Application DBA Maintenance Activities

Discussion panel

Disclaimer

- *Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.*
- *The info provided is in this presentation on a best effort basis. Although extensively verified it is not guaranteed to be correct nor complete.*

Applicication DBA Maintenance Activities

Back-up

Reorg, Statistics collection, (re)bind

Triggering/scheduling/running utilities

Schema changes

Ad hoc recovery

Replies from panel members on the overview foils in this color:

- **KBC**
- **Volvo**
- **IS4F**
- **BNPPBFortis**

Backup Process

Backup process at KBC

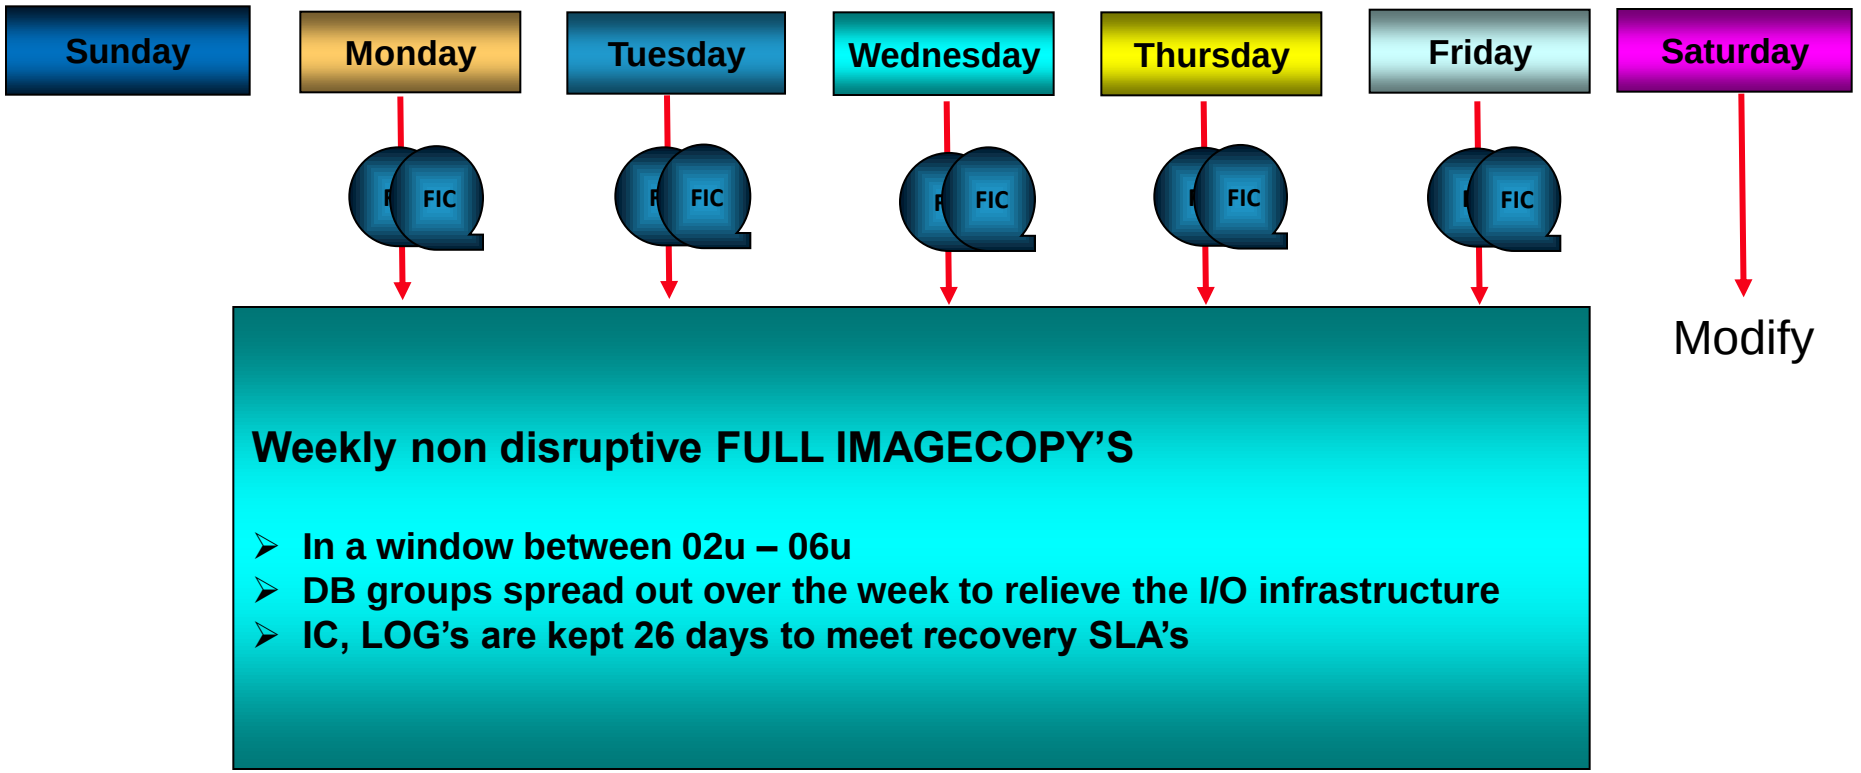

Backup process at KBC

- DB2 Catalog and Directory: daily instead of weekly
- All TS (partition level), except TS registered in COPY_EXCL table
- No Index copies
- Backup is kept on DASD for 24 u, managed by DFSMS
- PPRC mirror to other DC + DFHSM backup of each copy
- Utility jobs are generated via KBC application and scheduled in batch window (run history is kept in a DB2 table)
- Daily check: are all TS recoverable?
- Modify weekly to clean up SYSCOPY

KBC

BNPPBFortis Volvo
IS4F KBC

Image Copy

Frequency	%DB2 data All = 100%	Type •Reference •Change	•Concurrent copy •Incremental •IC of indexes •Flashcopy	Availability •Online •Batch window	Retention period	Reasons Policy	
Several a day	1	C	Concurrent copy	Online	15	Catalog	Every 6 hours
Daily	20	Change	Full		35 days	Criteria	
	40	C	Concurrent copy	online	15	> 1 update	
	Depends	Change	-	Online	70 days	RTS based	
