Db2 for z/OS Data Sharing: Performance Considerations

Mark Rader mrader@us.ibm.com

IBM Z Worldwide Systems Center (WSC)

May 13th, 2026

Agenda

Quick review:

- Components of Parallel Sysplex and Db2 data sharing
- Db2 data sharing processes

Performance considerations

Performance indicators

- Examples: displays and reports

Suggestions

Questions

Summary

Components of Parallel Sysplex and Db2 for z/OS data sharing

Coupling facility (CF) LPARs – hi-speed shared memory

- External (CF) or integrated (ICF – aka internal CF)
- CF control code (CFCC)
- Structures: lock, list, cache
- Links: external (short or long) or internal (IC – microcode)
- CF resource management (CFRM) policy

z/OS components

- Cross-system extended services (XES)
- Cross-system coupling facility (XCF)

Db2 components

- CF structures: lock (*dsgrpnm*_**LOCK1**), list (*dsgrpnm*_**SCA**), cache (*dsgrpnm*_**GBPx**)
- Shared data – 1 Db2 Catalog/Directory, user data
- Read access to all logs

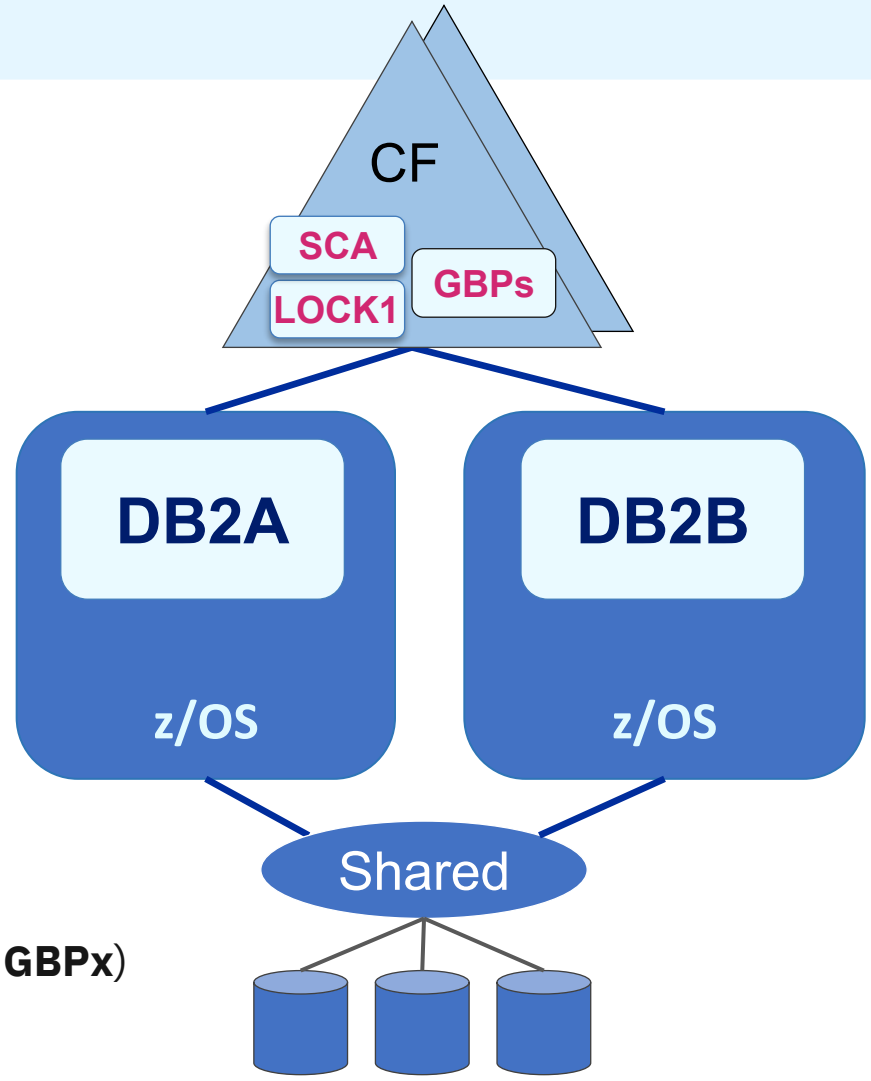

Db2 data sharing processes (1|4)

Global locking – IRLM and LOCK1

- P-locks determine inter-Db2 read/write interest
  - Coherency: resource held in 2 members, one member wants to change
  - Negotiable
- L-locks manage concurrency
  - Serialization: many readers OR one writer

Group-wide status – MSTR and SCA

- Not a significant performance concern if defined well

Buffer coherency – DBM1 and GBPs

- GBPs have directory entries to track inter-Db2 interest in buffer pages
- GBPs have data elements to cache changed pages
  - Caching depends on GBP and page set definition
  - GBP definition (GBPCACHE = YES|NO)
  - Page set definition (GBPCACHE = ALL|CHANGED|SYSTEM|NONE)

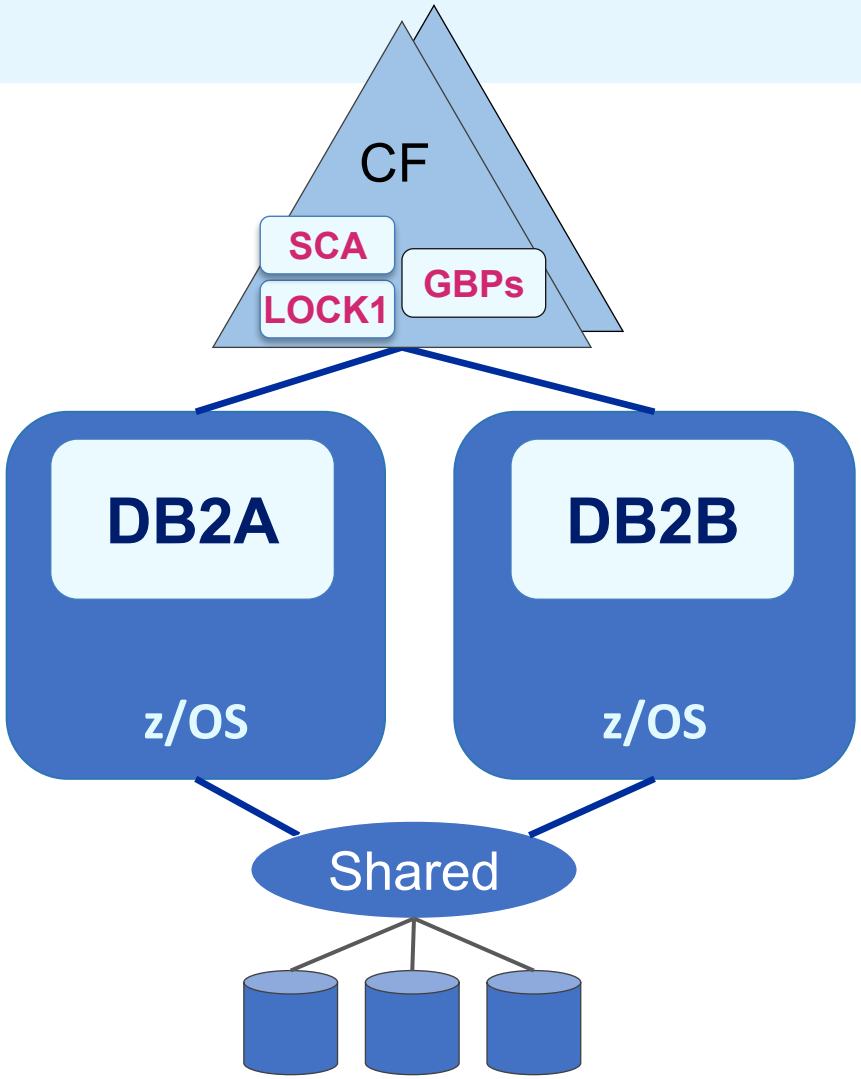

Db2 data sharing processes (2|4)

