

DB2 11 Charts My Top 20 Features

Andy Ward – Senior Principal Consultant

December 2014



Agenda

- Countdown of my top 20 features
 - 20 to 1
- Provide a brief overview of each capability
 - Not enough time to cover each feature in detail
 - Some of them could take the hour alone!
- Will you agree with my number 1 feature?

Large Storage Frame Sizes...1

Number 20

- What benefit do I get from large frame sizes?
 - TLB misses reduced
 - Potential CPU reduction
 - Now pageable (when using Flash Express)
- LFAREA Parameter
- 1MB page frames now supported for
 - DB2 execution code
 - z/OS 2.1 + Flash Express required
 - Log Buffers
- 2GB Frame Support (OA40967)
 - Requires z/OS 1.13 and EC12

Large Storage Frame Sizes...1

Number 20

- Large Frames and BPs
 - 1MB frames are now used by default if defined for BPs
 - Page fixing no longer required when using Flash Express (and >z/OS 1.13)
 - Useful if real storage availability is an issue
 - Available in DB211 CM
 - Retrofit to DB2 10 – PM85944 & PM90486
 - IBM tests showed a 2.1% CPU decrease
 - Mainly down to less TLB misses

- As always, paging is never a good thing
 - Monitor behaviour after any changes

Accelerator STATIC SQL Support

Number 19

- Prior to Analytics Accelerator 4.1 only dynamic SQL was supported
- As for dynamic, only cursor queries and INSERT FROM SELECT are offloaded
- New BIND/REBIND options
 - QUERY ACCELERATION NONE | ENABLE | ENABLE WITH FAILBACK | ELIGIBLE | ALL
 - Default is not specified
 - GET ACCEL ARCHIVE NO | YES
 - Default is not specified

Accelerator STATIC SQL Support

Number 19

- New table SYSACCEL.SYSACCELERATEDPACKAGES

Optimiser Cost Adjustments

Number 18

- It isn't only statistics that cause EXPLAIN differences between environments
- In DB2 11 the optimizer has reduced sensitivity to CPU speed
- NLEVELS sensitivity also reduced
- DB2 11 new features added to the cost model
 - Greater Sparse Index support
 - Index duplicate skipping
 - Correlated sub query cache

Partition Scalability Improvements

Number 17

- Version 8 provided us with a maximum of 4096 partitions
 - Likely to increase in later versions
- RELEASE(COMMIT) sensitive to number of partitions
 - Irrespective of how many were actually accessed
 - Problem evident for both UTS and non-UTS
 - The more partitions the greater the problem
 - Noticeable at 200+ partitions
- DB2 11 resolves this issue
 - CM no REBIND required

DGTTs – NOT LOGGED Support

Number 16

- Previously only CGTTs didn't log
 - But also didn't support indexes
- In tests using 5M rows
 - Significant CPU savings 20%
 - INSERTing/DELETEing/UPDATEing
 - 60% elapsed time reduction when ROLLBACK is called for both when PRESERVEing or DELETEing rows

VIEW Temporal Support

Number 15

- In DB211 VIEWS now support temporal constructs

```
– CREATE VIEW v0 (EMPNO, SALARY, COMM) AS SELECT  
  EMPNO, SALARY, COMM FROM EMP_TEMPORAL ;  
– SET CURRENT APPLICATION COMPATIBILITY =  
  'V11R1' ;  
SELECT * FROM v0  
FOR SYSTEM_TIME AS OF TIMESTAMP '2013-08-25  
23:55:00' ;
```

Group Buffer Pool Write Around

Number 14

- Heavy updaters and utilities can result in sustained GBP write activity
 - Application slowdowns, worst cases pages written to the LPL
- DB2 11 helps to alleviate this issue using write around
 - Writing pages directly to DASD
 - Still using the GBP to send cross-invalidation requests
- Only for deferred writes (asynchronous)
 - COMMIT and index splits are synchronous writes and not eligible
- GBPOOLT and CLASST used to invoke Write Around
- Significant elapsed time improvement
 - Slight increase in CPU time

Group Buffer Pool Write Around

Number 14

- DISPLAY GROUPBUFFERPOOL(GBP0) **MDETAIL**
- z/OS 1.12, CFCC level 17 or 18, >z196, DB2 11 CM

VCAT Translation for RESTORE SYSTEM

Number 13

- DB2 10 Log Apply restriction when using RESTORE SYSTEM in the same SYSPLEX
 - Making cloning problematic
- RESTORE SYSTEM
 - SWITCH VCAT
 - SYSVALUEDDN(ddname)

