



Db2 for z/OS Data Sharing Performance Analysis and Tuning

John Campbell
Distinguished Technical Consultant
Broadcom

John.Campbell@Broadcom.com



Disclaimer

Certain information in this presentation may outline Broadcom's general product direction. This presentation shall not serve to (i) affect the rights and/or obligations of Broadcom or its licensees under any existing or future license agreement or services agreement relating to any Broadcom software product; or (ii) amend any product documentation or specifications for any Broadcom software product. This presentation is based on current information and resource allocations as of April 21, 2026, and is **subject to change or withdrawal by Broadcom at any time without notice. The development, release and timing of any features or functionality described in this presentation remain at Broadcom's sole discretion.**

Notwithstanding anything in this presentation to the contrary, upon the general availability of any future Broadcom product release referenced in this presentation, Broadcom may make such release available to new licensees in the form of a regularly scheduled major product release. Such release may be made available to licensees of the product who are active subscribers to Broadcom maintenance and support, on a when and if-available basis. The information in this presentation is not deemed to be incorporated into any contract.

Broadcom may use any feedback provided by you related to a Broadcom product or this presentation for any Broadcom business purposes (including but not limited to, preparation, reproduction, and distribution of derivative works based upon such feedback), without any obligation to you including consent or payment.

Copyright © 2026 Broadcom. All rights reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. Broadcom, the pulse logo, Connecting everything, CA Technologies and the CA Technologies logo are among the trademarks of Broadcom.

THIS PRESENTATION IS FOR YOUR INFORMATIONAL PURPOSES ONLY. Broadcom assumes no responsibility for the accuracy or completeness of the information. TO THE EXTENT PERMITTED BY APPLICABLE LAW, BROADCOM PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON INFRINGEMENT. In no event will Broadcom be liable for any loss or damage, direct or indirect, in connection with this presentation, including, without limitation, lost profits, lost investment, business interruption, goodwill, or lost data, even if Broadcom is expressly advised in advance of the possibility of such damages.



Agenda

- Introduction
- Db2 Data Sharing Technical Architecture
- Global locking
 - Db2 LOCK1 structure
 - Global lock contention
 - Resizing LOCK1 Structure
- Group Buffer Pools
 - Db2 Group Buffer Pool structures
 - GBP reads
 - GBP writes
 - GBP castout
 - GBP tuning

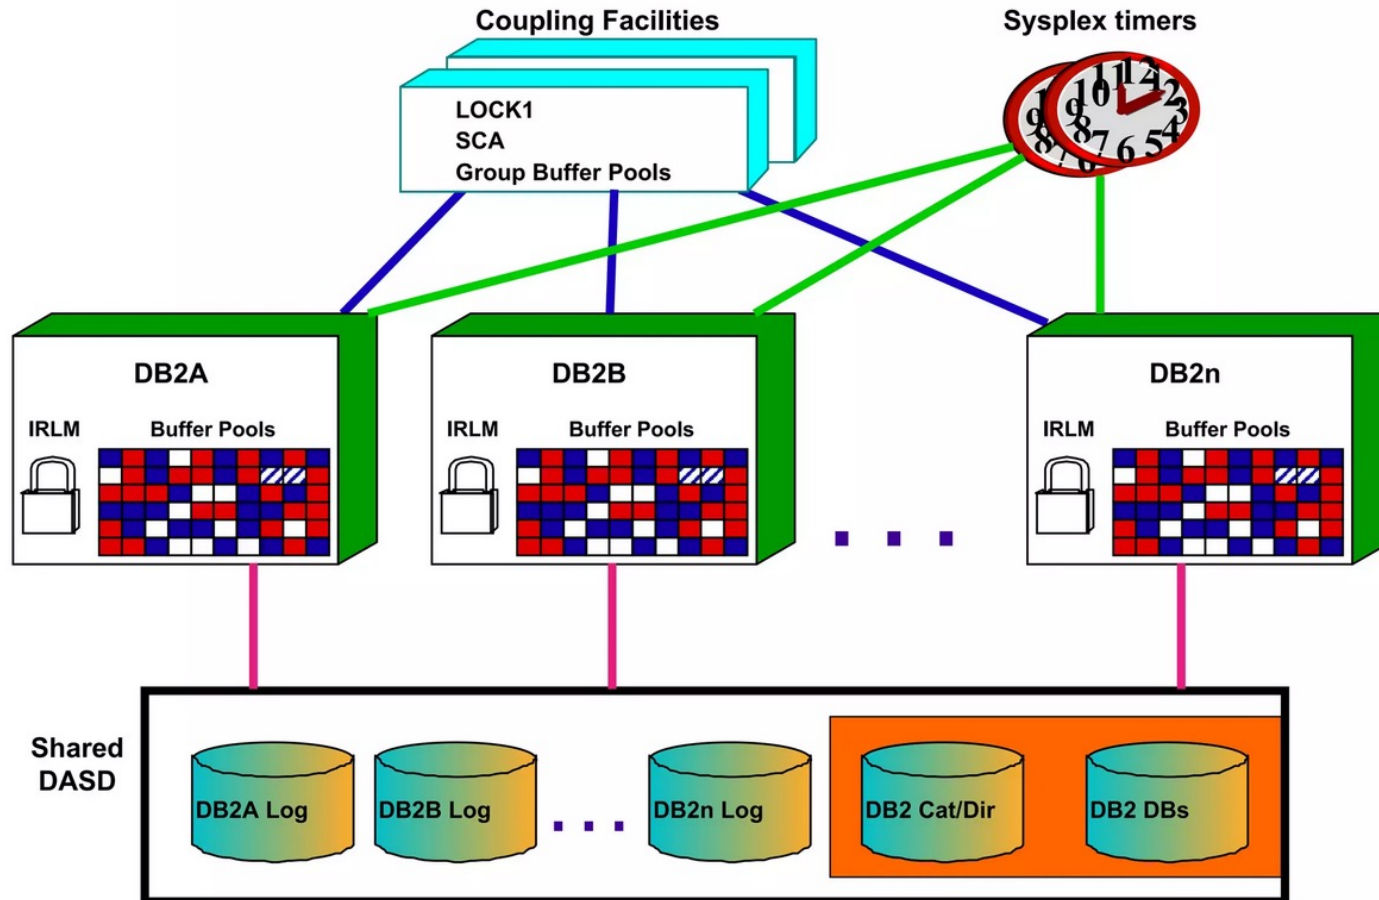


Introduction

- Earlier database clustering (shared data) solutions had issues
 - Excessive performance cost
 - Increased latency due to message passing
 - Seriously degraded throughput
 - Very poor scale out
- Contributory factors
 - Global locking protocol using message passing even when lock granted without contention
 - Block refresh from shared DASD where heavy write-write interest
 - No “pay-as-you-go (PAYG)” protocols
- Examples
 - IMS DB Data Sharing Mk. 1
 - Oracle RAC Cluster

