



DB2 for z/OS – An A-Z of Performance

Phil Grainger
CA

The Challenge

- Applications accessing DB2 will, in general, do one (and only one) of the following at one time:
 - Use CPU
 - Wait for something
 - I/O
 - CPU
 - Locked DB2 Object
 - Latch, drain, dataset operation etc.

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies.



The Solution

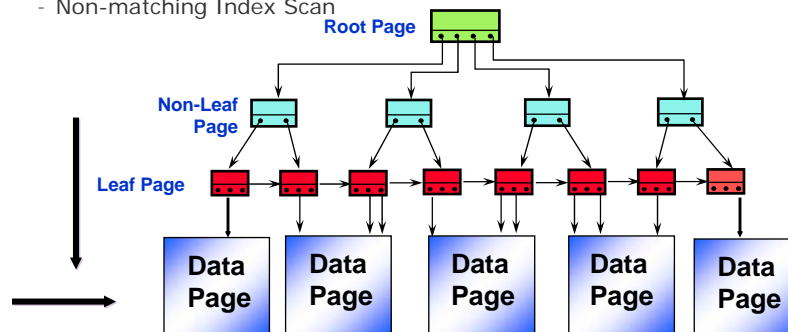
- Tuning is therefore EASY
- STOP all the waits
 - And you will have the best performance possible
-
- Let's take an alphabetical journey through the areas contributing to "wait"

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Access Paths

- Segmented Tablespace Scan
- Partition Scan
- Non-segmented Tablespace Scan
- Matching Index Scan
- Non-matching Index Scan
- Multiple Index Access
- Index Only
- One-Fetch Index



Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Access Paths

- Always changing
 - Hopefully being improved
- e.g.
- Use of indexes in subselect:

```
SELECT FROM TAB1  
WHERE TAB1.COL1 IN  
(SELECT FROM TAB2  
WHERE WHATEVER.....)
```

```
Index on TAB1.COL1  
ACCESSPATH MATCHING  
IN LIST INDEX ACCESS
```

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Access Path Hints

- CAN influence the optimizer's access path selection
- ACCESSPATH column in the PLAN_TABLE
 - Be careful
- Not only can hints be provided as BIND parameters
 - Hint ids may also (in v8) be specified on the statements themselves

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Bind Parameters

- ACQUIRE
 - USE (only value for packages)
 - Acquires IS lock on 1st Use
 - Escalates to IX at 1st Update
 - ALLOCATE
 - Acquires IX lock when application starts
- RELEASE - of Intent TS & T Locks and other resources
 - COMMIT
 - PLAN w/packages in use
 - DEALLOCATE
 - Online app that reuses threads
 - Batch w/regular commits and I/U/D
- OPTHINT
 - You can try and persuade DB2 to take a particular choice of access path

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Bufferpools

- The sole purpose of bufferpools is to reduce I/O
 - Or even eliminate it altogether
- Make them as big as makes sense
- It IS possible to make them too big
 - And even if not TOO big, they can still be bigger than necessary
 - This uses storage that could be utilised elsewhere

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Bufferpool Exploitation

- Bufferpools can be associated with a Dataspace (prior to V8, that is.....)
 - VPTYPE = Primary or Dataspace
 - 32 Gig with 4 K page size
 - Should be backed with ALL Real Storage
 - Small Lookaside buffer still in DBM1 for management
 - This is a VERY good choice for v7 on 64-bit hardware
- Default BP for Tablespace and Index in DSNZPARM
 - Also INDEXBP on CREATE DATABASE
- 8K, 16K and 32K Page Sizes
 - Random vs. Sequential access consideration

NOTE From V8, dataspace bufferpools are NOT supported

- All bufferpools are in the DBM1 address space

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Bufferpools - Read Processing


- Hit Ratio for random access should be watched
- Prefetch engines increased from 300 to 600
- Watch Workload Changes
- Particularly:
 - DM Critical Threshold Reached
 - Prefetch Disabled - No Buffer

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Bufferpools - Write Processing