Global locking: No contention example

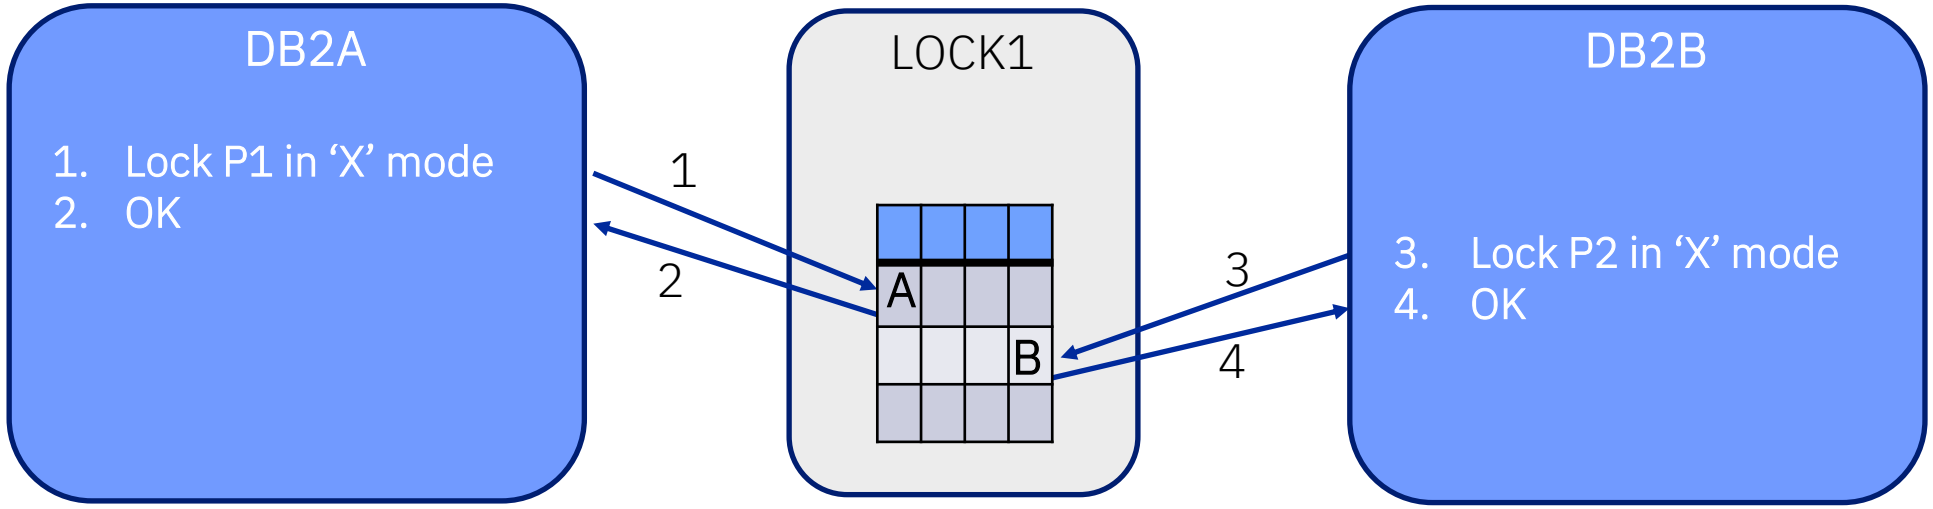

CF messages occur in microseconds

Db2 data sharing processes (3|4)

Global locking: contention example

– Contention results in suspend time for one of the members

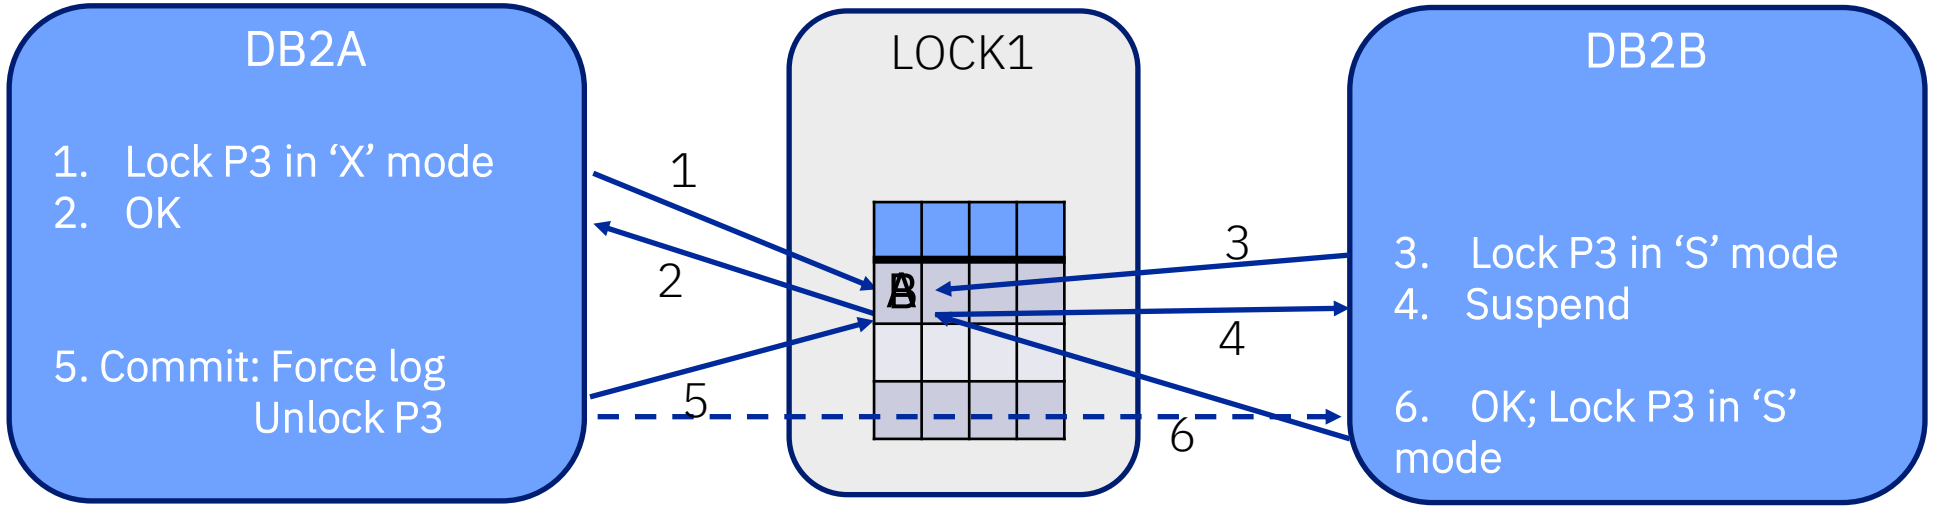

CF messages occur in microseconds

Db2 data sharing processes (4|4)