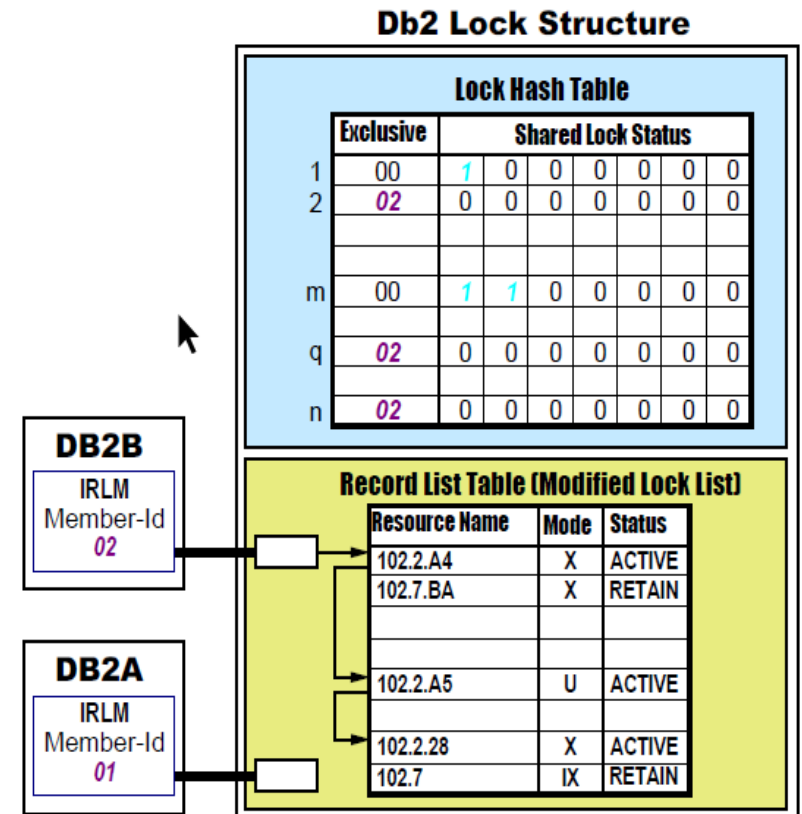


Db2 Data Sharing Architecture

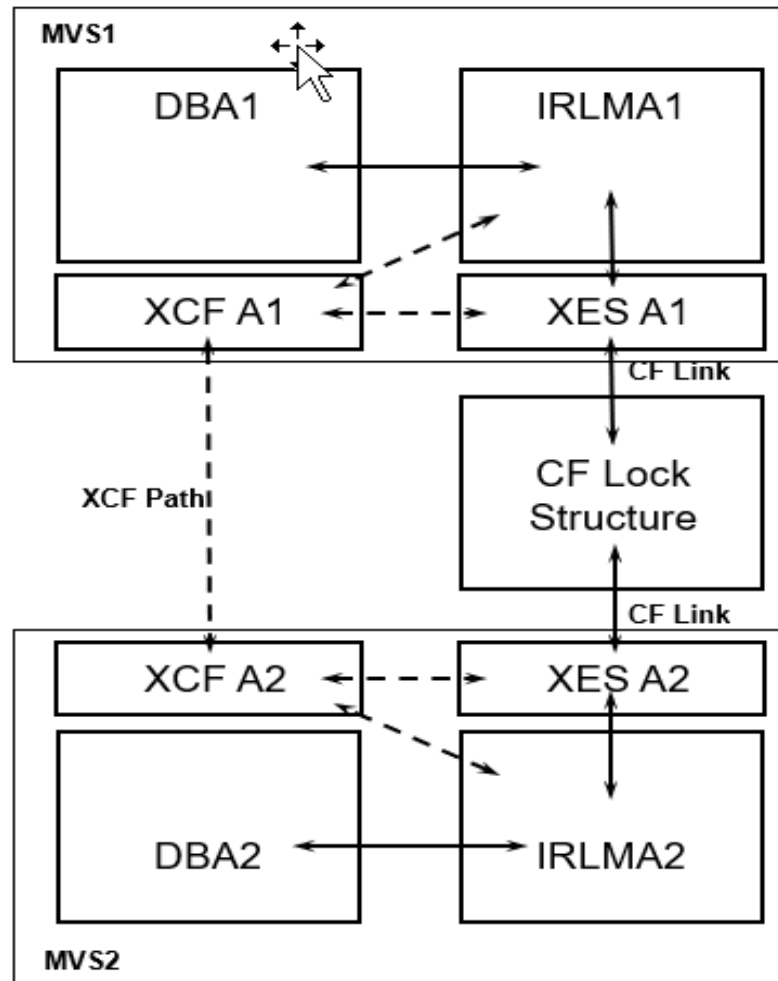


Db2 LOCK1 structure

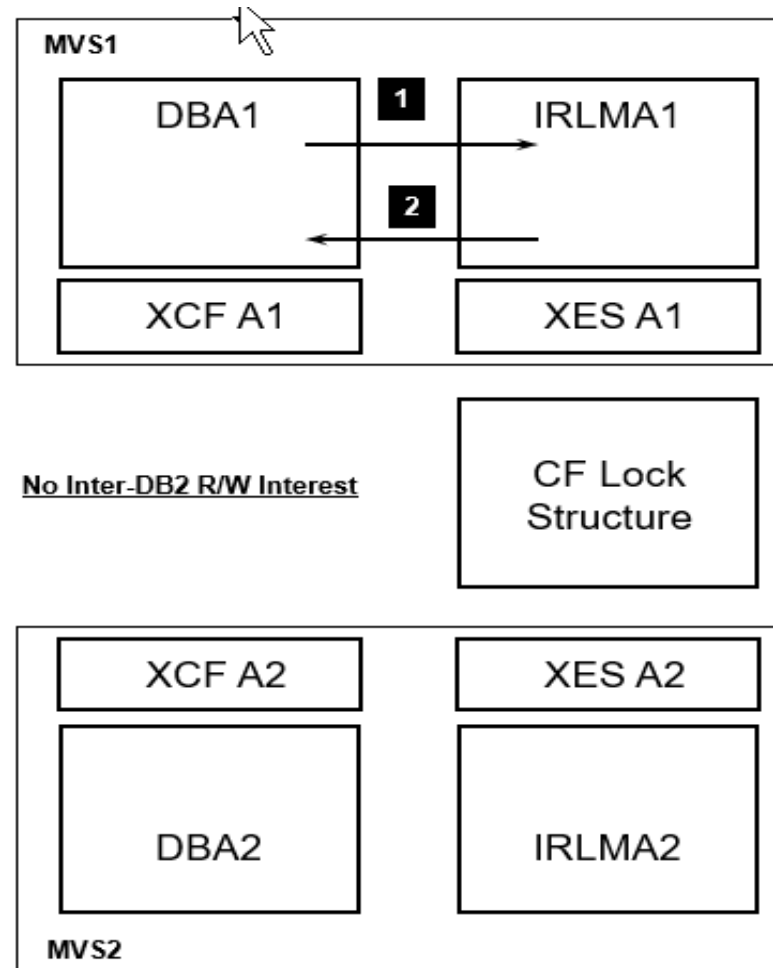
- What is it used for?
 - Fast inter-system global lock conflict detection
 - Optimization for fast grant of global lock where no contention
 - Fast inter-system page latch conflict for pages where sub-page concurrency is allowed