	DB2 catalog and directory	Change	We only use DB2 Full Image Copy of table spaces (partition level). PPRC mirror to other DC.	Batch Window	26 days	Recovery SLA= 14 days ; Daily check: is all data recoverable.	Modify recovery is run weekly to clean up syscopy.
Every 7 days	40	C	Concurrent copy	online	15	No update	
Weekly	100	Change	Full		35 days		
	20		Concurrent copy	Batch windows	15	Load	
	All	Change	-	Online	70 days		
	99,9% (in principle everything, but some exceptions)	Change	We only use DB2 Full Image Copy of tablespaces (partition level). PPRC mirror to other DC.	Batch Window	26 days	Recovery SLA= 14 days ; Daily check: is all data recoverable.	Modify recovery is run weekly to clean up syscopy.
Monthly	100	Reference	Full		35 days	Refresh	

Backup process - Discussion items

- Almost always shrlevel change image copies
  - Do you use QUIESCE after the IC to establish point of consistency ?
- C+D image copied more frequently than other objects (every 6 hours, daily)
- Concurrent copy (DFDSS copy ?) by IS4F
  - Why preference for concurrent copy ?
- No flashcopy image copies being used
- BACKUP SYSTEM not used
- Nobody is doing FIC of indexes
- No incremental image copies
- Varying retention periods
  - Are they based on SLAs ?

Backup – V11 Enhancements

Partition-level inline image copy

- Faster partition-level recovery from inline image copy
- Create partition-level inline image copies if using TEMPLATE with &PA or &PART
  - No new option or keyword on REORG
  - PM93611 (V11):
    - Support substring notation with &PA as long as substring ensures uniqueness
    - Support writing to tape as long as STACK YES not specified

- **RECOVER of single partition of a 20 partition table space**
  - ET reduced by 28%
  - CPU reduced by 49%

