery
  - Application focus
  - Subsystem parameters:
    - URCHKTH
    - URLGWTH

Data sharing availability: common issues

- Reliance on manual restart processes (‘emergency bridge call’)
  - Applies to same-system restart or cross-system restart
    - Most customers 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 Sysplex without operator interaction
  - Manual processes or operation responding to system prompts contribute to delays in completing Db2 restart

Data sharing availability: common issues

- 2-CEC / 2-ICF configuration without system managed duplexing
- Risk of ‘double failure’ if CEC outage
  - LOCK1 and an IRLM, or SCA and a MSTR
    - If all IRLMs survive, they can rebuild LOCK1
    - If all MSTRs survive, they can rebuild SCA
  - If LOCK1 *and* one IRLM affected by outage, Db2 data sharing group fails
  - If SCA *and* one MSTR affected by outage, Db2 data sharing group fails
  - In diagram at right, outage of one CEC has 50% chance of data sharing group outage
  - Some sites accidentally have LOCK1 in ICF01 and SCA in ICF02
    - Outage of one CEC has 100% chance of data sharing group outage!

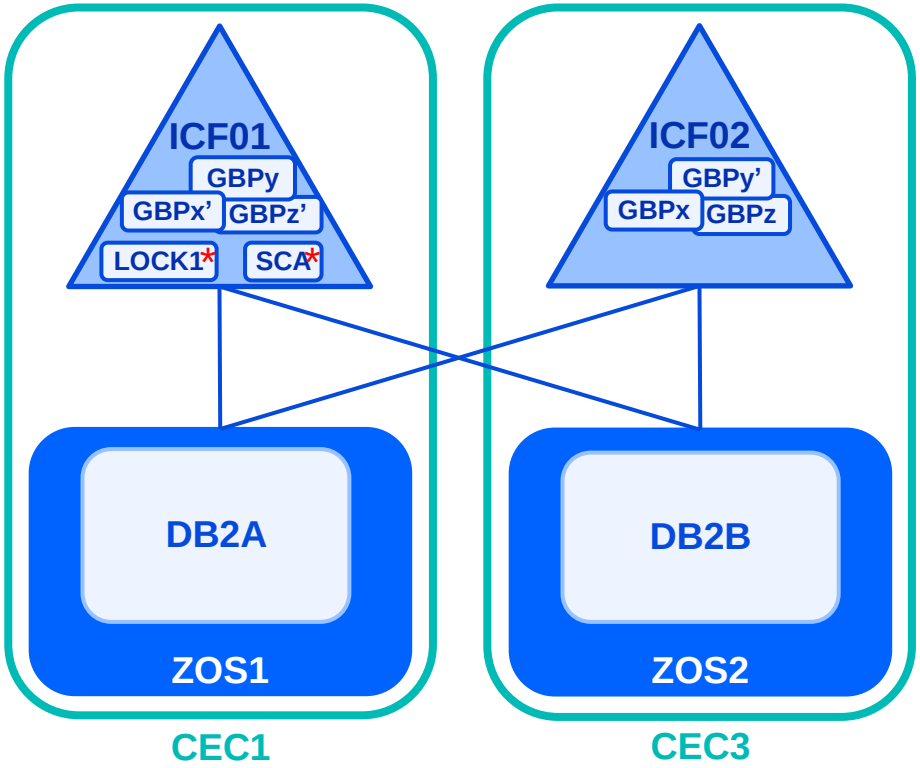

Solution: system managed duplexing (SMD) either asynchronous or synchronous

Data sharing availability: common issues

- 2-CEC / 2-ICF configuration without system managed duplexing
- Risk of ‘double failure’ if CEC outage
  - LOCK1 and an IRLM, or SCA and a MSTR
    - If all IRLMs survive, they can rebuild LOCK1
    - If all MSTRs survive, they can rebuild SCA
  - If LOCK1 *and* one IRLM affected by outage, Db2 data sharing group fails
  - If SCA *and* one MSTR affected by outage, Db2 data sharing group fails
  - In diagram at right, outage of one CEC has 50% chance of data sharing group outage
  - Some sites accidentally have LOCK1 in ICF01 and SCA in ICF02
    - Outage of one CEC has 100% chance of data sharing group outage!