 - Row level locking
 - Space map pages
 - Index leaf pages
 - Dynamic tracking of inter-system read/write interest for Db2 objects
 - Retained locks in the case of Db2 member failure



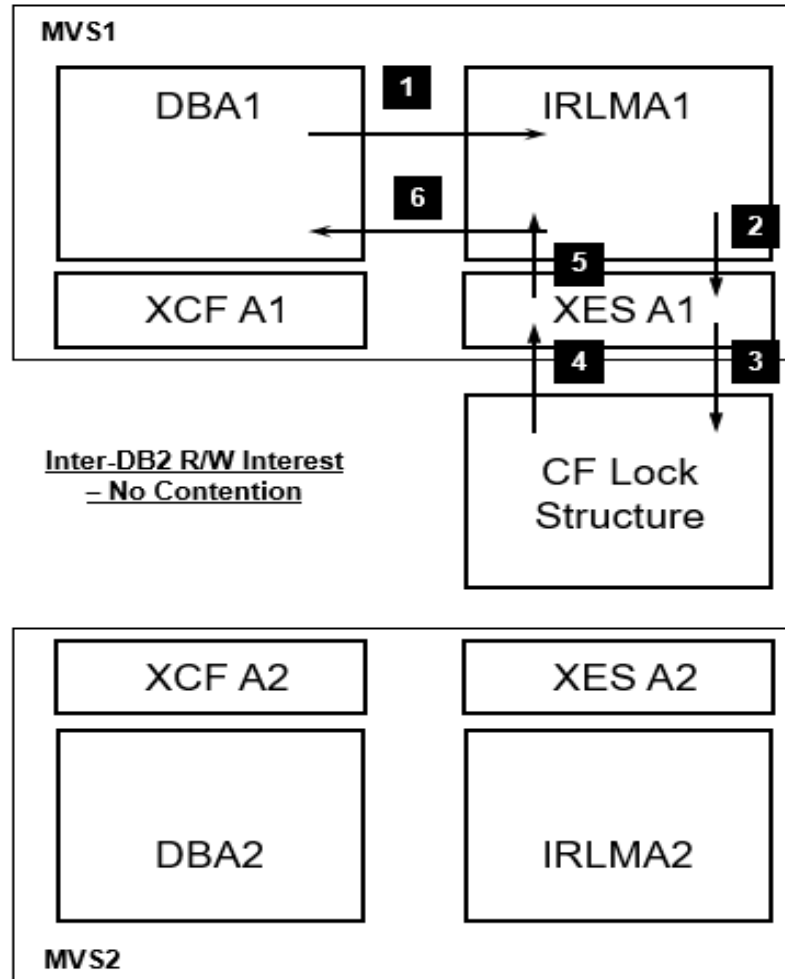
Locking components



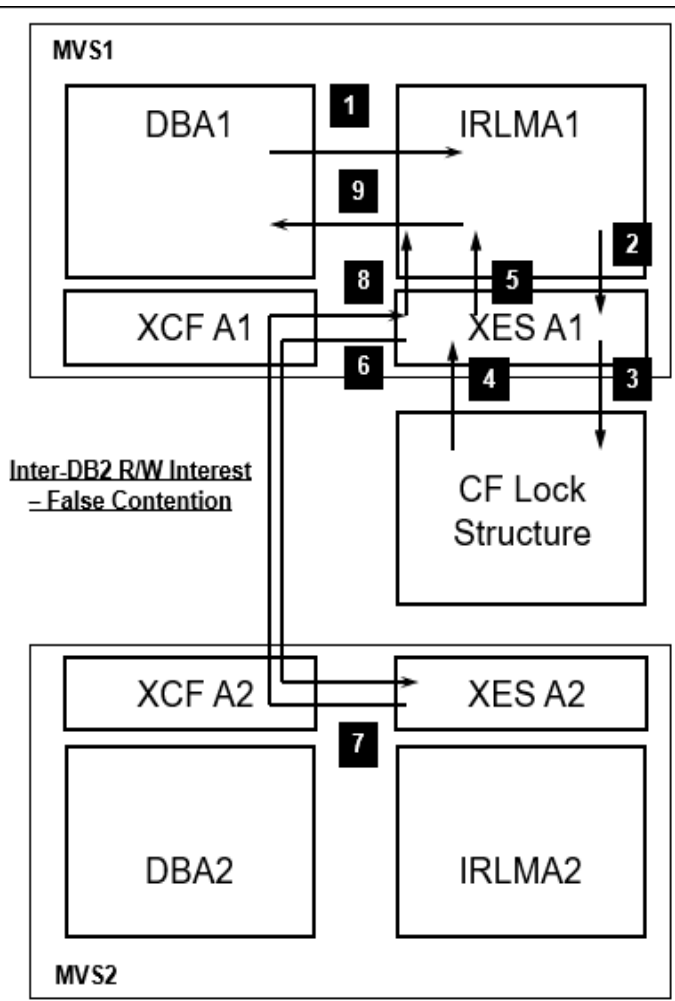
Lock processing



Lock processing ...



Lock processing ...



Db2 LOCK1 structure ...