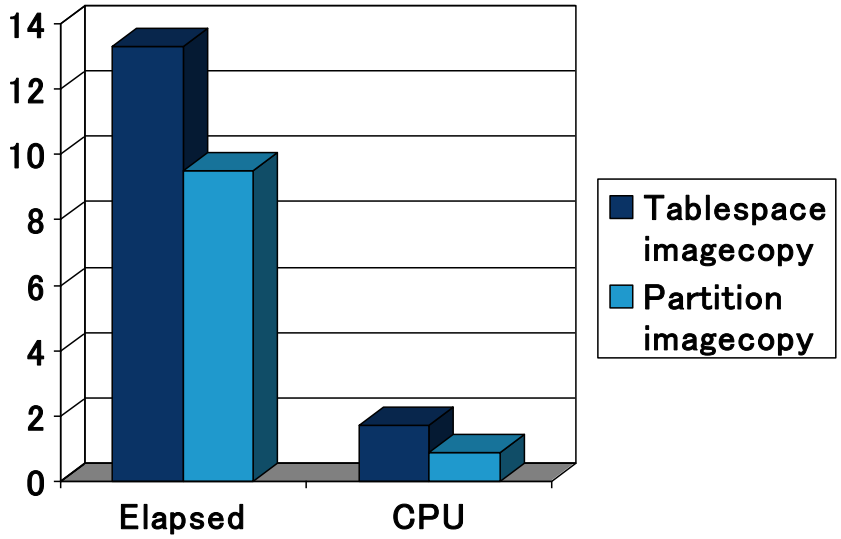

REORG Process

Reorg process at KBC

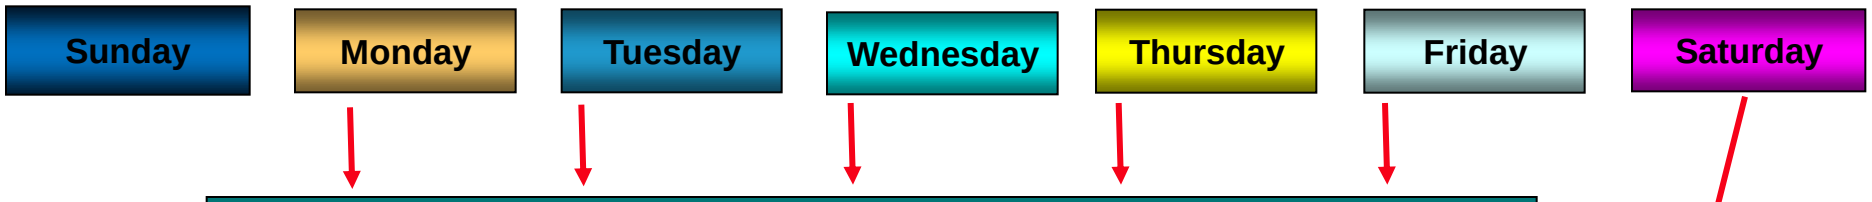

Weekly REORG procedure

- Non critical DB2 objects
- Based on RTS criteria
- Run in a window 02u – 06u
- DB groups spread out over the week to relieve the I/O infrastructure
- SHRLEVEL CHANGE REORG

Release REORG procedure

- Critical objects
- Based on RTS criteria
- During 'down-window' (application releases: 22u – 06u)
- SHRLEVEL CHANGE REORG

Reorg process at KBC

- RTS reorg criteria checked once a week:
  - clustering < 75%
  - relocated rows > 10%
  - non optimal leafs > 20%
  - unclustered inserts > 30%
  - pseudo deleted rids > 5%
  - extents > 15
  - indexlevel increase
- All TS an IX, except objects registered in REORG_EXCL table
- Inline stats
- Utility jobs generated via KBC application and scheduled in batch window (run history is kept in a DB2 table)

KBC

Reorg process at KBC

- REORG runs day after backup (recovery time, max 24 u log apply)
- DB2 Catalog and Directory: +/- once a year

BNPPBFortis Volvo
IS4F KBC

Reorg

Frequency	% DB2 data All = 100%	Type Reference Change	IX only	Availability requirements Batch window	Reasons Policy
Daily					
Weekly	30%	Change	YES	Weekly Maintenance window	Criteria
	20	Ref	No	Batch windows	performance
Monthy					
Yearly					
Statistics driven	2	Share	No	Online	DBA trigger
	ALL daily	Change	No	Online	RTS based
	95%	Change	Yes also	Batch window; Some critical TS only in UOC window	Criteria: 1) clust > 75% 2) relocated rows > 10% 3) non optimal leafs > 20% 4) unclusterd inserts > 30% 5) pseude deleted rids > 5% 6) extents > 15 7) indexlevel increase

Reorg – Discussion items

- No SHRLEVEL NONE REORG
- Mostly SHRLEVEL CHANGE
  - No problems with switch phase ?
- RTS driven in many cases

- Why Reorg ?
  - Just daily/weekly/monthly
  - For peformance ?
- Catalog and Directory Reorg
  - Good before migration to DB2 10 NFM
    - ENFM process likely to be faster
    - Clean up potential ‘corpses’ before migration might hit them
- REORG LOB

  - Peformance ok ?