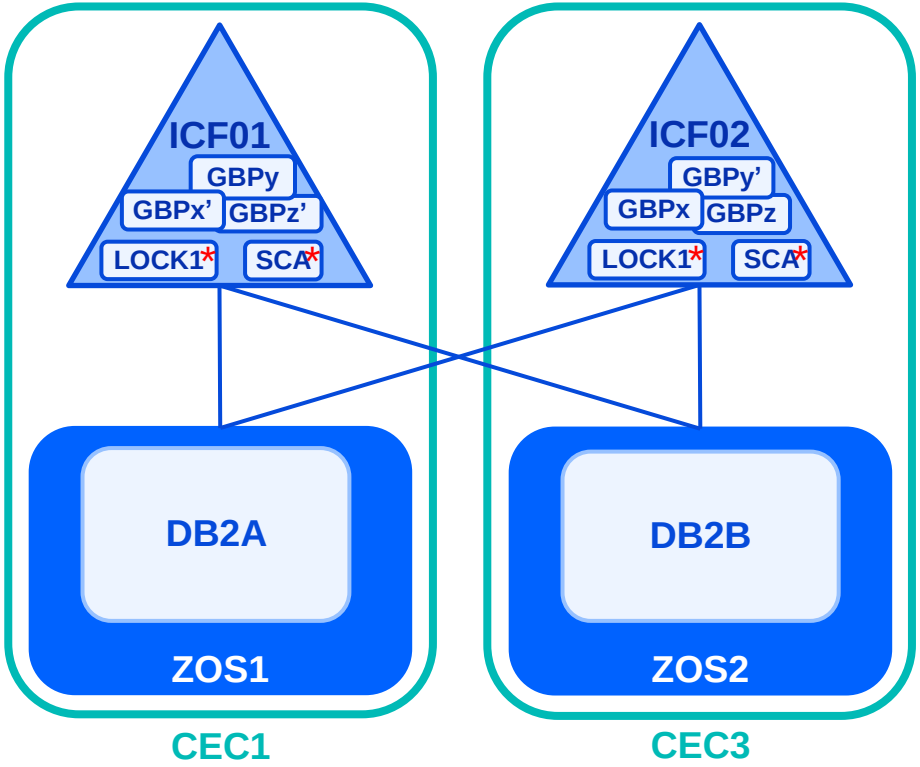

Solution: system managed duplexing (SMD) either asynchronous or synchronous

Data sharing availability: common issues

- Not implementing GBP duplexing
  - If CF outage affects GBPs, Db2 stays active
  - But applications may not be able to access data
  - Application impact could be many minutes of no access to data
- Not cloning applications and routing around failures
  - Applications should be able to execute anywhere in the Parallel Sysplex
  - CICS, IMS, WAS, MQ, DDF all have facilities to route around failures
- Long running units of recovery (UoRs) that impact service restoration
  - Programs that access critical data must commit frequently
  - If a failure requires lengthy backout, other applications may not be able to access data

Common issues

Workload balancing

Parallel Sysplex: single system image

- Single DDF location for distributed access
  - Location alias to access a subset of the group
- Local attach using group attachment name
  - DB2P in the example environment
  - Subgroup attach to limit local attach to certain members
- Group scope:
  - Display output
  - Performance monitors
  - DDL, BIND, utilities
  - Authorization
  - WLM application environments
    - &IWMSSNM symbolic allows same WLM application environment to execute on any member of the data sharing group
  - Log merging for replication (IFCID 306)