- VDWT-Vertical Deferred Write Threshold(0 % - 90 %) is for a single dataset
 - Sum of IN-USE and Updated pages
 - Default 10 %, consider Zero or One
 - Can also be a number of pages
 - Useful for very large bufferpools
- Pages Written divided by Asynchronous Writes
 - Max of 32 pages per Write
 - If Low, lower Write Thresholds
- Most frequent writes occur at DB2 checkpoints
- PGSTEAL Option - LRU or FIFO
 - FIFO can reduce Chain Latch Maintenance significantly
 - But can affect re-reading of pages



Bufferpools - A Strategy

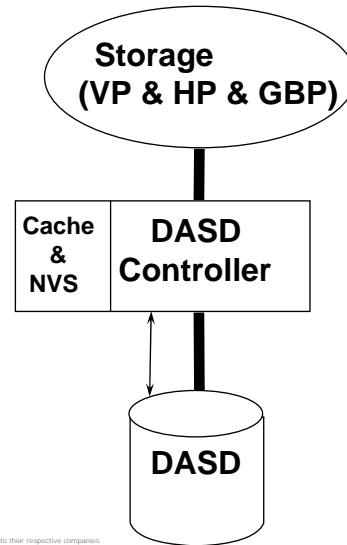
- Separating of production DB2 objects across bufferpools can offer several advantages:
 - This Object/Bufferpool separation delivers isolation
 - DB2 performance improves as objects can be tuned based on processing
- The Basic Plan
 - BP0 - Catalog & Directory
 - BP1 - Indexes
 - BP5 - Tablespace
 - BP7 - DSNDB07(Work Files)
- Expand by separating like-processed objects
 - Random ones
 - Sequential ones

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies.



DASD & Datasets

- Performance of Storage
 - DASD vs. Memory
- Hardware Differences
- Open Dataset Management
 - Uses Multiple Service Tasks
- Recalling datasets DOESN'T!



Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Direct Row Access

- ROWID Data Type
- PRIMARY_ACESSTYPE = 'D'
 - Will use ACESSTYPE if needed
- Could use index on ROWID
- ROWID also used for LOB support

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Dispatching Priority

- Nothing has Changed
 - VTAM
 - IRLM
- (Your monitor goes HERE!)
- DB2 System Services
- DB2 Database Services
- DB2 DDF

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Dynamic SQL

- Dynamic SQL Caching
 - Reuse and sharing of Prepared SQL
 - Parameter CACHEDYN = Yes/No
 - KEEP DYNAMIC(YES)
 - MAX KEPT DYN STMTS is the maximum number of prepared dynamic statements saved
- Predictive Governor
 - Estimation cost for a static or dynamic SQL prior to execution done with Explain

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.




Dynamic Statement Cache - Sizing

- DSNZPARM parameters now allow independent sizing of EDMPOOL items:
- EDMPOOL
 - Size of EDMPOOL itself
- EDMSTMTC
 - How space for dynamic statement caching
 - Default = *EDMPOOL*
- EDMDBDC
 - How much space for DBD caching
 - Default = *EDMPOOL*

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



EDM Pool



- Available Pages are allocated first for new pages
 - Inactive Pages are stolen when needed, If no room, Error returned
 - Contiguous storage limited to 32KB size
 - Addresses EDM POOL OUT OF SPACE
- FREE PAGES IN FREE CHAIN vs. PAGES IN EDM POOL(too Big, Too Small)
 - Increased Response time & I/O for loading of SKCT, SKPT, and DBD
 - Preparation on Caching Dynamic SQL
- Hit Ratio for REQUESTS to NOT IN EDM
 - DBD - Get as close to 100 %
 - Cursor and Package Sections - 85 % to 95 %
- Also, make sure your DBDs are "chunked" (v6)
 - Addresses some of the contiguous storage issues

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



EDMPOOL - Sizing

- DSNZPARM parameters now allow independent sizing of EDMPOOL items:
 - EDMPOOL
 - Size of EDMPOOL itself
 - EDMSTMTC
 - How space for dynamic statement caching
 - Default = *EDMPOOL*
 - EDMDBDC
 - How much space for DBD caching
 - Default = *EDMPOOL*
- These are ALSO separate areas of storage
 - NOT part of one contiguous "EDMPOOL"