- LOCK1 structure size is defined in the CFRM policy

```
STRUCTURE NAME(DSNDB2_LOCK1)
SIZE(1024M)
INITSIZE(512M)
ALLOWAUTOALT(NO)
PREFLIST(COUPLE01, COUPLE02)
```

By default, Db2 tries to do a 50-50 split between Lock Hash Table and Record List Table

Lock Hash Table (LTEs) = 256MB
Record List Table (RLEs) = 256MB

LTE size depends on number of data sharing members
2 bytes for 1-7 members
4 bytes for 8-23 members
8 bytes for 24-32 members

```
STRUCTURE NAME(DSNDB2_LOCK1)
SIZE(1024M)
INITSIZE(768M)
ALLOWAUTOALT(NO)
PREFLIST(COUPLE01, COUPLE02)
```

The Lock Hash Table has to be a power of 2

Lock Hash Table (LTEs) = 256MB
Record List Table (RLEs) = 512MB



The record table is susceptible to storage shortages if the structure is too small or if the allocation of the lock table leaves too little storage for the record table



Db2 LOCK1 structure ...

- Shortage of RLEs

```
DXR170I  ir7mx THE LOCK STRUCTURE wwwwww IS zz% IN USE      ← Alert at 50%, 60%, 70%
DXR142E  ir7mx THE LOCK STRUCTURE wwwwww IS zzz% IN USE     ← Alert at 80%, 90%, 100%
```

- At 80% full, data sharing continues with no restrictions, but storage is approaching a critical threshold
- At 90% full, only 'must complete' type of requests that require lock structure storage are processed
- At 100% full, any request that requires lock structure storage is denied with an 'out of lock structure storage' error



Global lock contention

DATA SHARING LOCKING	QUANTITY	/SECOND	/THREAD	/COMMIT
GLOBAL CONTENTION RATE (%)	0.64			
FALSE CONTENTION RATE (%)	0.11			
...				
SYNCH.XES - LOCK REQUESTS	227.5M	10.6K	1368.75	458.86
SYNCH.XES - CHANGE REQUESTS	1340.7K	62.24	8.07	2.70
SYNCH.XES - UNLOCK REQUESTS	225.8M	10.5K	1358.14	455.30
ASYNCH.XES - CONVERTED LOCKS	3485.41	0.48	3.87	0.00
...				
SUSPENDS - IRLM GLOBAL CONT	34192.00	1.59	0.21	0.07
SUSPENDS - XES GLOBAL CONT.	6076.00	0.28	0.04	0.01
SUSPENDS - FALSE CONT. MBR	21575.00	1.00	0.13	0.04



Global Contention should be less than 3-5% of XES IRLM Requests

Global Contention = SUSPENDS - IRLM + XES + FALSE (A)
 XES IRLM Requests = (SYNCH. XES - LOCK + CHANGE + UNLOCK)
 + ASYNCH.XES - CONVERTED LOCKS + (SUSPENDS - IRLM + XES + FALSE) (B)
 Global Contention Rate (%) = (A)/(B)*100



False Contention should be less than 1-3% of XES IRLM Requests

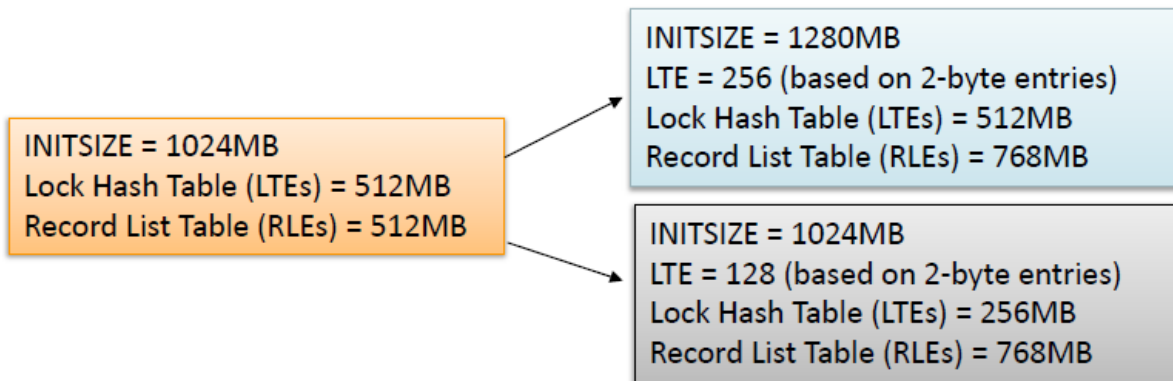
False Contention = SUSPENDS - FALSE (C)
 False Contention Rate (%) = (C)/(B)*100

False Cont = false contention on lock table hash anchor point
 Could be minimised by increasing the number of LTEs in the
 Lock Hash Table



Resizing Db2 LOCK1 Structure

- Increase size the LOCK1 structure if high rate of false contention and/or shortage of RLEs
- If shortage of RLEs but false contention rate is OK, consider using IRLM option **LTE=** to control the size of the Lock Hash Table
 - Specify a value for the LTE= parameter in the IRLMPROC or issue F irlmproc SET,LTE= command
 - Requires a REBUILD of the LOCK1 structure



For LTE=	2-byte entries	4-byte entries
8	16 MB	32 MB
16	32 MB	64 MB
32	64 MB	128 MB
64	128 MB	256 MB
128	256 MB	512 MB
256	512 MB	1024 MB
512	1024 MB	2048 MB
1024	2048 MB	4096 MB
2048	4096 MB	



Reducing Global lock contention

- Use a light-weight locking protocol (isolation level) and exploit read only lock avoidance
 - Benefits of lock avoidance
 - Increased concurrency by reducing lock contention
 - Decreased lock and unlock activity and avoid associated CPU overhead
 - Decrease number of CF requests
 - Minimize the impact of retained locks
 - Use ISOLATION(CS) with CURRENTDATA(NO) or use ISOLATION(UR)
- Commit frequently
 - Reduce lock contention
 - Improve effectiveness of global lock avoidance
- Avoid serialisation points e.g., single row table used as a ‘counter’
 - Use an IDENTITY column or pull value from SEQUENCE column, with CACHE option
- Exploit table and index partitioning