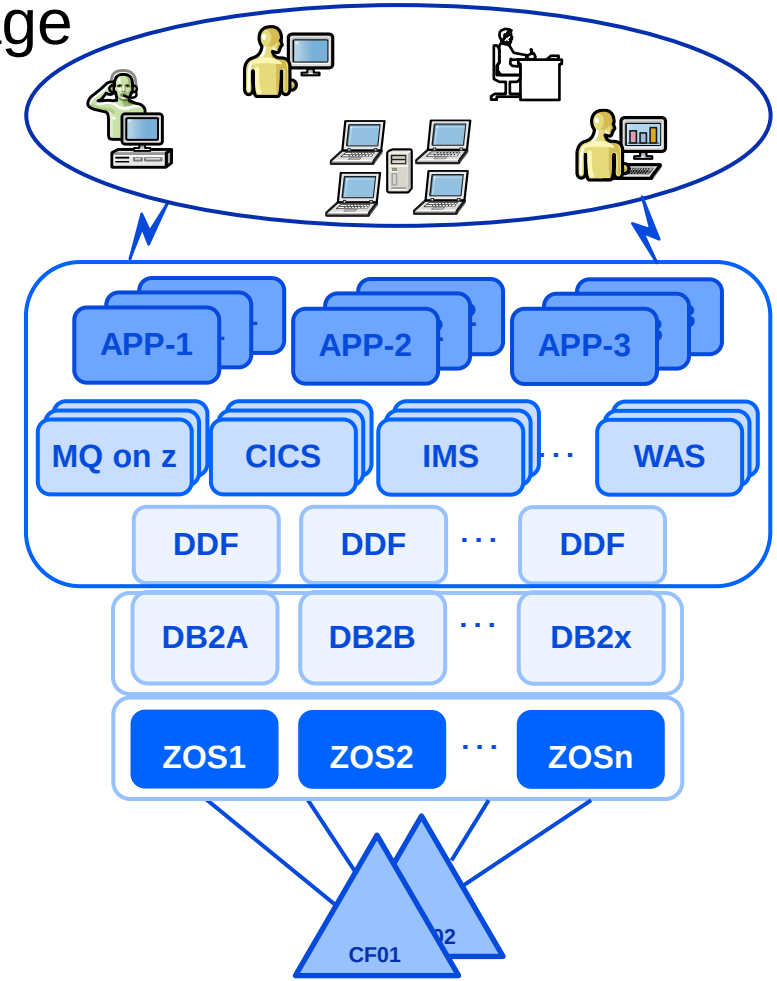

Dynamic workload balancing

- Workload balancing is primarily about availability and application access
- Dynamic reaction to workload changes and available resources

z/OS CommServer: Sysplex Distributor

MQ Shared Queues

CICSPlex System Manager (CPSM) routes workload between CICS regions

IMS Transaction Manager (IMS/TM): shared message queues, batch message processors (BMPs)

Websphere: WAS or WAS Liberty on z/OS

Db2 Distributed Data Facility (DDF) using dynamic virtual IP addressing (DVIPA)