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Explain

- Don't forget to always Explain everything
- The PLAN_TABLE is where all your tuning starts
 - A new table, DSN_STATEMENT_TABLE, can also be populated when Explain is run

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Explain STMTCACHE ALL/token

- There is also the DSN_STATEMENT_CACHE_TABLE
 - Which is populated by EXPLAIN STMTCACHE ALL/token
 - This also includes a CLOB column of the actual SQL statement

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Filter Factors & Host Variables

- Proportion of rows for predicate is true
 - COL1 = Literal - FF = 1/25 or .04
- Filter Factor for Uniform Distribution
 - Positive Value in CARDF in SYSCOLUMNS
- Filter Factors for Non-Uniform Distribution
 - Columns in SYSCOLDIST & SYSCOLDISTSTATS
- Reoptimization at Run Time - REOPT(VARS)
 - Addresses impact of Host Variables and Filter Factors
 - Reoptimization occurs when Cursor is opened
 - Look for big differences in Data Distribution

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Indexes

- Each Index causes 4 - 10 percent overhead on Delete and Insert processing
- Utility Impact of Indexes
- Matching Columns VS Index Screening
 - List Prefetch will use Index Screening

INDEXA

```
COL1,  
COL2,  
COL3,  
COL4
```

EXAMPLE # 1

```
SELECT COL5, COL6  
FROM TABLE_1  
WHERE COL1 = :HV  
AND COL2 = :HV  
AND COL3 = :HV  
AND COL4 = :HV
```

EXAMPLE #2

```
SELECT COL5, COL6  
FROM TABLE_1  
WHERE COL1 = :HV  
AND COL3 = :HV  
AND COL4 = :HV
```

- Periodically review indexes for unnecessary ones
 - Each new DB2 gets better at index utilisation

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Hiperpools

- Up to v7, Hiperpools enable bigger bufferpools
 - Effective limit of 1.6Gb for ALL bufferpools in a subsystem
- Unlike dataspace bufferpools, data cannot be read/written directly
- Hiperpools are in effect a bufferpool cache
 - All data flow is always through the bufferpools
 - The saving is all I/O
 - If a page is in the hiperpool, it can be fetched from there
 - Dataspace bufferpools CAN be addressed directly
- Version 8 removes Hiperpools completely
 - V8 limit is 1Tb for all bufferpools in a subsystem

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



ONLY with Bufferpools

- GETPAGE

- Is Page in Bufferpool already?
 - Yes – Return to Application
 - No - Is there a free page in the bufferpool?
 - Yes – Read page into bufferpool
Return page to application
 - No – Free page from bufferpool
Read page into bufferpool
Return page to application

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



With Bufferpools and Hiperpools

- GETPAGE

- Is Page in Bufferpool already?
 - Yes – Return to Application
 - No - Is page in Hiperpool?
 - Yes – Is there a free page in the bufferpool?
 - Yes – Move page from Hiperpool to bufferpool
Return page to application
 - No – Free page from bufferpool
Move page from hiperpool to bufferpool
Return page to application
 - No - Is there a free page in the bufferpool
 - Yes – Read page into bufferpool
Return page to application
 - No – Free page from bufferpool
Read page into bufferpool
Return page to application

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



With Dataspace Bufferpools

- GETPAGE

- Is Page in Bufferpool already?

- Yes – Return to Application

- No - Is there a free page in the bufferpool?

- Yes – Read page into bufferpool
Return page to application

- No – Free page from bufferpool
Read page into bufferpool
Return page to application

- Look familiar??