Global lock contention



DATA SHARING LOCKING	QUANTITY	/SECOND	/THREAD	/COMMIT
...				
SYNCH.XES - LOCK REQUESTS	227.5M	10.6K	1368.75	458.86
SYNCH.XES - CHANGE REQUESTS	1340.7K	62.24	8.07	2.70
SYNCH.XES - UNLOCK REQUESTS	225.8M	10.5K	1358.14	455.30
ASYNCH.XES - CONVERTED LOCKS	1315.00	0.06	0.01	0.00
...				
PSET/PART P-LCK NEGOTIATION	18037.00	0.84	0.11	0.04
PAGE P-LOCK NEGOTIATION	2863.00	0.13	0.02	0.01
OTHER P-LOCK NEGOTIATION	9067.00	0.42	0.05	0.02
P-LOCK CHANGE DURING NEG.	15991.00	0.74	0.10	0.03

P-lock Negotiation should be less than 3-5% of XES IRLM requests

$$\text{P-lock Negotiation} = \text{PSET/PART P-LCK NEGOTIATION} + \text{PSET/PART P-LCK NEGOTIATION} + \text{OTHER P-LCK NEGOTIATION (A)}$$

$$\text{XES IRLM Requests} = (\text{SYNCH.XES - LOCK} + \text{CHANGE} + \text{UNLOCK}) + \text{ASYNCH.XES - CONVERTED LOCKS} + (\text{SUSPENDS - IRLM} + \text{XES} + \text{FALSE}) \text{ (B)}$$

$$\text{P-lock Negotiation Rate (\%)} = (A)/(B)*100$$

- **P-lock contention and negotiation can cause IRLM latch contention, page latch contention, asynchronous GBP write, active log write, GBP read**
 - Page P-lock contention by one thread causes Page Latch contention for all other threads in the same member trying to get to the same page




Global lock contention ...

- Breakdown of page P-lock type in Db2 Statistics

GROUP	TOTAL	CONTINUED	QUANTITY	/SECOND	/THREAD	/COMMIT
PAGE P-LOCK LOCK REQ			877.4K	122.88	14.91	3.64
SPACE MAP PAGES			83552.00	11.70	1.42	0.35
DATA PAGES			10663.00	1.49	0.18	0.04
INDEX LEAF PAGES			783.2K	109.69	13.31	3.25
PAGE P-LOCK UNLOCK REQ			926.8K	129.80	15.75	3.84
PAGE P-LOCK LOCK SUSP			8967.00	1.26	0.15	0.04
SPACE MAP PAGES			593.00	0.08	0.01	0.00
DATA PAGES			143.00	0.02	0.00	0.00
INDEX LEAF PAGES			8231.00	1.15	0.14	0.03
PAGE P-LOCK LOCK NEG			10285.00	1.44	0.17	0.04
SPACE MAP PAGES			8.00	0.00	0.00	0.00
DATA PAGES			10.00	0.00	0.00	0.00
INDEX LEAF PAGES			10267.00	1.44	0.17	0.04

For insert-intensive workloads with high page P-lock contention on space map pages, consider MEMBER CLUSTER (optionally combined with APPEND + LOCKSIZE ROW)

 Do not use APPEND or LOCKSIZE ROW without MEMBER CLUSTER for an INSERT-at-the-end intensive workload → may result in excessive page p-lock contention on data pages and space map pages

For heavily concurrent insert activity (many concurrent threads) with high page P-lock contention on data pages caused by space search and false leads, consider INSERT ALGORITHM 2 (aka Fast Un-clustered INSERT)

If data page P-lock contention on small tables with LOCKSIZE ROW, consider MAXROWS=1 and LOCKSIZE PAGE to 'simulate' row locking and reduce spacemap free space update

If index tree P-lock (high index splits), consider

- Freespace tuning (PCTFREE, FREEPAGE)
- Exploit data and index partitioning to 'dilute' hot spot
- Increase index page size – warning: could also aggravate contention!



Global lock contention ...

- Db2 Accounting for more granular information

CLASS 3 SUSPENSIONS	AVERAGE TIME	AV. EVENT	TIME/EVENT
LOCK/LATCH(DB2+IRLM)	0.000097	0.58	0.000167
IRLM LOCK+LATCH	0.000004	0.21	0.000021
DB2 LATCH	0.000092	0.37	0.000251
..			
PAGE LATCH	0.000595	1.89	0.000314
NOTIFY MSGS	0.000000	0.00	N/C
GLOBAL CONTENTION	0.004844	9.26	0.000523

GLOBAL	CONTENTION	L-LOCKS	AVERAGE TIME	AV. EVENT	GLOBAL	CONTENTION	P-LOCKS	AVERAGE TIME	AV. EVENT
L-LOCKS			0.000011	0.05	P-LOCKS			0.004833	9.21
PARENT (DB, TS, TAB, PART)			0.000000	0.00	PAGESET/PARTITION			0.000000	0.00
CHILD (PAGE, ROW)			0.000000	0.00	PAGE			0.004790	9.16
OTHER			0.000011	0.05	OTHER			0.000043	0.05

LOCKING	AVERAGE	TOTAL
...		
LOCK REQUEST	139.05	2642
UNLOCK REQUEST	34.63	658
QUERY REQUEST	56.26	1069
CHANGE REQUEST	13.32	253
OTHER REQUEST	0.00	0
TOTAL SUSPENSIONS	0.32	6
LOCK SUSPENSIONS	0.00	0
IRLM LATCH SUSPENS.	0.32	6
OTHER SUSPENS.	0.00	0

DATA SHARING	AVERAGE	TOTAL
GLOBAL CONT RATE(%)	3.88	N/A
FALSE CONT RATE(%)	0.14	N/A
...		
LOCK REQ - XES	128.00	2432
UNLOCK REQ - XES	78.21	1486
CHANGE REQ - XES	5.63	107
SUSPENDS - IRLM	8.26	157
SUSPENDS - XES	0.00	0
CONVERSIONS- XES	0.68	13
FALSE CONTENTIONS	0.32	6

GROUP TOT4K	AVERAGE	TOTAL
...		
PG P-LOCK LOCK REQ	65.26	1240
SPACE MAP PAGES	6.95	132
DATA PAGES	16.11	306
INDEX LEAF PAGES	42.21	802
PG P-LOCK UNLOCK REQ	58.26	1107
PG P-LOCK LOCK SUSP	8.84	168
SPACE MAP PAGES	0.95	18
DATA PAGES	1.84	35
INDEX LEAF PAGES	6.05	115



Global lock contention ...