Package Based Continuous Block Fetch

Number 12

- SQL Based Continuous Block Fetch
 - Single connection
 - No other SQL can use the connection whilst the cursor is being fulfilled
 - Application changes required

- Package Continuous Block Fetch
 - Secondary connection for each R/O cursor
 - Dedicated to the cursor
 - Data flows until the cursor is exhausted or CLOSEd
 - No application changes
 - Just BIND with DBPROTOCOL(DRDACBF)
 - Big CPU and elapsed improvements
 - Tests show up to 78% on the server and 20% on the requester

ALTER TABLE DROP COLUMN

Number 11

- Tablespace must be a UTS
- Long list of restrictions
- Pending definition change
 - ALTER does not have an effect at execution time
 - An entry is added to SYSPENDINGDDL
 - AREOR state is flagged

Increased zIIP Usage

Number 10

- IBM Utilities
 - IBM testing shows 14-45% additional eligibility
 - Inline statistics collection
 - LOAD, REORG & REBUILD INDEX
 - RUNSTATS COLGROUP collection
 - LOAD REPLACE index management

- System Tasks
 - DB2 10 made prefetch reads and deferred writes zIIP eligible
 - DB2 11 adds asynchronous log reading and writing, castout processing, XML cleanup and pseudo delete index processing



Extended RBA/LRSN

Number 9

- 6 Byte RBA/LRSN Primer

- RBA

- 256TB addressing capacity
 - Problems with logging capacity being exhausted

- LRSN

- TOD Clock value
 - Up to the year 2042
 - But Deltas during data sharing conversion may cause the LRSN to run out much faster than this (operating in the future)
 - Granularity of 16 microseconds
 - Caused CPU ‘spinning’
 - Some relief for this in DB2 9 & 10

Extended RBA/LRSN

Number 9

- New 10 byte RBA/LRSN
 - DB2 11 always displays 10 byte addresses whether conversion has taken place or not
 - RBA
 - 1YB of addressing capacity (10^{24})
 - Should take decades to exhaust even with increasing processor speeds
 - LRSN
 - Supports over 30000 years

Externalising RTS Statistics On Demand

Number 8

- RTS are updated every 30 minutes by default
- An extension to the ACCESS DATABASE command can drive on demand externalisation
 - ACCESS DB(DB) SP(TS) MODE(**STATS**) PART(n)
- Useful for
 - Heavily changing objects
 - Understanding object usage (CRUD analysis) for effective design

Reduced Need for Reorg...1

Number 7

- Pseudo Deleted Index Keys
 - Exclusive access required to the index page to clean these up as part of a DELETE
 - Can result in more GETPAGEs and LOCK requests
 - RID reuse can cause DEADLOCKS
 - REORG INDEX required to clean up

- DB2 11 delivers an automated cleanup function
 - zIIP eligible
 - RTS used to identify candidates
 - User control over
 - Number of threads to execute (maximum 128)
 - Which objects are eligible (SYSINDEXCLEANUP)

Reduced Need for Reorg...2

Number 7

- Indirect Reference Avoidance

- Prior to DB2 11 the solution involved MAXROWS and FREESPACE
- DB2 11 provides an autonomic solution involving RTS

- New SQL construct, catalog columns and ZPARM

- REORGUPDATESIZE in SYSTABLESPACESTATS
- `CREATE TABLESPACE TS1 FREEPAGE 0 PCTFREE 20 FOR UPDATE -1`
 - -1 causes DB2 to learn the best value (a hard value can be used too)
 - PCTFREE used by LOAD & REORG FOR UPDATE INSERT, LOAD & REORG
- PCTFREE_UPD ZPARM (AUTO, like -1, or value)

Array Data Type

Number 6

- DB211 introduces the capability to define array data types
- Ordinary arrays
 - `CREATE TYPE INTARRAY AS INTEGER ARRAY[100];`
- Associative arrays
 - `CREATE TYPE CHARARRAY AS CHAR(10) ARRAY[VARCHAR(10)];`
- Built-in functions to control arrays
 - `ARRAY_DELETE, ARRAY_FIRST/LAST/NEXT/PRIOR, CARDINALITY, MAX_CARDINALITY, TRIM_ARRAY`