Workload manager (WLM) provides input to routing choices

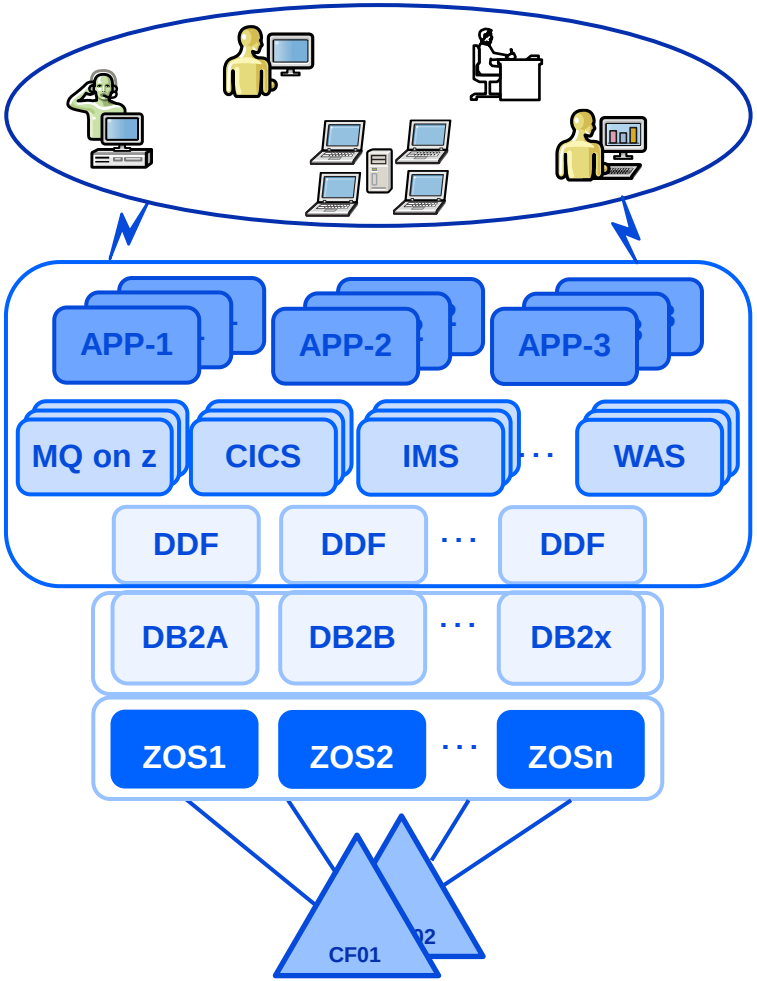

Websphere and sysplex workload balancing: example

- IBM Data Server Driver: JDBC type 4
  - Data server driver ‘transports’ = Db2 for z/OS ‘connections’
    - Db2 CONDBAT parameter
- Enable sysplex workload balancing: application threads can switch between Db2 members on a transaction boundary
  - Availability
  - Performance
  - Balancing

