

Bart Steegmans

Db2 for z/OS Performance – March 2019

Db2 for z/OS (Maintenance) Update



Agenda

- Phil focused on What' new in Db2 12 since GA
- This (short) session focussed on
 - What is new besides continuous delivery in Db2 12 ?
 - What is new in Db2 11
- Limited to the highlights since the last GSE meeting-Oct 2018

What is new besides continuous delivery?



APAR Rating - ZSAx keyword

- New keyword added to PE and HIPER APARs
- In addition to other keywords that are being used
- Introduced to provide a “rating” to try and help customers determine when to apply a particular HIPER/PE PTF
- The ZSA keyword is under APAR Error Description, usually as an 'Additional Keyword'
- ZSA ratings, ZSA1 (low), ZSA2, ZSA3, ZSA4, ZSA45, ZSA5 (highest) described below

APAR Rating - ZSAX keyword – Per Type

- System Outage: 4.5 and 5 is an example for adding the consideration for the number of customer who has hit the problem, it should be the same for other HIPER category
 - 4: system outage should automatically get a 4
 - 4.5: If there are 1-5 customer already hit the problem
 - 5: If there are more than 10 customer already hit the problem
- Data Loss:
 - 4: non-recoverable or pervasive, common
 - 3: recoverable, incorrout output, but with few conditions
 - 2: recoverable, incorrout, but fairly rare to hit it
 - 1: super rare cases
- Function Loss:
 - 4: pervasive causing application outages
 - 3: likely common
 - 2/1: rare
- Performance:
 - 4: looping indefinitely
 - 3: degrades $\geq 5\%$
 - 2: degrades $< 5\%$
 - 1: not noticeable
- Pervasive:
 - 4: automatically
- MSysplex:
 - 4: automatically

Function level **independent** enhancements – 1

- Applies to V11, V12, V11+ V12
- PH00637 – DDL and concurrent dynamic SQL
- PI98649 – Updates to REST Services (including versioning)
- PH00572 – Suppress warning messages in DSNTEP2 etc. TOLWARN(QUIET) suppress FETCH messages if ≥ 0
- PH01728 – New DSNZPARM FLASHCOPY_XRCP
- PH02437 – DSNB233I warning to convert to extended LRSN
- PH03263 – New Timestamp format support
- [PH05193] – [Async Cross-invalidation (XI) support]
- PH05030 – zHyperlink Write for Db2 active logs

Function level **independent** enhancements – 2

- **PH01123** – zHyperlink Read support (retrofit to V11)
- **PH00317** – New DSNZPARM to control NOPAD default for REORG UNLOAD EXTERNAL or DISCARD
- **PH02362** – DSNREXX support for SET CURRENT APPLICATION COMPATIBILITY register
- **PH03678** – New RUNSTATS keyword for DSNZPARM STATCLGSRT (consider also PH09074)
- **PH04336** – SQLCA in MQSeries LISTENER if SQLcode ≠ 0 (instead of just SQLCODE)
- **PH06082** – RLF enhancements, e.g. monitor update activity on the active RLF tables and automatically refresh the in-memory RLF tables (consider also PH07658)
- **IRLM+Db2** - IRLM deadlock detection enhancement