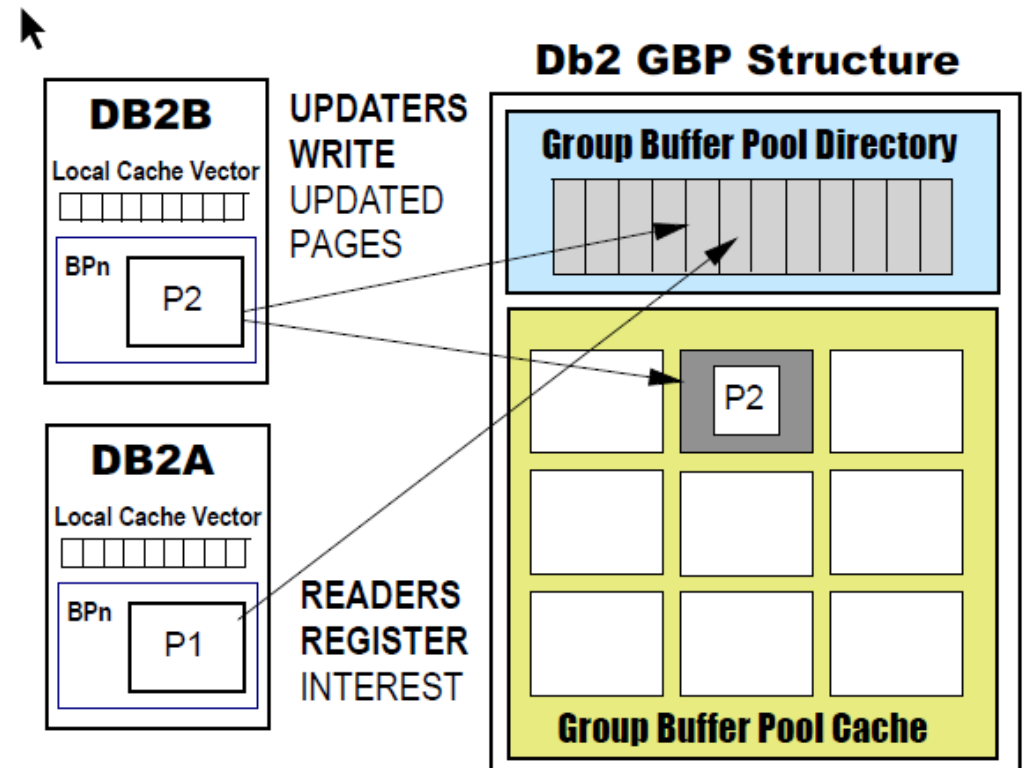
- Use Lock suspension report to identify 'hot spots'
 - For detailed analysis, start the following Db2 Performance trace for short periods of time during peak processing
 - -START TRACE(P) CLASS(30) IFCID (44,45,105,107,213,214,215,216,226,227)
 - Sample Db2 OMPE report to generate a CSV file that can be easily loaded into a spreadsheet

```
LOCKING
      REPORT
      LEVEL(SUSPENSION)
      DDNAME(LORPTDD1)
      SPREADSHEETDD(SPSHDD)
      ORDER(DATABASE-PAGESET)
```



Db2 Group Buffer Pool structures

- What are they used for?
 - Register buffers for cross invalidation (XI)
 - List' option provided for prefetch
 - Write changed pages and send XI signals
 - Hardware instruction to test vector bits for buffer XI
 - Fast refresh of XI'ed buffers
 - Store-in cache by default
 - Force-at-commit database write protocol for writes to CF



Cross Invalidations

- Two reasons for cross invalidations
 1. Perfectly normal condition in an active data sharing environment with inter-system read/write interest
 2. Directory entries reclaims which you want to tune away and eliminate
 - CPU and I/O overhead if there is not enough directory entries
 - Extra CF access and Sync I/O Read
- Db2 –DISPLAY GROUPBUFFERPOOL(*) TYPE(GCONN) GDETAIL(INTERVAL)

```
DSNB787I - RECLAIMS
           FOR DIRECTORY ENTRIES      = 0
           FOR DATA ENTRIES          = 0
           CASTOUTS                   = 0

DSNB788I - CROSS INVALIDATIONS
           DUE TO DIRECTORY RECLAIMS  = 0
           DUE TO WRITES              = 0
           EXPLICIT                   = 0
```



Reclaims for Directory Entries should be minimised
Cross Invalidations due to Directory Reclaims should be zero



GBP reads

GROUP BP14	QUANTITY	/SECOND	/THREAD	/COMMIT
...				
SYN.READ(XI)-DATA RETURNED	1932.00	0.09	0.01	0.00
SYN.READ(XI)-NO DATA RETURN	39281.6K	1823.66	236.31	79.22
SYN.READ(NF)-DATA RETURNED	22837.00	1.06	0.14	0.05
SYN.READ(NF)-NO DATA RETURN	6955.8K	322.93	41.85	14.03



Sync.Read(XI) miss ratio should be < 10%

TOTAL SYN.READ(XI) (A) = SYN.READ(XI)-DATA RETURNED
+ SYN.READ(XI)-NO DATA RETURN

Sync.Read(XI) miss ratio = SYN.READ(XI)-NO DATA RETURN / (A)

- Local Buffer Pool search -> GBP search -> DASD I/O
- SYN.READ(NF) = Local Buffer Pool miss
- SYN.READ(XI) = Local Buffer Pool hit but cross invalidated buffer
 - Data should be found in GBP
 - If not, GBP may be too small, or pages have been removed because of directory entry reclaims



GBP writes

- Changed Pages Sync Written to GBP / force write
 - Commit
 - P-lock negotiation
- Changed Pages Async Written to GBP
 - Deferred Write (hitting VDWQT or DWQT)
 - System checkpoint
- Pages in Write-Around
 - Applies only to Pages Async Written to GBP
 - DB2 conditionally enables and disables GBP write-around protocol based on % unavailable pages
 - Turn on at GBP level threshold 50%, GBP Class Threshold 20%
 - Turned off at GBP level threshold 40%, GBP class threshold 10%
 - Pages are written directly to DASD instead of to the GBP
 - Cross invalidation signals sent to local BPs after DASD write I/O is complete



GBP castout

- GBP castout thresholds are similar to local BP deferred write thresholds
 - Encourage Class castout (CLASST) threshold by lowering its threshold
 - More efficient than GBP castout threshold (notify to pageset/partition castout owner)
 - CLASST threshold checked on by each GBP write
 - GBPOOLT threshold check by GBP castout structure owner based on timer (1 sec default)
 - Default threshold lowered in V8
 - With Db2 11, CLASST can go below 1%



GROUP BP14	QUANTITY	/SECOND	/THREAD	/COMMIT
-----	-----	-----	-----	-----
PAGES CASTOUT	2224.8K	103.28	13.38	4.49
UNLOCK CASTOUT	58868.00	2.73	0.35	0.12
...				
CASTOUT CLASS THRESHOLD	26835.00	1.25	0.16	0.05
GROUP BP CASTOUT THRESHOLD	594.00	0.03	0.00	0.00
GBP CHECKPOINTS TRIGGERED	45.00	0.00	0.00	0.00

	V7	V8/V9/V10/V11
VDWQT (dataset level)	10%	<u>5%</u>
DWQT (buffer pool level)	50%	<u>30%</u>
CLASST (Class_castout)	10%	<u>5%</u>
GBPOOLT (GBP_castout)	50%	<u>30%</u>
GBPCHKPT (GBP checkpoint)	8	<u>4</u>



GBP castout ...