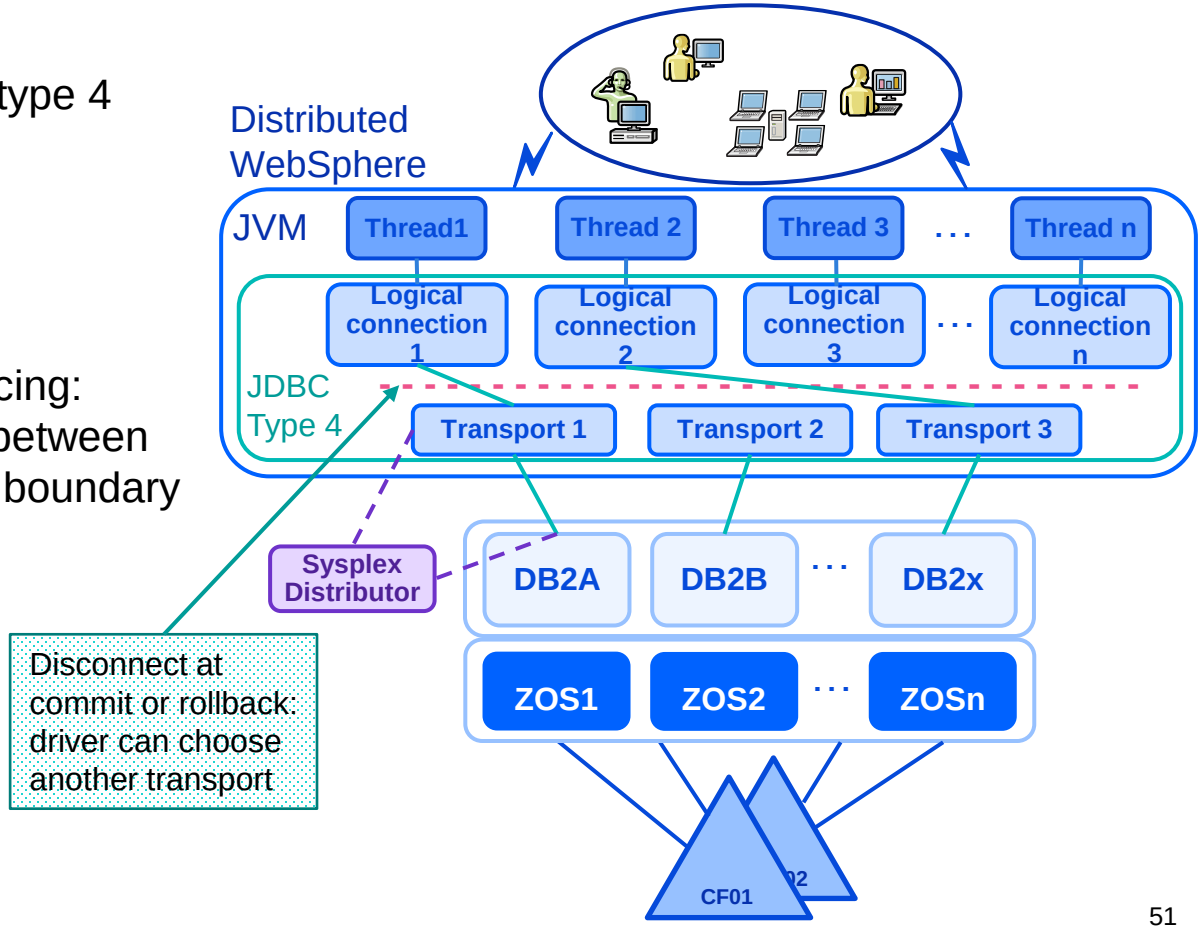

Websphere and sysplex workload balancing: example

- IBM Data Server Driver: JDBC type 4
  - Data server driver ‘transports’ = Db2 for z/OS ‘connections’
    - Db2 CONDBAT parameter
- Enable sysplex workload balancing: application threads can switch between Db2 members on a transaction boundary
  - Availability
  - Performance
  - Balancing
- Over time, connections to all members