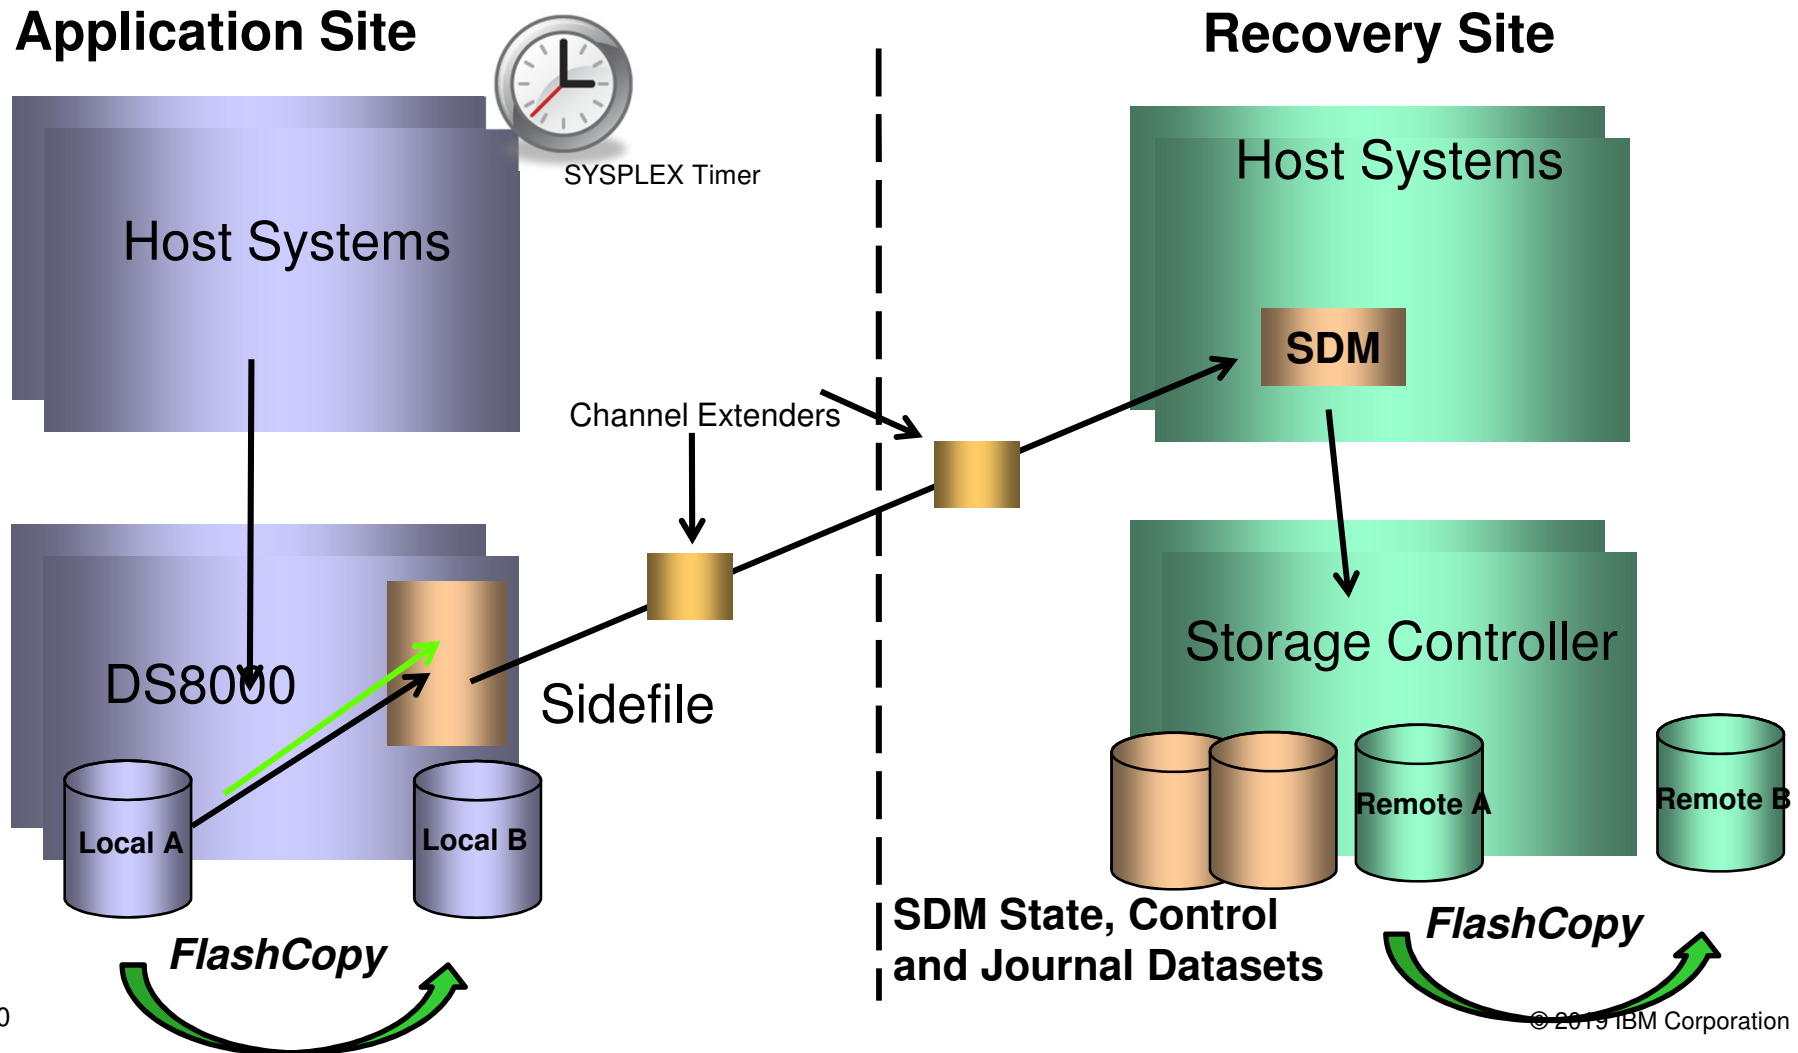
DSNTEP2, DSNTIAUL etc.