Reorg – V11 Enhancements

Improve performance of part-level REORG with NPSIs

- New option to defer shadow index build until all keys passed through sort
- New parm keyword (SORTNPSI) and the zparm (REORG_PART_SORT_NPSI) to govern the behavior
  - AUTO/ENABLE/DISABLE options
- Retrofit to DB2 9 & 10 in PM55051
- Result:
  - Customer test of REORG of 40% of partitions showed 55% ET reduction & 22% CPU increase
  - DB2 Sort gives additional ET reduction & cuts CPU to less than original starting point

SWITCH phase impact relief – reduced application impact

- Easier drain acquisition
- Prevent new claims on all target partitions whilst waiting for drains
  - Faster drain acquisition for part-level REORG
- New DRAIN_ALLPARTS option to momentarily drain all data parts
  - Eliminates claim-drain “deadlocks” for part-level REORG with NPSIs
- Restructure SWITCH phase processing for outage reduction
  - SWITCH phase ET reduction of 91% measured when reorging 20 parts
  - Retrofitted back to DB2 10 via PI09303

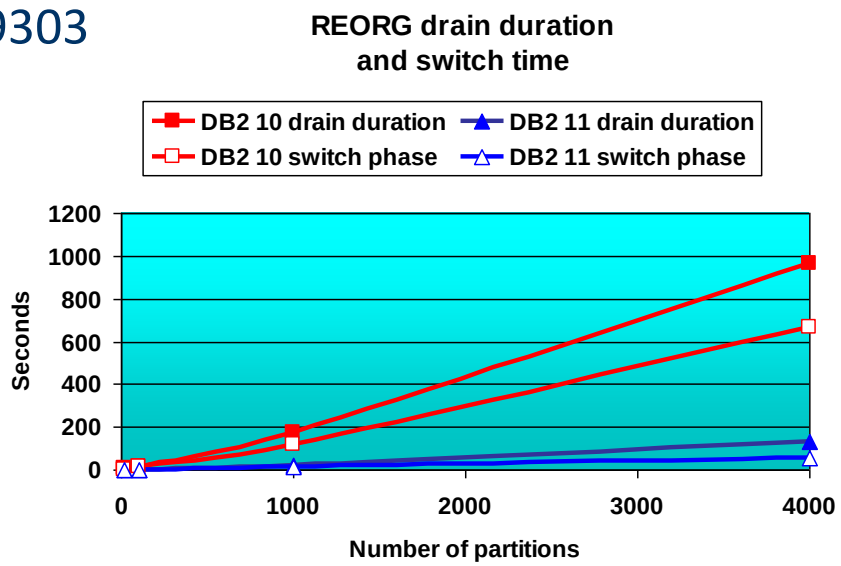

Timing of SWITCH phase with MAXRO DEFER

- Govern timing of drain and switch for long-running REORGs without the need to schedule separate –ALTER UTILITY command
- New SWITCHTIME parameter to determine earliest point at which drain processing will be attempted