Buffer coherency and GBP access

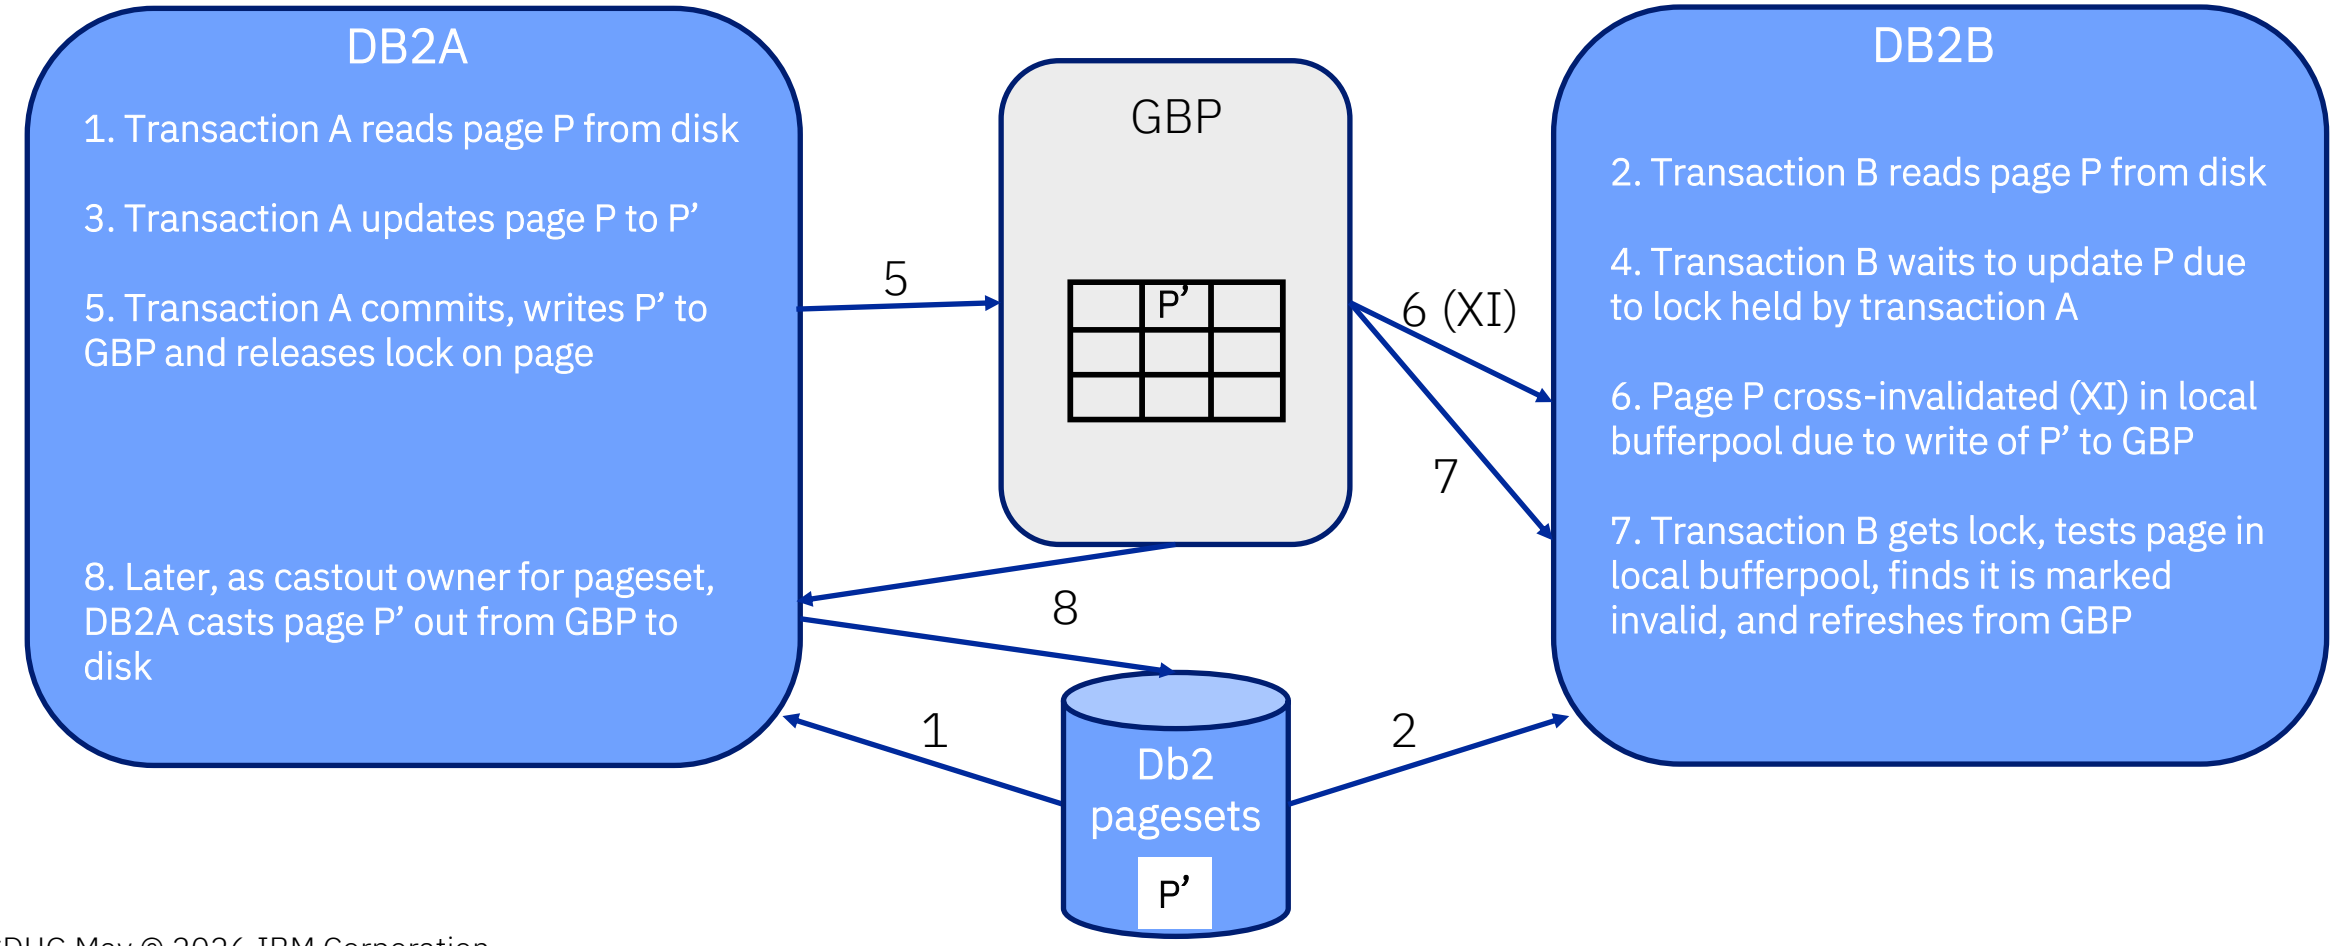

Data sharing: critical factors

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
  - One report: > 1,000,000 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

Synchronous messages

Most CF messages for Db2 or IRLM are synchronous

- Host CPU cost for duration of round trip to CF -> CPU busy

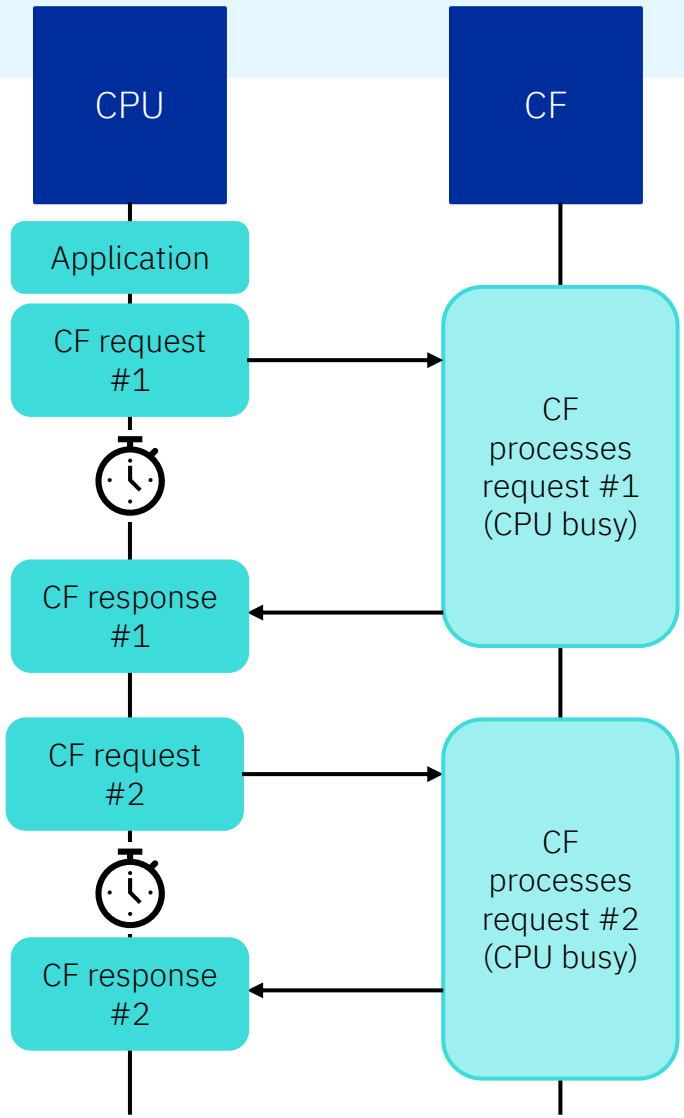

Performance considerations

Performance considerations (1|2)