- TOLWARN(NO), the current default, can result in incomplete output (or none at all) for a query that gets an unexpected SQL warning from OPEN or FETCH. It's easy to miss this outcome because the run terminates with rc=4.
- TOLWARN(YES) allows processing of a query to continue after an unexpected SQL warning from OPEN or FETCH. (but still produces the warning msg). The run terminates with rc=4 unless something more severe is encountered.
- **TOLWARN(QUIET) is like TOLWARN(YES) except that diagnostic messages for unexpected SQL warnings from OPEN and FETCH are suppressed.**

Some FlashCopy terms

- FlashCopy target can now be in a XRC relationship (global mirror)
 - DS8882F
 - OA55068
 - PH01728
 - DSNZPARM FLASHCOPY_XRCP generates FCTOXRCP keyword in DFSDS command: COPY, REORG, LOAD, REBUILD, RECOVER, CHECK (SHRLEVEL CHANGE)
 - New RACF profile STGADMIN.ADR.COPY.FCTOXRCP
 - FCTOXRCP is not compatible with FASTREPLICATION(NONE): Hence appropriate DSNZPARM settings, e.g. REC_FASTREPLICATION=NONE
 - BACKUP/RESTORE SYSTEM
 - New definitions in the COPYPOOL, e.g.
 - FRBACKUP to XRC Primary Volumes allowed ... (N/Y/blank)

XRC System – Remote Pair FlashCopy (RPFC) feature



Newly supported timestamp string formats

- The fix for APAR PH03263 provided support for the new timestamp string format shown below (shown without and with time zone information):

`yyyy-mm-dd hh:mm:ss`

What is new is this dash - before, if you wanted colons in the time part, there had to be a blank (versus a dash) between the date and time parts of the timestamp string

`yyyy-mm-dd hh:mm:ss+th:tm`

- New format can be used whenever a string needs to be converted to at timestamp

Asynchronous XI (cross-invalidation)

- Allows XI requests to process asynchronously for cache structures
 - Reduce transaction elapsed time
- Exploiters like Db2 must “synch up” before committing a transaction to ensure that all outstanding XI notifications have been delivered
- New message (DSNB818I) and IFCID fields to reflect this new processing

Asynchronous XI – How to get it

- Db2 requirements
 - PH05193 (function level independent)
 - Data pages in table space not using row-level locking
 - Invalidation is done when writing updated pages during **commit**
 - At this time - enabling by explicit service request only

- z/OS and hardware requirements
 - OA54688 (since 2.2)
 - CFLEVEL 23
 - Long CF link distances are required to get benefit

zHyperLink

- What is the purpose:
 - Make synchronous I/O operations much faster: Reduce to ~20-40µs
- How to get it:
 - **z14 with zHyperLink Express adapter (FC #0431) installed**
 - zHyperLink Express uses optical cable
 - DS888x with I/O bay planar board and firmware level 8.3
 - z/OS 2.1: OA50653 + more (e.g. OA51450, ...)
 - Db2:
 - PI82575 (V12) / PH01123 (V11): Data base synch read I/Os
 - PH05030: Log writes for Db2 active logs
 - Set new ZPARM ZHYPERLINK:
 - Turning off: DISABLE
 - Turning on data base read: DATABASE
 - Turning on log write: ACTIVELOG
 - Turning on data base read & log write: ENABLE

zHyperLink – Additional considerations

- Eligible if
 - Synchronous data base read I/O and synchronous log write I/O
 - CPU will spin similar as in CF requests
 - Distance between z14 and DS888x <150m
 - If distance is longer, then I/O processing as usual
 - Data is in the disk controller cache
- How to plan for eligible workload:
 - zBNA tool (System z Batch Network Analyzer tool) at ASCB level
 - Application level via accounting class 3
- Consider: zHyperLink write support requires zHyperWrite support in PPRC environment

Additional RUNSTATS enhancements

- Improved RUNSTATS COLGROUP performance
 - Statistics are collected from the index, if matching
 - Up to 15% elapsed time, up to 25% CPU reduction
 - Avoid data sorts for single column COLGROUPs
 - New keyword STATCLGMEMSRT (PH03678) to specify sort usage
 - 0: indicates that a data sort will always be done
 - 1-4096: memory in MB for the utility to use to avoid the data sort
 - Default by ZPARM STATCLGSRT
 - Default: 10
 - That ZPARM was introduced by Db2 11: PI74408; Db2 12: PI76730

Db2 12 RLF enhancement

- Monitor update activity on the active RLF tables and automatically refresh the **in-memory RLF tables** if something has changed
- Allow the `-START RLIMIT ID =xx` command where `xx` is the suffix of the table currently in use (active), in-memory table will be updated to reflect table changes
- Support the `SCOPE(GROUP)` option on the `-START RLIMIT` command so that the command can take effect on all members of the data sharing group

IRLM–Db2 Deadlock Detection Enhancement

- As the number of waiters increase on multiple members, especially when those waiters are for a few resources, the deadlock WAIT GRAPH becomes very large and complex.
- This results in requiring large amounts of CPU time to complete while the IRLM holds the main latch, preventing other requests from processing, often resulting in workload stall or slow down.
- This enhancement speeds up timeout processing to reduce the long list of waiters on a single resource lock
- IRLM APARs PH08427 , PH08429 , PH08430 & PH08431 (single PTF UI61778)
- Db2 APAR PH08708 (PTFs: UI61736(v12) / UI61737(v11))