- GBP memory shortages may impact application performance and ultimately become an availability issue
 - When GBP memory fills up, Db2 starts pacing for the GBP writes at Commit which results in slower response times

```
DSNB750I  -DP11 DISPLAY FOR GROUP BUFFER POOL GBP11 FOLLOWS
...
DSNB786I  -DP11  WRITES
           CHANGED PAGES                = 688259863
           CLEAN PAGES                   = 0
           FAILED DUE TO LACK OF STORAGE  = 3630
           CHANGED PAGES SNAPSHOT VALUE  = 44474
```

GROUP BP14	QUANTITY	/SECOND	/THREAD	/COMMIT
CASTOUT CLASS THRESHOLD	26835.00	1.25	0.16	0.05
GROUP BP CASTOUT THRESHOLD	594.00	0.03	0.00	0.00
GBP CHECKPOINTS TRIGGERED	45.00	0.00	0.00	0.00
WRITE FAILED-NO STORAGE	0.00	0.00	0.00	0.00

- After repetitive write failures, page will be put on the LPL
 - If the write failures are against an index, the entire index might be put on the LPL
- Recovery actions are then necessary



GBP castout ...

- As transaction and data volumes grow, the GBP can become stressed
 - Pages written to GBP faster than castout engines can process
 - GBP congested with changed pages
 - Can cause GBP full condition

```
21.02.15 s0012052 *DSNB325A -DP1A DSNB1CNE THERE IS A CRITICAL SHORTAGE OF SPACE IN GROUP BUFFER POOL GBP11
...
21.07.37 s0012052 DSNB327I -DP1A DSNB1CNE GROUP BUFFER POOL GBP11 HAS ADEQUATE FREE SPACE
```

- Problems are often precipitated by update-intensive batch jobs and/or utilities run against GBP dependent objects
 - Intense, sustained GBP page write activity can lead to shortage of GBP memory
 - Automatic GBP ALTER via XES Autoalter can respond and increase allocated GBP size up to SIZE
 - Provided there is sufficient head room



WRITE FAILED-NO STORAGE < 1% of TOTAL CHANGED PGS WRTN

Reduce castout thresholds, and/or
Reduce GBP checkpoint timer and/or
Increase GBP size



GBP Tuning

- Db2 –DIS GBPOOL(*) TYPE(GCONN) GDETAIL(INTERVAL)

```
DSNB750I }-PR4B DISPLAY FOR GROUP BUFFER POOL GBP2 FOLLOWS
...
DSNB757I -PR4B MVS CFRM POLICY STATUS FOR DSNPR0B_GBP2      = NORMAL
           MAX SIZE INDICATED IN POLICY                    = 614400 KB
...
DSNB758I -PR4B      ALLOCATED SIZE                          = 614400KB
...
DSNB759I -PR4B      NUMBER OF DIRECTORY ENTRIES            = 384147
           NUMBER OF DATA PAGES                          = 116010
...
DSNB786I -PR4B  WRITES
           CHANGED PAGES                                  = 2882842576
           CLEAN PAGES                                    = 0
           FAILED DUE TO LACK OF STORAGE                  = 71
           CHANGED PAGES SNAPSHOT VALUE                  = 10642
DSNB787I -PR4B  RECLAIMS
           FOR DIRECTORY ENTRIES                         = 2495178
           FOR DATA ENTRIES                             = 3290663329
           CASTOUTS                                       = 4018446743
DSNB788I -PR4B  CROSS INVALIDATIONS
           DUE TO DIRECTORY RECLAIMS                     = 586960
           DUE TO WRITES                                  = 877165837
           EXPLICIT                                       = 2
```

Note:
Make sure to collect Statistics Class 5 (IFCID 230)
→ Additional GBP stats including Reclaims and XIs

OMPE stores the GBP stats in two different tables:
DB2PM_STAT_GBUFFER
DB2PMSYSPAR_230



GBP Tuning

- GBP size is defined in the CFRM policy

```
STRUCTURE NAME(DSNDB2_GBP1)  
  SIZE(1024M)  
  INITSIZE(512M)
```

- Additionally, specify a RATIO using the ALTER GBPOOL command to indicate ratio of Directory Entries (ENTRIES) to Data Page (Elements)



APAR PH13045 introduces 2 changes for Db2 12 users:

Default value of RATIO: 5 → 10

Limit of RATIO on the ALTER GROUPBUFFERPOOL command: 255 → 1024

NEW



GBP Tuning

- Example based on Db2 –DIS GBPOOL command output

GBPOOL	SIZE (MB)	ALLOC_SIZE (MB)	DIR_ENTRIES	DATA_PAGES	RATIO ⁽¹⁾	SUM VPSIZES + DATA_PAGES
GBP0	150	90	74285	14857	5	34857
GBP1	400	300	318108	45444	7	225444
GBP2	600	550	225969	116010	1.9	356010
GBP8K0	990	990	403011	106257	3.8	456257
GBP8K1	600	500	554987	36998	15	436998
GBP16K0	120	60	15792	3156	5	5156
GBP32K	220	120	48499	3030	16	73030

GBPOOL	% COVERAGE	FAIL_LACK_OF_STG	DIR_RECLAIMS	XI_DIR_RECLAIMS
GBP0	213%	0	0	0
GBP1	141%	0	0	0
GBP2	63%	71	2495178	586960
GBP8K0	88%	4231	35733124	12267527
GBP8K1	127%	0	0	0
GBP16K0	306%	0	0	0
GBP32K	66%	0	0	0



GBP tuning

- Targeted tuning #1: GBP with large number of directory reclaims (but no or minimal writes failures) -> GBP2

– Tuning

- Keep number of data pages the same to avoid aggravating write failures
- Increase structure INITSIZE and RATIO to cover the maximum number of directory entries that could ever be required (1 for each local buffer and 1 for each GBP page)
 - Note: RATIO can be a decimal value with 1 digit after the decimal point (e.g., 5.2)

– Changes required

- New DIR ENTRIES = Sum VPSIZE across all Db2 members + GBP DATA PAGES
- New RATIO = New DIR ENTRIES / GBP DATA PAGES
- New INITSIZE = (GBP DATA PAGES * PAGE SIZE (KB) + New DIR ENTRIES * 410 bytes / 1024) / 1024
- New SIZE = 1.3 to 2x New INITSIZE

↑
The size of a Directory Entry can vary by CF level
For a rough estimate, use 410-430 bytes per entry on CF Level 17



GBP tuning ...