Db2 CF requests are synchronous relative to the:

- Application for which the CF request is made
- Host processor (GCP or zIIP) that initiated the request
  - Busy for duration of request

Data sharing performance variables

- Hardware configuration
- 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

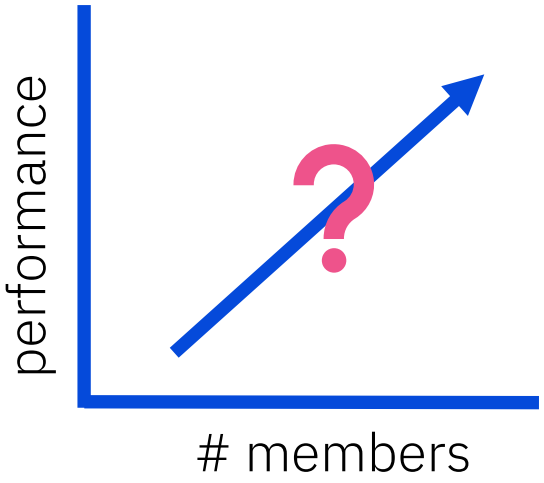

Performance considerations (2|2)

CF link capacity

CF processor capacity

- If CF processors too busy, longer CF request service times and busier host processors

Db2 CF structure allocation and definition

- LOCK1
  - Lock table large enough to minimize false contention
  - Enough record list entries to track change activity
- GBPs
  - Enough directory entries to track the aggregate buffer pools
  - Enough data elements to accommodate high write activity

Distance between systems

- CF requests to remote CFs require about 10 µsec / km round trip

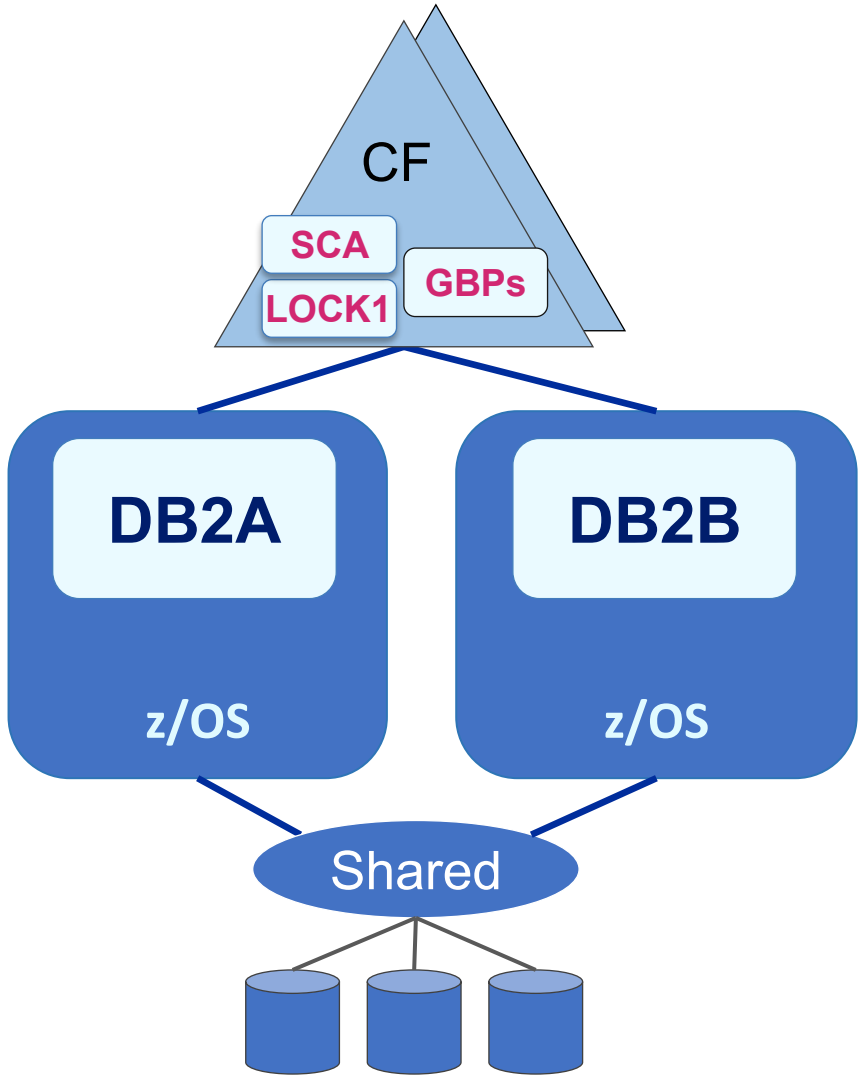

Performance indicators

CF performance indicators

Coupling facility displays and reports
- RMF Monitor III
- RMF CF activity reports

CF request rates by structure

CF processor busy

Structure detail
- Synchronous service times by LPAR
- Asynchronous conversion
  - XES detects service times above a certain level and converts requests to async
  - Frees host CP to perform other tasks
- Delayed requests
- Contention and false contention

CF channel activity (link traffic)