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



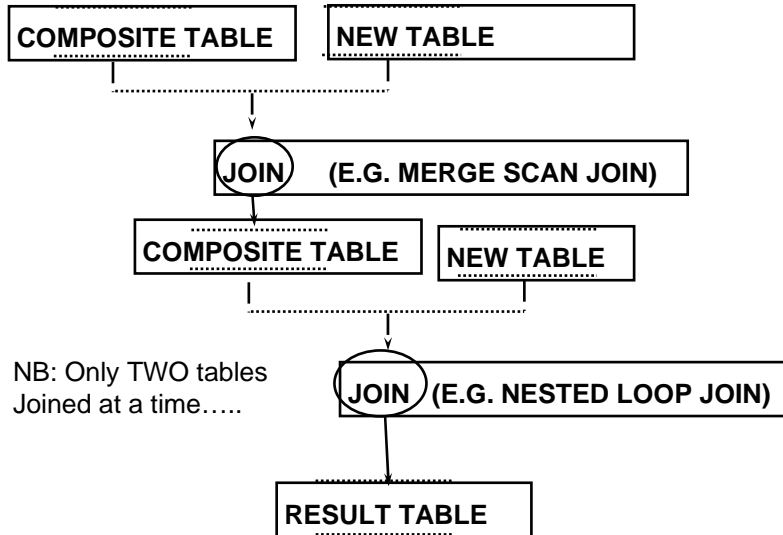
Ready for a Break Yet??



Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Joins



Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Joins

- **SELECT FROM TAB1, TAB2** **INDEX ON TAB2.COL1**
WHERE TAB1.COL1 = TAB2.COL1 **TAB1.COL1 CHAR(8)**
 TAB2.COL1 CHAR(22)

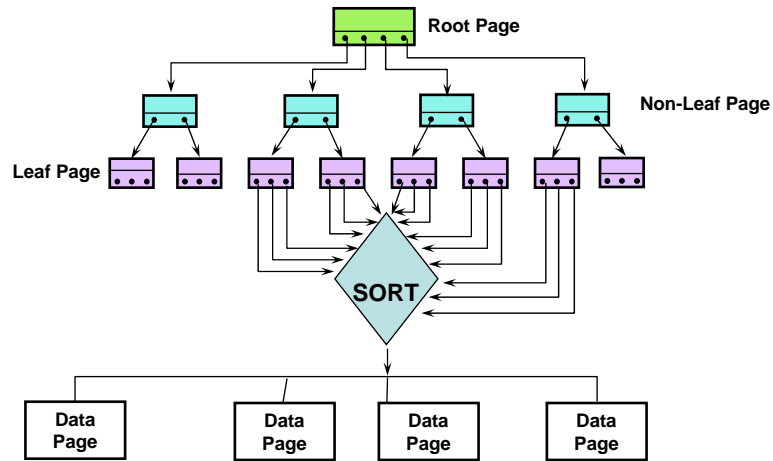
INDEX WILL BE CONSIDERED

- Outer Join and Transitive Closure
 - TAB1.COL1 = TAB2.COL1 AND TAB2.COL1 = 'ABC'
 - THEN TAB1.COL1 IS EQUAL TO 'ABC' as well
- Beware of "too many" tables
 - V8 "most reasonable" vs "best" access path

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



List Prefetch



- RIDPOOL - Concurrent RID Processes X AVG # of RIDS X 2 X 5 Bytes

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Locking

- MAXROWS
- PAGESIZE
- LOCK SIZE
- CURRENTDATA YES/NO
- ISOLATION (UR)
 - See later

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Logging

- Active Log Dataset write time
 - DASD Speed - As fast as 2 ms
 - As always, watch I/O Contention
- UNAVAILABLE OUTPUT LOG BUFFERS
 - Hopefully Zero
 - Size of OUTPUT LOG BUFFER
 - Benefit Recover and Restart
- COMMIT Logic can have impact on Performance(Rollback)
- Fast Log Apply for Recover and Restart
 - LOGAPSTG
 - Suggest minimum of 100M

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Multi-Row Access

- This is a BIG performance boost
 - Can be of the order of 30% less cpu usage
- Of course programs need to change
 - ROWSET cursors and multi-row inserts from arrays
 - GET DIAGNOSTICS
- But the benefits will probably be worthwhile
- DSNTIAUL can now use multi-row fetch

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Page Size

- 4K, 8K, 16K, 32K
 - More Choices
- Bufferpools for Different Page Size
- When Page Size is greater than 4K
 - What is the best page size ?
 - Example: 5K Object
 - If Processed Sequentially - 32 K would be better - 6 objects in 32 K
 - If Processed Randomly - 8 K or 16K might be better
- Also, be careful of table-based partitioning
 - A bigger pagesize can mean potentially more partitions
 - pto