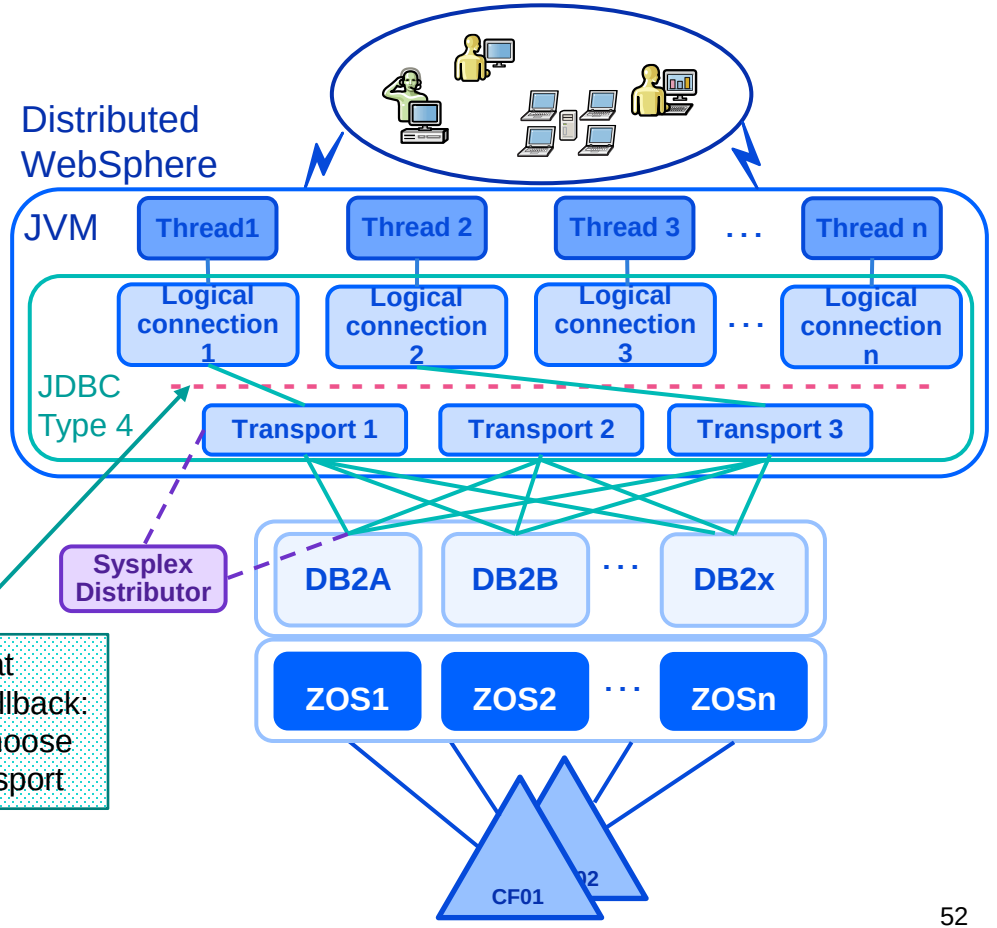

Workload balancing: common issues

- Distributed access getting stuck on one member
  - Db2 10 introduced new system parameters (DSNZPARM)
    - MAXCONQW time waiting for active data base access thread (DBAT)
    - MAXCONQN queue depth waiting for active DBAT
  - Db2 will close connections if either value exceeded
    - Client will connect to another Db2 member
- Runaway requesters consume all DBATs (MAXDBAT) or connections (CONDBAT)
  - ‘Denial of service attack’
  - Db2 10 introduced system profile monitoring to manage connections or threads
    - Variety of filters or qualifiers
    - Lower granularity than system parameters CONDBAT or MAXDBAT

Workload balancing: common issues

DDF applications accessing single Db2 member, rather than any member of the group

DDF applications not using Sysplex workload balancing

Most common issue: not enabling applications and subsystems for Parallel Sysplex

Applications depending upon system resilience rather than using available resources in the Sysplex

- Hard to manage planned outages
- Workload spikes can lead to performance problems when plenty of capacity still available

Workload balancing: common issues

- Some think ***workload balancing*** means even distribution
  - Not intended for even split (50-50, or 33-33-33, etc.)
  - Intended for availability
  - “Sysplex workload balancing” based on numerous inputs to WLM
    - WLM uses Db2 input to indicate ‘health’ of Db2 member
- SW charge ‘penalty box’
  - Limit certain workloads to one environment
    - Manage / limit SW charges
  - Lose availability benefits

Recent developments in Db2 data sharing

Continuous delivery in Db2 12 and new capabilities in Db2 13

Db2 12 continuous delivery for data sharing

- IRLM deadlock process improvements
- Retry of automatic recovery for logical page list (LPL) and GBP recovery pending (GRECP)
- Asynchronous cross-invalidation of GBPs
  - Especially valuable with distance between CFs and LPARs
- -ALTER GROUPBUFFERPOOL (-ALT GBPOOL)
  - Changes for ratio of directory entries to data elements
    - Limit of RATIO option raised from 255 to 1024
    - Default for RATIO option changed from 5 to 10
- Sysplex support for multi-factor authentication (MFA) and IBM RACF PassTickets

Db2 13 enhancements

- IRLM can dynamically trigger increase of LOCK1 size, producing more RLEs
- GBP residency time
  - Directory entries
  - Data elements
  - Residency time available in IFCIDs, -DIS GBPOOL, SMF
- GBP castout speed – reduce delays associated with GBP full conditions
  - Check more often for GBP full to drive group level castout
  - Retry GBP writes more frequently
- Reduced P-lock contention for PBR table spaces with relative page numbers

Dynamically alter CF lock structure

FL 100

Db2 lock structure (*dsgrpnm*_LOCK1) allocated based upon CFRM policy

– INITSIZE = initial size, SIZE = max size

– Lock entry table = $2^n \leq 0.5$ x INITSIZE

CFRM ALLOWAUTOALT YES

– XES can dynamically adjust structure allocation

  - FULLTHRESHOLD > 0 enables monitoring and dynamic rebuild; ‘full’ applies to modify lock list record list entries (RLEs) only
  - Expansion of LOCK1 results in more RLEs

Under heavy workload IRLM does not request expansion; XES expansion not responsive enough