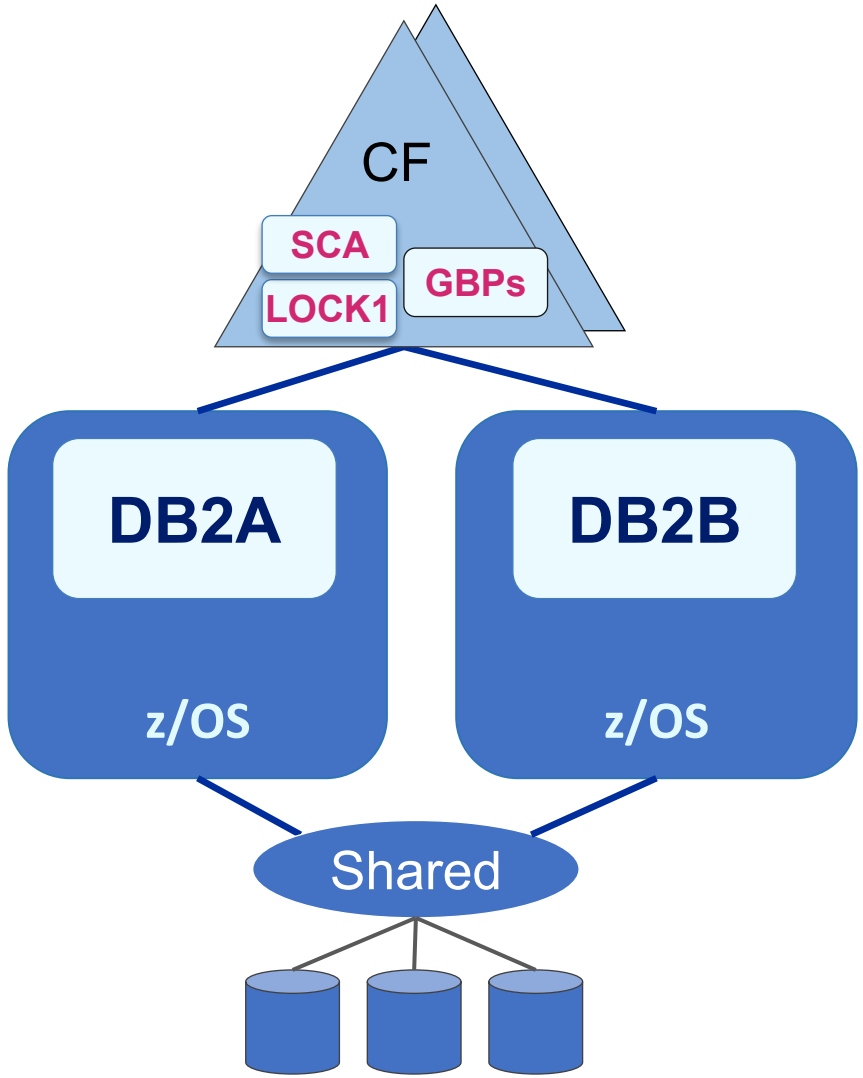

Synchronous Request

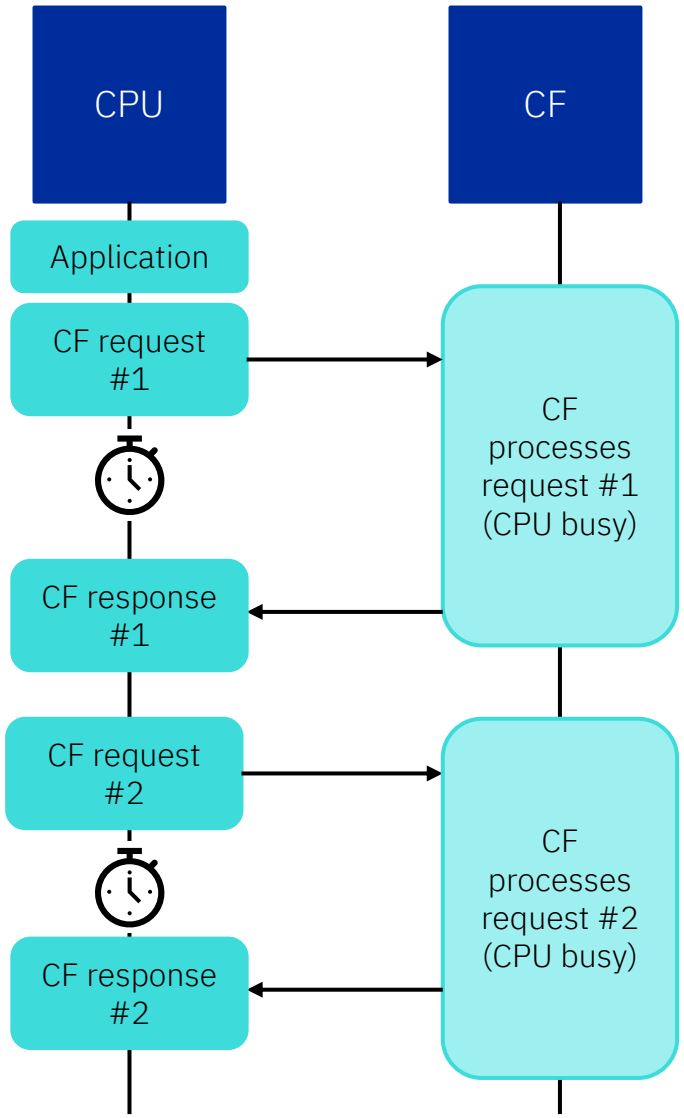

Asynchronous Request

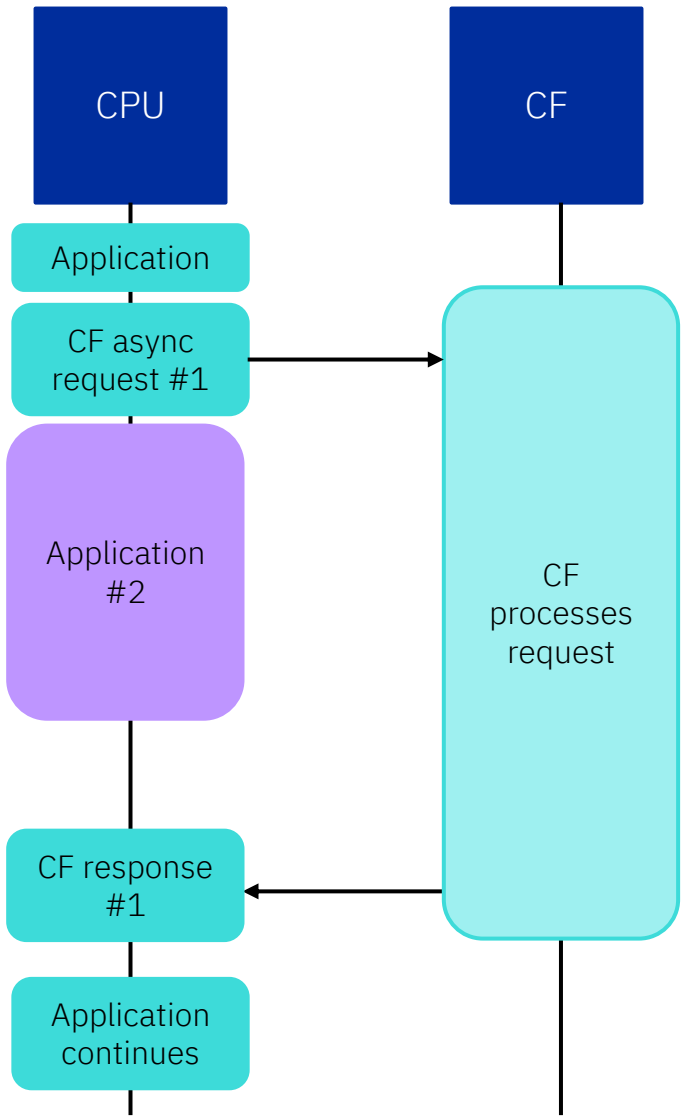

CF activity report: structure summary

Sample Db2 structures

TYPE	STRUCTURE NAME	STATUS	CHG	ENC	ALLOC SIZE	% OF CF STOR	# REQ	% OF ALL REQ	% OF CF UTIL	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC/ DIR REC XI'S
LIST	CTMPRDS_CONTROLM	ACTIVE		NO	49M	0.2	14425	0.1	0.1	8.01	16K 6	31K 16	N/A N/A	N/A N/A
	CTMPRDS_CTMCMEM	ACTIVE		NO	50M	0.2	634	0.0	0.0	0.35	47K 26K	47K 26K	N/A N/A	N/A N/A
	CTMPRDS_XAELIST	ACTIVE		NO	49M	0.2	0	0.0	0.0	0.00	40K 0	79K 0	N/A N/A	N/A N/A
	DSNDBPG_SCA	ACTIVE		NO	99M	0.3	135205	0.5	12.2	75.11	99K 823	198K 3038	N/A N/A	N/A N/A
LOCK	CTMPRDS_XAELOCK	ACTIVE		N/A	49M	0.2	0	0.0	0.0	0.00	0 0	0 0	N/A 0	N/A N/A
	DSNDBPG_LOCK1	ACTIVE		N/A	147M	0.5	9371K	36.8	11.0	5206.0	125K 3824	0 0	34M 48K	N/A N/A
	DSNDBPG_GBP2	ACTIVE PRIM		NO	2G	5.9	5516K	21.7	32.6	3064.3	3310K 627K	158K 157K	N/A N/A	0 0

CF activity report: summary section notes

requests over the RMF interval

Average requests per second

Values at right are millions (M) or thousands (K)

- M^=1024* 1024 ; K^=1024

For GBPs:

- Directory reclaims on upper line
- Directory reclaims resulting in cross-invalidation on lower line
  - Non-zero values to be avoided
  - Can cause ‘unnecessary’ I/O activity
  - Details available with –DIS GBPOOL GDETAIL or in SMF 100

CF activity report: processor summary

Average CF utilization can be a key factor

Production CFs engines should be dedicated to the CF LPAR

```
PROCESSOR SUMMARY
-------------------------------------------------------------------------------------------------------------------

COUPLING FACILITY         8562        MODEL T02      CFLEVEL  24           DYNDISP OFF

AVERAGE CF UTILIZATION  (% BUSY)            2.6      LOGICAL PROCESSORS:   DEFINED   3    EFFECTIVE  3.000
                                                                           SHARED    0    AVG WEIGHT   0.0
```