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Page Size and Partitioning

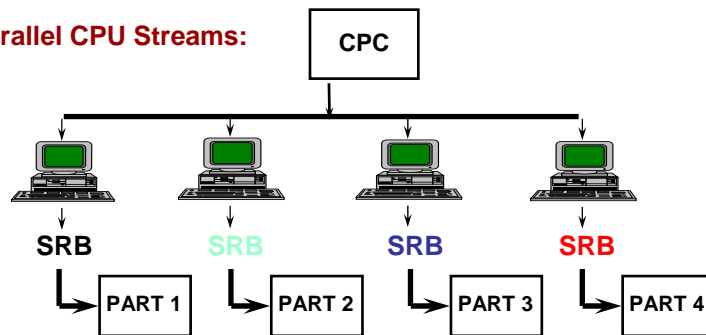
DSSIZE	4K pages	8K pages	16K pages	32K pages
1-4 Gb	4,096	4,096	4,096	4,096
8Gb	2,048	4,096	4,096	4,096
16Gb	1,024	2,048	4,096	4,096
32Gb	512	1,024	2,048	4,096
64Gb	256	512	1,024	2,048

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Parallelism

Parallel CPU Streams:



- PARAMDEG - Upper DEGREE limit
- Controls Too Many Parallel Tasks

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Predicates

- Indexable predicates allow the optimizer to consider using an index when available
- STAGE 1 predicates are evaluated at the time the data rows are retrieved (SARGABLE). Performance advantage in using STAGE 1 PREDICATES because this stage eliminates ROWS passed to STAGE 2 via the Data Manager.
- STAGE 2 predicates are evaluated after data retrieval via the relational (NON-SARGABLE) data services which is more expensive than the Data Manager
- Please don't use "STAGE 3" predicates!

- Oh, and this changes with EVERY new release of DB2

EXAMPLES:

COL1 BETWEEN :HV AND :HV
COL1 NOT BETWEEN :HV AND :HV
COL1 <> :HV

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Predicates

- Identical predicates are processed in sequence starting with EQUAL(=) predicates

```
WHERE COL1 = :HV
AND COL2   = :HV
AND DEPT   = :HV
AND GENDER = :HV
```

1.....

```
WHERE COL1 = :HV
AND COL2   = :HV
AND DEPT BETWEEN :HV AND :HV
AND GENDER = :HV
```

Or 2

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



PLAN_TABLE

- Understand steps within a SQL
- Retrieve rows in proper sequence
- Keep your PLAN_TABLE definitions up to date

```
SELECT * FROM PLAN_TABLE
WHERE APPLNAME LIKE 'PROG%'
ORDER BY TIMESTAMP, QUERYNO, QBLOCKNO,
PLANNO, MIXOPSEQ
```

```
QUERYNO
  QBLOCKNO
    PLANNO
      MIXOPSEQ
```

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Runstats & Catalog Statistics

- Columns in Index that are Related only had Only have FIRSTKEYCARD & FULLKEYCARD
 - New RUNSTATS Parameters for New Stats
 - KEYCARD - Collect cardinality for index columns
 - FREQVAL - Controls collection of most frequent values
 - Columns in SYSCOLDIST & SYSCOLDISTSTATS
 - TYPE - C = Cardinality, F = Frequent Value
 - CARDF - # of Distinct Values, For "C" Type Stats
 - COLGROUPNO - Columns Stats gathered for
 - NUMCOLUMNS - Count of columns
- SAMPLE Keyword for Non-indexed Columns
- Many are Floating Point to cater for large numbers

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



SORT

- Sort record is all selected columns + sort key
 - Limit sort records and size

```
SELECT C1, C2, C3, C4, C5
FROM T1 WHERE....
ORDER BY C2, C4

SORT RECORD - C2, C4, C1, C2, C3, C4, C5
```
- Watch these statistics
 - MERGE PASSES DEGRADED
 - Should be zero
 - WORKFILE REQUESTS REJECTED
 - Should be zero
 - WORKFILE PREFETCH NOT SCHEDULED
 - Should be zero
 - Prefetch Quantity less than 4 or 5