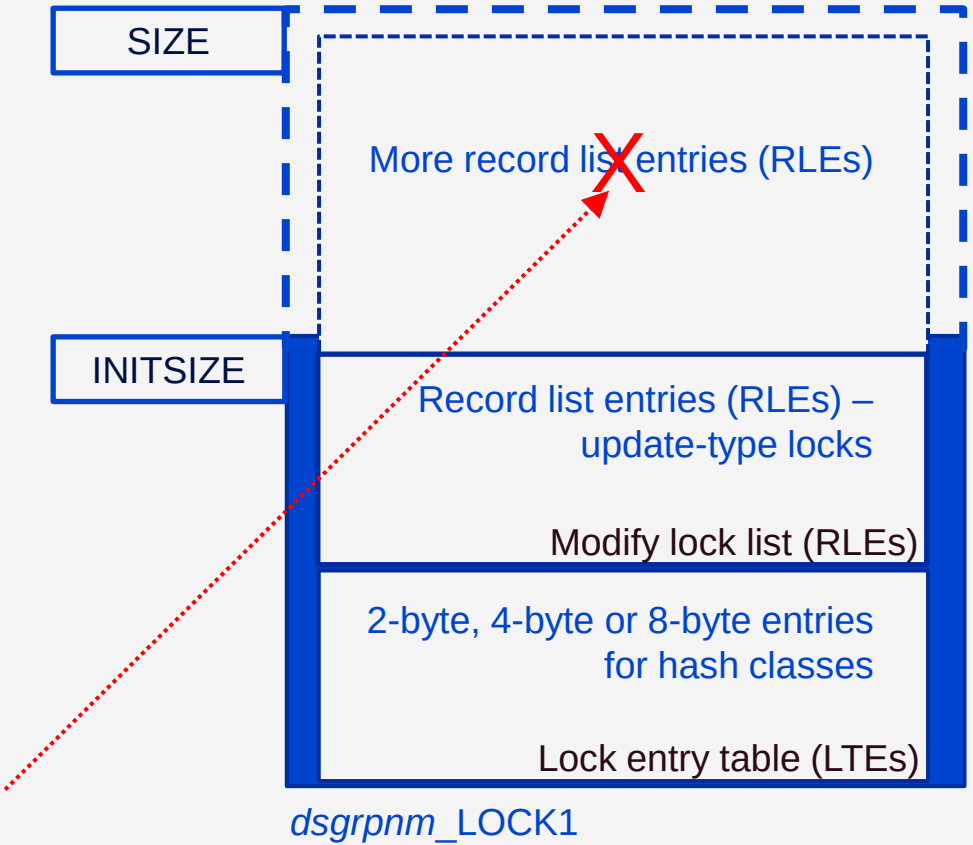

IRLM requests dynamic lock structure increase

FL 100

Db2 13: IRLM initiates dynamic alter of lock structure to avoid structure full conditions

- More responsive than AUTOALTER in CFRM policy
- New messages:

```
DXR189I  <irlmname> ALTERING LOCK STRUCTURE SIZE
DXR190I  <irlmname> ALTER LOCK STUCTURE COMPLETED

IXC530I START ALTER REQUEST FOR STRUCTURE DSNCAT_LOCK1 ACCEPTED
  TARGET  SIZE:                  8 M
IXC534I REQUEST TO ALTER STRUCTURE DSNCAT_LOCK1
COMPLETED.  TARGET ATTAINED.
CURRENT SIZE:                    8 M  TARGET:          8 M
CURRENT ENTRY COUNT:          4427    TARGET:       4427
CURRENT ELEMENT COUNT:           0    TARGET:          0
CURRENT EMC COUNT:               0    TARGET:          0
```

Closing

Summary and additional resources

Db2 data sharing: benefits summary

- Data sharing technology allows Db2 to deliver continuous availability and nearly unlimited scalability into the future
- Db2 data sharing is a proven technology
  - Many sites rely on Db2 data sharing to meet their production availability requirements
  - Parallel Sysplex and data sharing are the norm for enterprise applications and provide critical resources to the Hybrid Cloud
- Db2 development continues to deliver data sharing enhancements

Thank you!

mrader@us.ibm.com

Resources: Publications

- *Data Sharing: Planning and Administration*
  - IBM Documentation: Intro to Db2 data sharing (Db2 12)
  - IBM Documentation: Intro to Db2 data sharing (Db2 13)
- z/OS publications
  - *MVS Setting Up a Sysplex*
  - *z/OS Sysplex Services*
- *PR/SM Planning Guide:* SB10-7162

Resources: Redbooks

- DB*2 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
- *DB2 9 for z/OS: Resource Serialization and Concurrency Control:* SG24-4725-01
- *Parallel Sysplex Configuration Overview:* SG24-6485
- *Exploiting Parallel Sysplex Customer Experience:* SG24-7108
- *Parallel Sysplex Best Practices:* SG24-7817
- *Improving z/OS Application Availability by Managing Planned Outages*, SG24-8178