CF activity report: structure activity

Db2 lock structure

```
STRUCTURE NAME = DSNDBPG_LOCK1     TYPE = LOCK    STATUS = ACTIVE                    ENCRYPTED = N/A
          # REQ    -------------- REQUESTS -------------   -------------- DELAYED REQUESTS -------------
SYSTEM    TOTAL                    #      % OF  -SERV TIME(MIC)-    REASON     #     % OF  ---- AVG TIME(MIC) -----  EXTERNAL REQUEST
NAME      AVG/SEC                  REQ     ALL     AVG    STD_DEV             REQ     REQ   /DEL     STD_DEV    /ALL  CONTENTIONS
zOS1       8816K   SYNC    8816K   94.1     2.2      0.8     NO SCH     0     0.0     0.0         0.0     0.0   REQ TOTAL      9698K
           4898    ASYNC       0    0.0     0.0      0.0     PR WT      0     0.0     0.0         0.0     0.0   REQ DEFERRED     10K
                   CHNGD       0    0.0  INCLUDED IN ASYNC  PR CMP     0     0.0     0.0         0.0     0.0   -CONT            10K
                   SUPPR       0    0.0                                                                       -FALSE CONT    5729
zOS2        183K   SYNC     183K    1.9     2.6      1.1     NO SCH     0     0.0     0.0         0.0     0.0   REQ TOTAL       221K
          101.4    ASYNC       0    0.0     0.0      0.0     PR WT      0     0.0     0.0         0.0     0.0   REQ DEFERRED   3698
                   CHNGD       0    0.0  INCLUDED IN ASYNC  PR CMP     0     0.0     0.0         0.0     0.0   -CONT          3697
                   SUPPR       0    0.0                                                                       -FALSE CONT     955
zOS3        132K   SYNC     115K    1.2    26.9      3.9     NO SCH     0     0.0     0.0         0.0     0.0   REQ TOTAL       146K
          73.46    ASYNC     17K    0.2    57.6     41.9     PR WT      0     0.0     0.0         0.0     0.0   REQ DEFERRED   3747
                   CHNGD       0    0.0  INCLUDED IN ASYNC  PR CMP     0     0.0     0.0         0.0     0.0   -CONT          3713
                   SUPPR       0    0.0                                                                       -FALSE CONT    1763
zOS4        122K   SYNC      83K    0.9    27.6      3.5     NO SCH     0     0.0     0.0         0.0     0.0   REQ TOTAL       148K
          67.51    ASYNC     39K    0.4    71.0     51.9     PR WT      0     0.0     0.0         0.0     0.0   REQ DEFERRED    923
                   CHNGD       0    0.0  INCLUDED IN ASYNC  PR CMP     0     0.0     0.0         0.0     0.0   -CONT           811
                   SUPPR       0    0.0                                                                       -FALSE CONT     456
zOS5        119K   SYNC      97K    1.0    27.2      3.6     NO SCH     0     0.0     0.0         0.0     0.0   REQ TOTAL       151K
          66.11    ASYNC     22K    0.2    59.0     46.1     PR WT      0     0.0     0.0         0.0     0.0   REQ DEFERRED   3634
                   CHNGD       0    0.0  INCLUDED IN ASYNC  PR CMP     0     0.0     0.0         0.0     0.0   -CONT          3598
                   SUPPR       0    0.0                                                                       -FALSE CONT     645
-------------------------------------------------------------------------------------------------------------------------------------
TOTAL       9371K   SYNC    9293K   99.2     3.0      4.5     NO SCH     0     0.0     0.0         0.0     0.0   REQ TOTAL        10M
            5206    ASYNC     78K    0.8    64.6     48.7     PR WT      0     0.0     0.0         0.0     0.0   REQ DEFERRED    22K
                    CHNGD       0    0.0                      PR CMP     0     0.0     0.0         0.0     0.0   -CONT           22K
                    SUPPR       0    0.0                                                                       -FALSE CONT    9548
```

CF activity report: structure activity notes

Db2 lock structure

requests over the RMF interval and average per second

- Anything over 100,000 requests per second is high

Synchronous (SYNC) and asynchronous (ASYNC) service time

- 3 µsec or less is on same system with IC links
- In this example, longer service times due to remote CF a little over 2 km away
  - Significant percentage converted to async

External request / contention

- CONT = real contention
  - Aim to keep below 2% of external requests
- FALSE CONT = false contention: XES recognizes that two resources hashed to the same lock class
  - Aim to keep below 1% of external requests
  - If too high, increase LOCK1 size and rebuild structure (at least double size of structure, then rebuild, to build larger lock table)

CF activity report: structure activity

Db2 group buffer pool

```
STRUCTURE NAME = DSNDBPG_GBP2        TYPE = CACHE   STATUS = ACTIVE PRIMARY          ENCRYPTED = NO
          # REQ     -------------- REQUESTS -------------   -------------- DELAYED REQUESTS -------------
SYSTEM    TOTAL                 #     % OF  -SERV TIME(MIC)-    REASON     #    % OF  ---- AVG TIME(MIC) -----
NAME      AVG/SEC               REQ    ALL     AVG    STD_DEV              REQ    REQ   /DEL     STD_DEV   /ALL

zOS1       4768K   SYNC    4608K   83.5     5.0      5.1     NO SCH     0     0.0    0.0        0.0     0.0
           2649    ASYNC    160K    2.9    63.5     26.5     PR WT      0     0.0    0.0        0.0     0.0
                   CHNGD       0    0.0  INCLUDED IN ASYNC   PR CMP     0     0.0    0.0        0.0     0.0
                   SUPPR       0    0.0                      DUMP       0     0.0    0.0        0.0     0.0

zOS2      37177    SYNC      37K    0.7     9.8     12.9     NO SCH     0     0.0    0.0        0.0     0.0
          20.65    ASYNC      73    0.0    32.4     30.2     PR WT      0     0.0    0.0        0.0     0.0
                   CHNGD       0    0.0  INCLUDED IN ASYNC   PR CMP     0     0.0    0.0        0.0     0.0
                   SUPPR       0    0.0                      DUMP       0     0.0    0.0        0.0     0.0

zOS3       1598    SYNC     504    0.0    28.5      4.1     NO SCH     0     0.0    0.0        0.0     0.0
           0.89    ASYNC   1094    0.0    60.0     79.4     PR WT      0     0.0    0.0        0.0     0.0
                   CHNGD       0    0.0  INCLUDED IN ASYNC   PR CMP     0     0.0    0.0        0.0     0.0
                   SUPPR       0    0.0                      DUMP       0     0.0    0.0        0.0     0.0

zOS4        676K   SYNC    6495    0.1    39.9     15.0     NO SCH     0     0.0    0.0        0.0     0.0
          375.7    ASYNC    670K   12.1   101.0     46.5     PR WT      0     0.0    0.0        0.0     0.0
                   CHNGD       0    0.0  INCLUDED IN ASYNC   PR CMP     0     0.0    0.0        0.0     0.0
                   SUPPR       0    0.0                      DUMP       0     0.0    0.0        0.0     0.0

zOS5      32709    SYNC     427    0.0    28.8      4.1     NO SCH     0     0.0    0.0        0.0     0.0
          18.17    ASYNC     32K    0.6    52.8     32.5     PR WT      0     0.0    0.0        0.0     0.0
                   CHNGD       0    0.0  INCLUDED IN ASYNC   PR CMP     0     0.0    0.0        0.0     0.0
                   SUPPR       0    0.0                      DUMP       0     0.0    0.0        0.0     0.0

-------------------------------------------------------------------------------------------------------------
TOTAL      5516K   SYNC    4652K   84.3     5.1      5.4     NO SCH     0     0.0    0.0        0.0     0.0    -- DATA ACCESS ---
           3064    ASYNC    863K   15.7    92.2     46.1     PR WT      0     0.0    0.0        0.0     0.0    READS       24193
                   CHNGD       0    0.0                      PR CMP     0     0.0    0.0        0.0     0.0    WRITES    2691034
                   SUPPR       0    0.0                      DUMP       0     0.0    0.0        0.0     0.0    CASTOUTS  2481129
                                                                                                             XI'S        15382
```