Overriding Predicate Selectivity

Number 5

- DB2 uses default Filter Factors to estimate certain predicate types
- DB2 11 allows these to be overridden on a per statement and related predicate basis
 - Selectivity instance (up to 20)
 - Defined with a weighting
 - Multiple instances can be grouped and are known as a 'selectivity profile'
- Uses 3 EXPLAIN tables
 - DSN_USERQUERY_TABLE, DSN_PREDICAT_TABLE, DSN_PREDICATE_SELECTIVITY

Identifying Missing Statistics

Number 4

- Opening up the Optimizer a little more
- What statistics should I collect?
 - DB2 11 allows you to see what the Optimiser didn't find
- New catalog table `SYSSTATFEEDBACK`
 - Populated asynchronously via `BIND`, `REBIND` & `PREPARE`
 - Cleaned up by `RUNSTATS` as statistics are collected
- New `EXPLAIN` table
 - `DSN_STAT_FEEDBACK`

Global Variables

Number 3

- Prior to DB2 11 applications had to find creative ways to pass information between SQL statements
- Global variables can now be created and referenced
 - `CREATE VARIABLE BATCH_START_TS TIMESTAMP
DEFAULT CURRENT TIMESTAMP;`
 - `SELECT BATCH_START_TS, CURRENT TIMESTAMP FROM
SYSIBM.SYSDUMMY1;`
- The scope is the SQL connection
- New catalog tables
 - `SYSVARIABLES`, `SYSVARIABLEAUTH`, `SYSVARIABLES_TEXT` (auxilliary table) for `DEFAULTTEXT` column

Data Volumes are Increasing at an Alarming Rate

1 THE RAPID GROWTH OF GLOBAL DATA



The production of data is expanding at an astonishing pace. Experts now point to a 4300% increase in annual data generation by 2020. Drivers include the switch from analog to digital technologies and the rapid increase in data generation by individuals and corporations alike.

■ Size of Total Data ■ Enterprise Created Data
■ Enterprise Managed Data

2020: MORE THAN 1/3 OF THE DATA PRODUCED WILL LIVE IN OR PASS THROUGH THE CLOUD.

2012: CUSTOMERS WILL START STORING 1 EB OF INFORMATION.



Embracing Hadoop...1



Number 2

- Interfacing with IBM's BigInsights platform
- Read data stored in HDFS
 - JAQL_SUBMIT

```
▪ Select Jaql_submit('read(hdfs("receipts.json")) ->
transform { seller: $.issuer.id,
subtotal: $.subtotal } -> group by $seller =$seller
into {$seller, revenue: sum($.subtotal)} -> write
(del(location='/test1.csv'))', `
'http://kea.svl.ibm.com', `TIMEOUT=60 USER=scott
PASSWORD=tiger`)
From SYSIBM.SYSDUMMY1;
```

Embracing Hadoop

Number 2

- Read from BigInsights and populate DB2 object

- UDF – HDFS_READ

- ```
SELECT TX.*
FROM
TABLE(hdfsRead('http://172.16.134.134:8080/data/sample.csv') AS TX (COL_1 VARCHAR(35),
COL_2 CHAR(4),
COL_3 INTEGER));
```

- JSON & JAQL

- zDoop (Hadoop running on zLinux)

- Partnership between IBM and Veristorm

# Transparent Archiving...1

## *Number 1*

- DB2 10 provided STTs
  - Providing archiving in a temporal setting
  - A table cannot use both types of archiving now available
- DB2 11 delivers automated DELETE archiving for all (non STT) objects
  - Define two tables x and y
  - ALTER TABLE x ENABLE ARCHIVE USE y
  - ALTER TABLE x ADD COLUMN c will be replicated on table y
- Archiving controlled by two Global Variables
  - SYSIBMADM.GET\_ARCHIVE (Y or N)
  - SYSIBMADM.MOVE\_TO\_ARCHIVE (Y, N or E)

# Transparent Archiving...2

## *Number 1*

- No need to change existing SQL
- Works with dynamic SQL
  - Set GET\_ARCHIVE prior to EXECUTE IMMEDIATE or PREPARE
- Archiving can be disabled via an ALTER
- Significant performance improvement over DELETE TRIGGERS

# Summary

- Did you agree with me?
  - Favourites are pretty subjective for different people and different roles
- DB2 11 has added some rich and varied functionality
  - Across the spectrum, from performance to SQL to availability and everything in between
- Hopefully you have seen some things to whet your appetite
- Happy tinkering!





---

**Andrew Ward**  
Senior Principal Consultant  
andrew.ward@ca.com