Appendix

Application implications

Duplexing Db2 CF structures

Application implications

- SQL interface does not change
- 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

Db2 CF structure duplexing: 2 kinds

- “User managed” duplexing
  - GBPs *strongly recommended*
  - “User” = Db2 for z/OS; Db2 is aware of two structures in different CF LPARs
- “System managed” duplexing
  - LOCK1 and/or SCA
  - XES component of z/OS keeps two structures in different CF LPARs synchronized
    - Db2 is not aware of second structure
    - All updates to primary structure, then to secondary
  - Recommended if CFs are on ICFs and co-reside on CEC with Db2 members
    - Avoid ‘double failure’ scenario
    - Up to 4x cost for synchronous lock requests if synchronous system managed duplexing
  - Db2 12: asynchronous system managed duplexing for the Db2 LOCK1 structure greatly reduces the cost of synchronous lock requests

GBP duplexing

- Allocate secondary GBP on alternate CF
- Db2 writes changed pages to both primary and secondary GBP
- If loss of connectivity or loss of structure –
  - Switch to secondary
  - 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

GBP duplexing

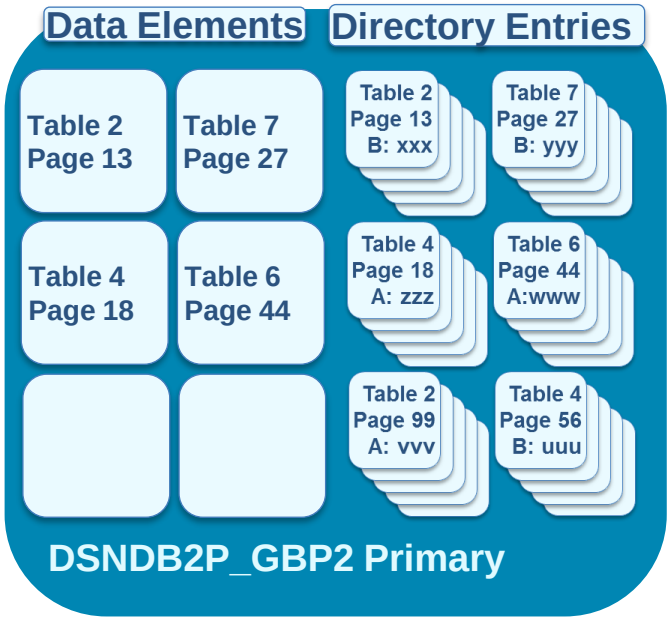

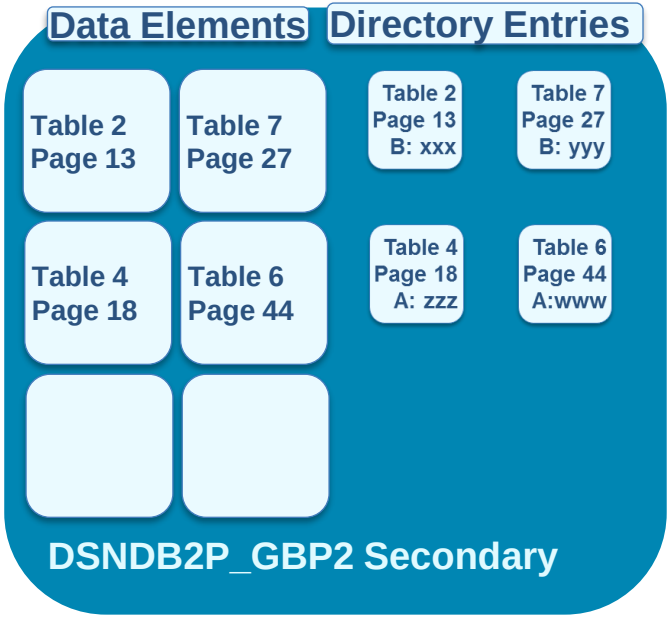

Db2 Version 5 GBP duplexing recovery