CF activity report: structure activity notes

Db2 group buffer pool

requests over the RMF interval and average per second

Synchronous (SYNC) and asynchronous (ASYNC) service time

- Expect sync service times for GBP to be 2-3 times greater than for LOCK1
- In this example, longer service times due to remote CF a little over 2 km away
  - Significant percentage converted to async

Data access

- If only writes are non-zero, indicates a secondary group buffer pool

CFRM policy entries for Db2

```
STRUCTURE NAME(DSNDB0P_LOCK1) SIZE(2G)
   INITSIZE(1G)
   MINSIZE(1G)
   REBUILDPERCENT(1)
   PREFLIST(PRODCF00, PRODCF01, PRODCF10, PRODCF11)
```

LOCK1

- SIZE (2G) indicates maximum structure size
- INITSIZE (1G) is initial allocation size
- SETXCF commands can increase allocation from 1G to desired value <= 2G
  - Increase only available for record list entries (RLEs) unless rebuild the lock structure
  - Size increase (double or more) plus rebuild required to increase lock table to reduce false contention rate
- ALLOWAUTOALT(YES) recommended to allow XES to adjust between INITSIZE and SIZE
  - Can decrease down to MINSIZE
  - Will only affect RLEs, and not designed for spikes
  - Db2 13: IRLM can trigger LOCK1 allocation increase to add RLEs; should be more responsive than XES increase
- FULLTHRESHOLD > 0 - XES will monitor structure size and consider size increase at value specified

CFRM policy entries for Db2

```
STRUCTURE NAME(DSNDB0P_LOCK1) SIZE(2G)
   INITSIZE(1G)
   MINSIZE(1G)
   REBUILDPERCENT(1)
   PREFLIST(PRODCF00, PRODCF01, PRODCF10, PRODCF11)
```

LOCK1

- Where is DUPLEX in this definition?
  - It defaulted to DUPLEX(DISABLED)
- DUPLEX(DISABLED) is appropriate for LOCK1 (and SCA) if allocated on external CF
  - Or at least on separate system (CEC) from Db2 members
- In 2-CEC 2-ICF configuration, DUPLEX(ENABLED) – system managed duplexing, recommended for LOCK1 and SCA for availability
  - Avoid single point of failure
- For LOCK1, DUPLEX(ENABLED) is very expensive: synchronous system managed duplexing
  - Asynchronous system managed duplexing performs much better
  - Consider DUPLEX(ENABLED,ASYNCONLY) for 2-CEC 2-ICF, or similar, configurations
    - More and more customer experience with async CF duplexing

CFRM policy entries for Db2

```
STRUCTURE NAME(DSNDB0P_GBP1) SIZE(4137M)
   INITSIZE(2121M)
   MINSIZE(2121M)
   ALLOWAUTOALT(YES)
   DUPLEX(ENABLED)
   PREFLIST(PRODCF00, PRODCF01, PRODCF10, PRODCF11)
```

GBPs

- DUPLEX (ENABLED) strongly recommended
  - Availability, not performance
  - Minimal performance impact
- ALLOWAUTOALT(YES) allows XES to adjust ratio of directory entries to data elements and to increase allocation up to SIZE
  - Can also decrease down to MINSIZE
- Changing ratio without ALLOWAUTOALT requires –ALTER GBPOOL command ***and*** manual structure rebuild
  - Probable increase in application elapsed time during manual rebuild

Db2 displays and reports

-DIS GROUPBUFFERPOOL at group level

- -DIS GBPOOL(*) TYPE(GCONN) GDETAIL(*)
  - Contains status and definition information as well as statistics
  - Reports statistics since GBP allocation
- -DIS GBPOOL(*) TYPE(GCONN) GDETAIL(INTERVAL)
  - To monitor an interval, execute this command before and after the desired interval
  - Output messages from second command will show GBP statistics for the interval
- Typical problems due to incorrectly defined GBP
  - Directory entry reclaims
  - XIs due to directory entry reclaims: increased synchronous I/O
  - Writes failed due to lack of storage: application delays, or worse

Db2 displays and reports

-DIS GBPOOL GDETAIL status

- This excerpt shows the first few messages from the DIS GBPOOL output for GBP0 in this case
- Information in these messages does not change frequently
  - Example: directory to data ratio

```
DSNB750I  -D2P1 DISPLAY FOR GROUP BUFFER POOL GBP0 FOLLOWS
DSNB755I  -D2P1 DB2 GROUP BUFFER POOL STATUS
             CONNECTED                                  = YES
             CURRENT DIRECTORY TO DATA RATIO            = 26
             PENDING DIRECTORY TO DATA RATIO            = 26
             CURRENT GBPCACHE ATTRIBUTE                 = YES
             PENDING GBPCACHE ATTRIBUTE                 = YES
DSNB756I  -D2P1   CLASS CASTOUT THRESHOLD                    = 5, 0
             GROUP BUFFER POOL CASTOUT THRESHOLD        = 25%
             GROUP BUFFER POOL CHECKPOINT INTERVAL      = 4 MINUTES
             RECOVERY STATUS                            = NORMAL
             AUTOMATIC RECOVERY                         = Y
DSNB757I  -D2P1 MVS CFRM POLICY STATUS FOR DSNP01G_GBP0      = NORMAL
             MAX SIZE INDICATED IN POLICY               = 1576960 KB
             DUPLEX INDICATOR IN POLICY                 = ENABLED
             CURRENT DUPLEXING MODE                     = DUPLEX
             ALLOCATED                                  = YES
DSNB758I  -D2P1     ALLOCATED SIZE                           = 788480 KB
               VOLATILITY STATUS                        = NON-VOLATILE
               REBUILD STATUS                           = DUPLEXED
               CFNAME                                   = CF01
```

Db2 displays and reports

-DIS GBPOOL GDETAIL write section

The three highlighted fields should be zero

- WRITES FAILED DUE TO LACK OF STORAGE
- RECLAIMS FOR DIRECTORY ENTRIES
- CROSS INVALIDATIONS DUE TO DIRECTORY RECLAIMS

```
DSNB786I  -D2P1   WRITES
               CHANGED PAGES                          = 63440609
               CLEAN PAGES                            = 0
               FAILED DUE TO LACK OF STORAGE          = 0
             CHANGED PAGES SNAPSHOT VALUE             = 0
DSNB787I  -D2P1   RECLAIMS
               FOR DIRECTORY ENTRIES                  = 0
               FOR DATA ENTRIES                       = 486093
             CASTOUTS                                 = 7680071
DSNB788I  -D2P1   CROSS INVALIDATIONS
               DUE TO DIRECTORY RECLAIMS              = 0
               DUE TO WRITES                          = 39422068
               EXPLICIT                               = 0
```

Db2 displays and reports

-DIS GROUPBUFFERPOOL at member level

–DIS GBPOOL(*) TYPE(GCONN) MDETAIL(INTERVAL)

- To monitor an interval, execute this command before and after the desired interval
- Output messages from second command will show GBP statistics for the interval

Performance information:

– Hit ratio for reads to the GBP when the local BP page is cross-invalidated (XI)

– Hit ratio for reads to the GBP when the page is not found in the local BP (NF)

– Lots of other interesting fields

Db2 displays and reports

-DIS GBPOOL MDETAIL excerpt

- Key member level information in the fields highlighted by the arrows
- GBP read because local BP page cross-invalidated (XI)
  - A – data returned
  - B – data not returned
  - A/(A+B) = XI hit ratio -> tuning target
    - Here 99.97%
- GBP read because local BP page not found (NF)
  - C – data returned
  - D – data not returned
  - C/(C+D) = NF hit ratio (lucky benefit)
    - Here 24.78%%

```
DSNB773I  -D2P1 MEMBER DETAIL STATISTICS
        SYNCHRONOUS READS
         DUE TO BUFFER INVALIDATION
          DATA RETURNED          = 1188514     A
          DATA NOT RETURNED      = 16676       B
DSNB774I  -D2P1    DUE TO DATA PAGE NOT IN BUFFER POOL
          DATA RETURNED          = 148271      C
          DATA NOT RETURNED      = 1702664     D
DSNB789I  -D2P1    REGISTER PAGE LIST       = 126048
         PAGES RETRIEVED          = 821086
DSNB776I  -D2P1   SYNCHRONOUS WRITES
         CHANGED PAGES            = 2160808
         CLEAN PAGES            = 0
DSNB777I  -D2P1   ASYNCHRONOUS WRITES
         CHANGED PAGES            = 520572
         CLEAN PAGES            = 0
        FAILED DUE TO LACK OF STORAGE  = 0
        WRITE-AROUND PAGES          = 0 C
```