- Targeted tuning #2: GBP with large number of write failures (with or without large number of directory reclaims and XI due to directory reclaims) -> GBP8K0
 - Tuning
 - Reduce CLASST and GBPOOLT to trigger more frequent castout
 - Increase the number of data pages to reduce critical space shortages and write failures
 - Adjust structure INITSIZE and RATIO to still cover the max number of directory entries that could ever be required (1 for each local buffer + 1 for each GBP data page)
 - Note: RATIO can be a decimal value with 1 digit after the decimal point (e.g., 5.2)
 - Changes required
 - New GBP DATA PAGES = 1.3 to 2x GBP DATA PAGES
 - New DIR ENTRIES = Sum VPSIZE across all Db2 members + New GBP DATA PAGES
 - New RATIO = New DIR ENTRIES / New GBP DATA PAGES
 - New INITSIZE = (New GBP DATA PAGES * PAGE SIZE (KB) + New DIR ENTRIES * 410 bytes / 1024) / 1024
 - New SIZE = 1.3 to 2x New INITSIZE



GBP tuning ...

- Simplified tuning when using XES AUTO ALTER
 - Very useful autonomic functionality to simplify GBP tuning
 - Tries to avoid Structure Full conditions
 - Tries to avoid Directory Reclaim conditions
 - Recommendation
 - Set CFRM policy properties
 - ALLOWAUTOALT(YES)
 - FULLTHRESHOLD = 80-90%
 - SIZE = 1.3-2x INITSIZE
 - MINSIZE = INITSIZE
 - Periodically review GBP actual allocations
 - If SIZE is reached, it limits the effectiveness of XES AUTO ALTER
 - Plan to update and activate the CFRM policy with new INITSIZE and SIZE and schedule a REBUILD to increase the structure



GBP tuning ...

- Simplified tuning when using XES AUTO ALTER ...
 - Changes when allocated size reaches SIZE - but no or minimal write failures
 - New DIR ENTRIES = SUM VPSIZE across all Db2 members + GBP DATA PAGES
 - New **RATIO** = New DIR ENTRIES / GBP DATA PAGES
 - New **INITSIZE** = (GBP DATA PAGES * PAGE SIZE (KB) + New DIR ENTRIES * 410 bytes / 1024) / 1024
 - New **SIZE** = 1.3 to 2x New INITSIZE
 - Changes when allocated size reaches SIZE with large number of write failures
 - New GBP DATA PAGES = 1.3 to 2x GBP DATA PAGES
 - New DIR ENTRIES = Sum VPSIZE across all Db2 members + New GBP DATA PAGES
 - New **RATIO** = New DIR ENTRIES / New GBP DATA PAGES
 - New **INITSIZE** = (New GBP DATA PAGES * PAGE SIZE (KB) + New DIR ENTRIES * 410 bytes / 1024) / 1024
 - New **SIZE** = 1.3 to 2x New INITSIZE

A more sophisticated approach would be to look at the peak rate of Changed Pages written to the GBP and calculate the average residency time → tuning target 30-60 seconds

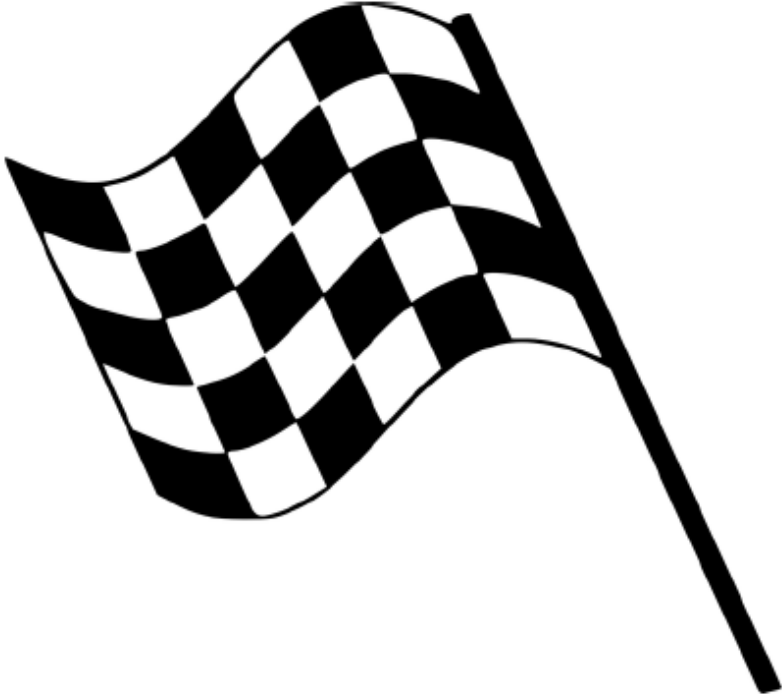


GBP tuning ...

- Increase of local buffer pool size on a 'healthy GBP'
 - Tuning
 - Keep number of data pages the same to avoid aggravating write failures
 - Increase structure INITSIZE and RATIO to cover the max number of directory entries that could ever be required (1 for each local buffer + 1 for each GBP data page)
 - Changes required:
 - New DIR ENTRIES = SUM New VPSIZE across all Db2 members + GBP DATA PAGES
 - New **RATIO** = New DIR ENTRIES / GBP DATA PAGES
 - New **INITSIZE** = (GBP DATA PAGES * PAGE SIZE (KB) + New DIR ENTRIES * 410 bytes / 1024) / 1024
 - New **SIZE** = 1.3 to 2x New INITSIZE



Questions





Thank You



BROADCOM[®]

MAINFRAME SOFTWARE