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Sort Work Database

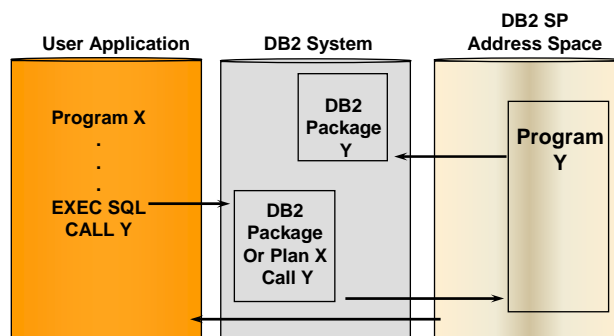
- DSNDB07 or Named DB for Data Sharing
- Minimum of 5 or 1/5 Max # of Partitions
- Placement of Datasets
 - Across Channels and Control Units
 - DON'T have multiple datasets on the same device
- Own Bufferpool
 - VPSEQT Set to 95 to 98
 - Watch DWQT and VDWQT
- New ZPARM SPRMOPT
 - Influences Optimizer to choose Sort

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Stored Procedures

- Addresses the "CLIENT/SERVER" Problem
- Statically Bound Package
 - Centrally Located, Benefits Security Issues
 - CLIENT Program Logic Simplified
- SPAS Under Control of WLM in V5



Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Temporary Database

- Database created with the attribute of AS TEMP
- Used for static scrollable cursors
- Declared Global Temporary Tables
- Size tablespaces carefully
 - In general, work will NOT span these tablespaces

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Uncommitted Read

- Isolation Level of UR
- Reading uncommitted data
 - Statistical Analysis
 - Sources Unknown
 - Don't Care
- WITH UR on a specific SQL
- Significant performance improvement

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Utilities

- RECOVER INDEX with IMAGE COPY
 - 5 - 10 times faster
 - COPY ENABLE on CREATE/ALTER INDEX
- REUSE on REORG and LOAD
 - Skips DELETE/DEFINE
- ONLINE REORG
 - Mapping Table Management Improvement
- COPY & RECOVER Lists
 - TABLESPACE & INDEX
 - From / To DASD
- REORG REBALANCE

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Wait Times

- IN DB2 TIME COMPARED TO IN DB2 CPU TIME
 - How much difference ?
 - What are we waiting for ?
- Service Task Wait Time Broken Down to:
 - Open / Close
 - Dataset Create/Extend/Delete
 - Syslgrnx Downlevel Detection
 - Phase 2 Commit/Abort
- ALWAYS have class 3 turned on
 - Gives much more wait time information

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



zAAPs – z9 Application Assist Processor

- Add on processors for the z9 (and later) families of mainframes
- Java “eligible” workload can run ONLY on a zAAP
 - Or can “crossover” to run on the main CPCs
- Of use only where Java is being used

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



zIIPs – z9 Integrated Information Processor

- Much more interesting to DB2 folks
- Eligible (SRB) workload can be executed on a zIIP instead of on a main CPC
- Minimum requirements of DB2 v8 and z/OS v1.6
- The eligible workload tends to be parts of distributed code
 - Requires special types of SRB processing
 - Eg as in DB2 Connect
- zIIP capacity does NOT add to MIPS count
 - Number of zIIPs cannot exceed number of CPCs

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



dsnZparm

- Many performance related items can be set in DSNZPARM
- We've seen the EDMPOOL sizing
- Maximum degree settings

- It's a shame that some of these aren't easier to change
- Still, at least we have –SET SYSPARM now

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Summary

- Remember Applications accessing DB2 will do one(and only one) of the following at one time:
 - Use CPU
 - Wait for something
 - I/O, CPU, Locked DB2 Object
- Watch SQL related areas
- Keep current on maintenance
- Manage and monitor changes
- Keep current on what YOUR version of DB2 can do!

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.



Phil Grainger

CA

phil.grainger@ca.com

Copyright ©2006 CA. All rights reserved. All trademarks, trade names, services marks and logos referenced herein belong to their respective companies.