Db2 displays and reports

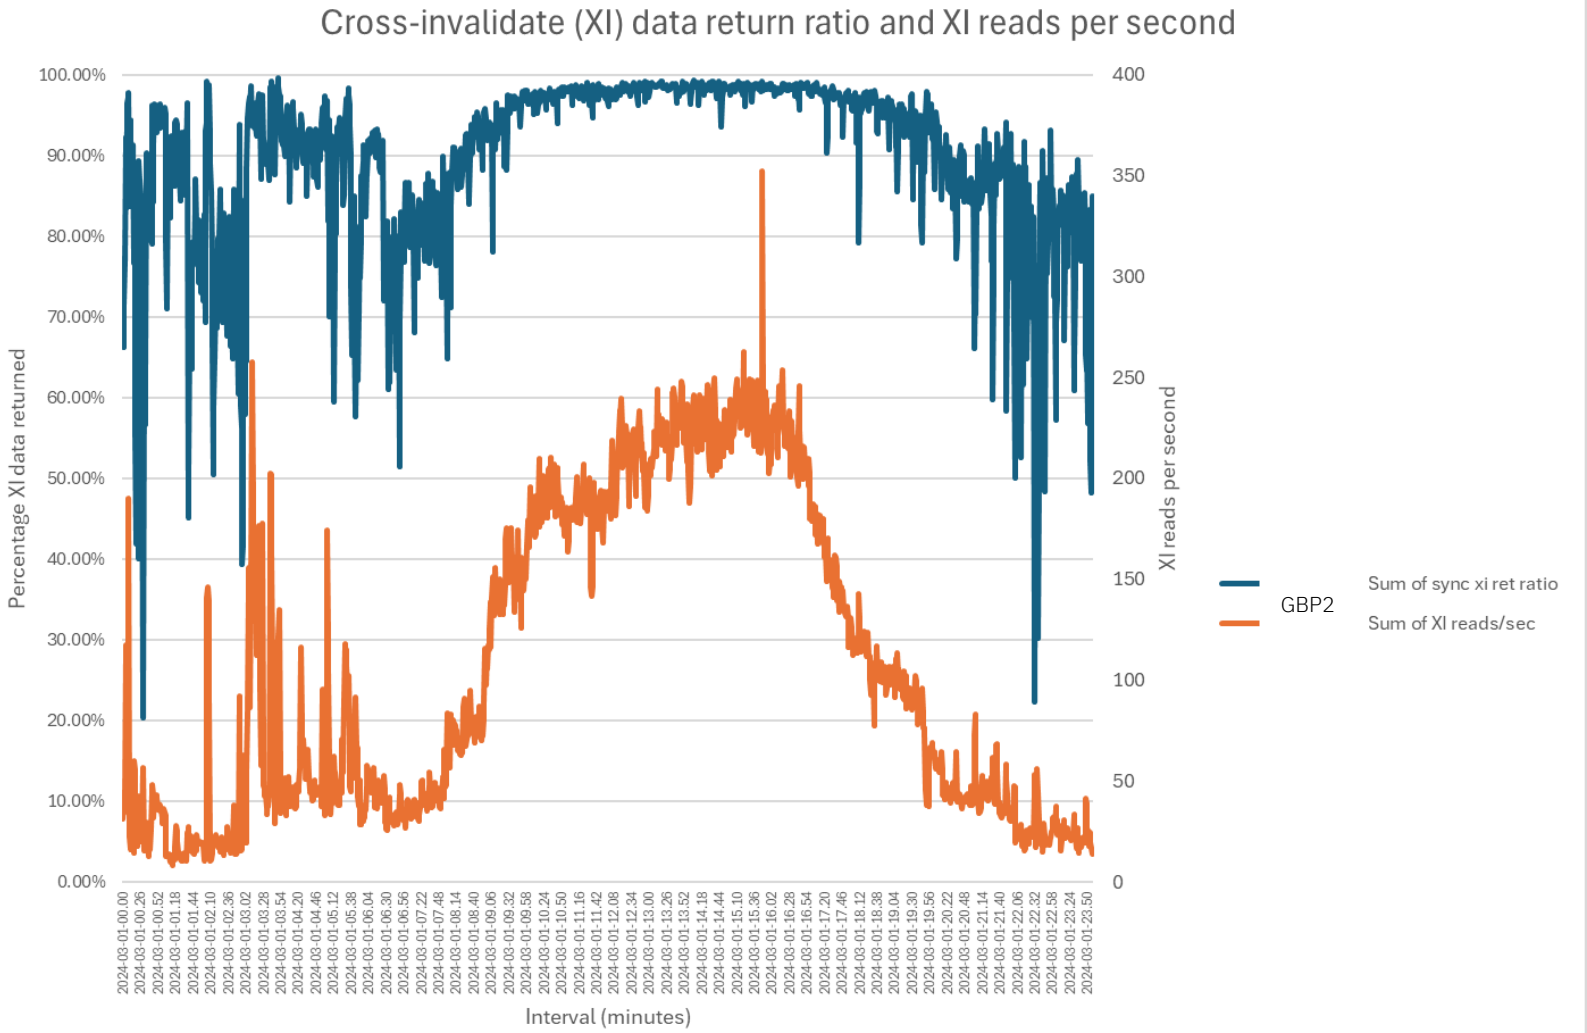

Graphical view of XI read hit ratio and XI reads per second

- Blue (left axis) is percentage of XI reads with data returned from GBP
- Orange (right axis) is XI reads per second to the GBP
- Notice high read hit ratio during online period
  - Inconsistent read hit ratio during batch periods

What if you want to improve your read hit ratios?

- Increase data element residency time
  - Db2 12: increase GBP size and measure
  - Db2 13 and CFCC 25: GBP residency times

Db2 13 and GBP residency time

IFCID 230 (part of Db2 statistics class 5)

-DIS GBPOOL GDETAIL

Higher numbers better

- The longer the data element (Db2 page) is resident in the GBP, the greater likelihood it will be there when requested

Exception:

- Zero (0) means the data element (or directory entry) has never been stolen, which is the best possible scenario

Suggestions

Suggestions: CFs and structures

Dedicated processors for production CF LPARs
- 2 or more, depending on workload

Use CFSizer tool: https://www.ibm.com/support/pages/cfsizer
  - Sunset end of 2026: use z/OSMF instead
- Take output as starting point, round up for INITSIZE
- Define SIZE up to 2 x INITSIZE

Use DUPLEX(ENABLED) and ALLOWAUTOALT(YES) for GBPs

Use ALLOWAUTOALT(YES) for LOCK1 and SCA

DUPLEX setting for LOCK1 and SCA depends upon CF and/or ICF configuration
- This is an availability, as well as performance, topic

Suggestions: Db2 structure monitoring

Monitor sync service times, async conversions and false contention for LOCK1

- Sync service times and async conversion can be dependent upon configuration
- Target: real contention < 2% of requests; false contention < 1% of requests

Monitor GBPs: sync service times, async conversion

- Group level: write failed due to lack of storage, directory reclaims, cross-invalidations due to directory reclaims
- Member level: XI hit ratio, NF hit ratio

Ensure Db2 team and z/OS team communicate to catch issues early

Questions?

Summary

Db2 data sharing adds topics to performance (and capacity) discussions

- Synchronous CF service times are an important factor
- Multiple views of LOCK1 and GBP performance
- CF reports, Db2 monitors or reports, Db2 displays

Db2 data sharing monitoring and tuning generally crosses organizational boundaries

- Ensure your teams communicate if they perceive an issue

Thank you!

Tori DiDonato
victoria.felt@ibm.com

Anna McKee
anna.mckeeh@ibm.com

Mark Rader
mrader@us.ibm.com