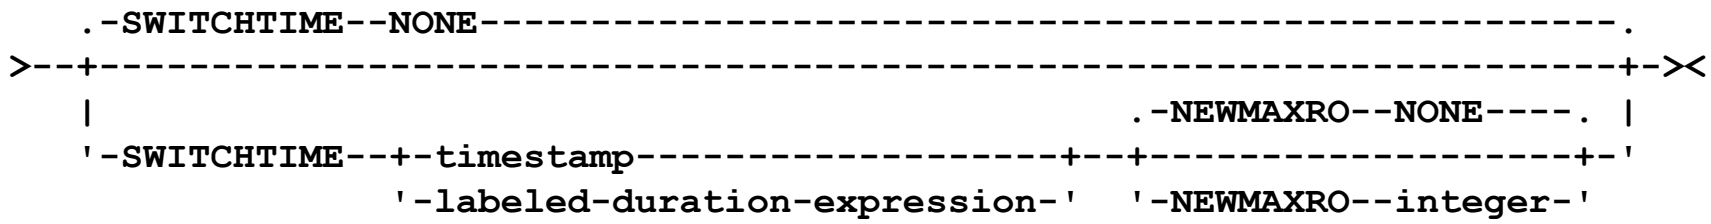

Physically delete empty PBG partitions

- Ability for REORG to physically delete empty PBG partitions
- New zparm REORG_DROP_PBG_PARTS
  - DISABLE – keep V10 behaviour (default)
  - ENABLE – Delete empty PBG partitions on table space-level REORG
- Considerations:
  - Cannot be specified on REORG statement
  - If PBG created using NUMPARTS or ALTER ADD partition used, REORG may prune to a lesser number of partitions
  - No PIT recovery to prior to a pruning REORG
    - No facility to resurrect deleted partitions

Automated mapping table handling

- Support mapping tables in PBG
  - Increases mapping index limit from 64Gb to 16Tb
  - Retrofitted to V9 and V10 via PM58177
- Mapping table DDL must change in 11 due to RBA/LRSN change
- Requirements to automate mapping tables
- So… New automated mapping tables in REORG
  - Automatically create new format mapping table if required
    - If mapping table specified & correct format then honour specification
    - Else if specified but incorrect format then create new in same db as original
    - Else if not specified and zparm DB specified then create in zparm DB
    - Else create in implicitly created DB
    - DROP at end of REORG or end of last REORG if multiple REORGs in job step
  - NFM requires new format mapping table, CM supports both old and new
  - No additional auth requirements necessary for creation of mapping tables

REORG without sorting data

- Increasingly REORGs are performed for reasons OTHER than to regain clustering of data, yet no ability to avoid cost of reclustering
- Prior to DB2 11, REORG SHRLEVEL CHANGE did not support SORTDATA NO
- DB2 11 supports SORTDATA NO with SHRLEVEL CHANGE
- New RECLUSTER YES/NO option on SORTDATA NO
  - RECLUSTER NO – Do not unload data through clustering index and do not sort data records in clustering order

Improved REORG LISTDEF processing

- PARALLEL YES/NO option introduced in APAR in V9
  - NO – Prevent REORG from processing multiple partitions in single REORG when input is partlevel LISTDEF
  - Zparm REORG_LIST_PROCESSING at zparm level
- Need compromise option for customers who want to take advantage of REORG parallelism but cannot afford to shadow many partitions at a time
- New option LISTPARTS n to limit # of partitions to be processed in a single REORG if input is a part-level LISTDEF
- In DB2 11, PARALLEL YES/NO is superseded by LISTPARTS
  - PARALLEL YES/NO IS deprecated but still supported in 11
- Retrofitted back to DB2 10 via PM52012

REBALANCE enhancements

- Improved availability & failure prevention
- Support REORG SHRLEVEL CHANGE REBALANCE
  - Complements online ALTER of partition limit keys
  - Retrofitted back to V10 (and V11 CM) via PI11839
- Improve resiliency with enhanced distribution algorithm & improved handling of empty partitions
- Build compression dictionary for all partitions
  - Previously, partitions that were empty at the start of REORG would not have a dictionary built, requiring a subsequent REORG to gain compression
- New SORTCLUSTER option to sort data in clustering as well as partitioning order to avoid AREO*
  - Occurred when partitioning key not a superset of clustering key
  - SORTCLUSTER YES – sort in partitioning and clustering order, avoid AREO*
  - No specification – keep existing behaviour, set AREO*
  - SORTCLUSTER NO – keep existing behaviour, but do not set AREO*

REORG of LOB data

- Support REORG of LOB data even though aux index is unavailable
  - Problem in V10 if LOB tablespace is REORP and index is RBDP
    - LOBs can’t be reorged and index can’t be rebuilt
- REORG SHRLEVEL NONE for LOBs changed to RC8 from 11 CM onwards
  - Not supported in 10 NFM, but returns RC0 with MSGDSNU126I
- REORG LOB Performance enhancement
  - PI17945 (V11 and V10)

Improved REORG serviceability

- Need ability to use online REORG even when SYSLGRNX cannot be relied upon
- Support LOGRANGES NO option for REORG SHRLEVEL CHANGE

REORG change of defaults to match best practices

- Change default options:
  - DRAIN WRITERS to DRAIN ALL
  - DISCARD to DISCARD NOPAD YES
  - UNLOAD EXTERNAL to UNLOAD EXTERNAL NOPAD YES

Statistics Collection Process

Runstats process at KBC

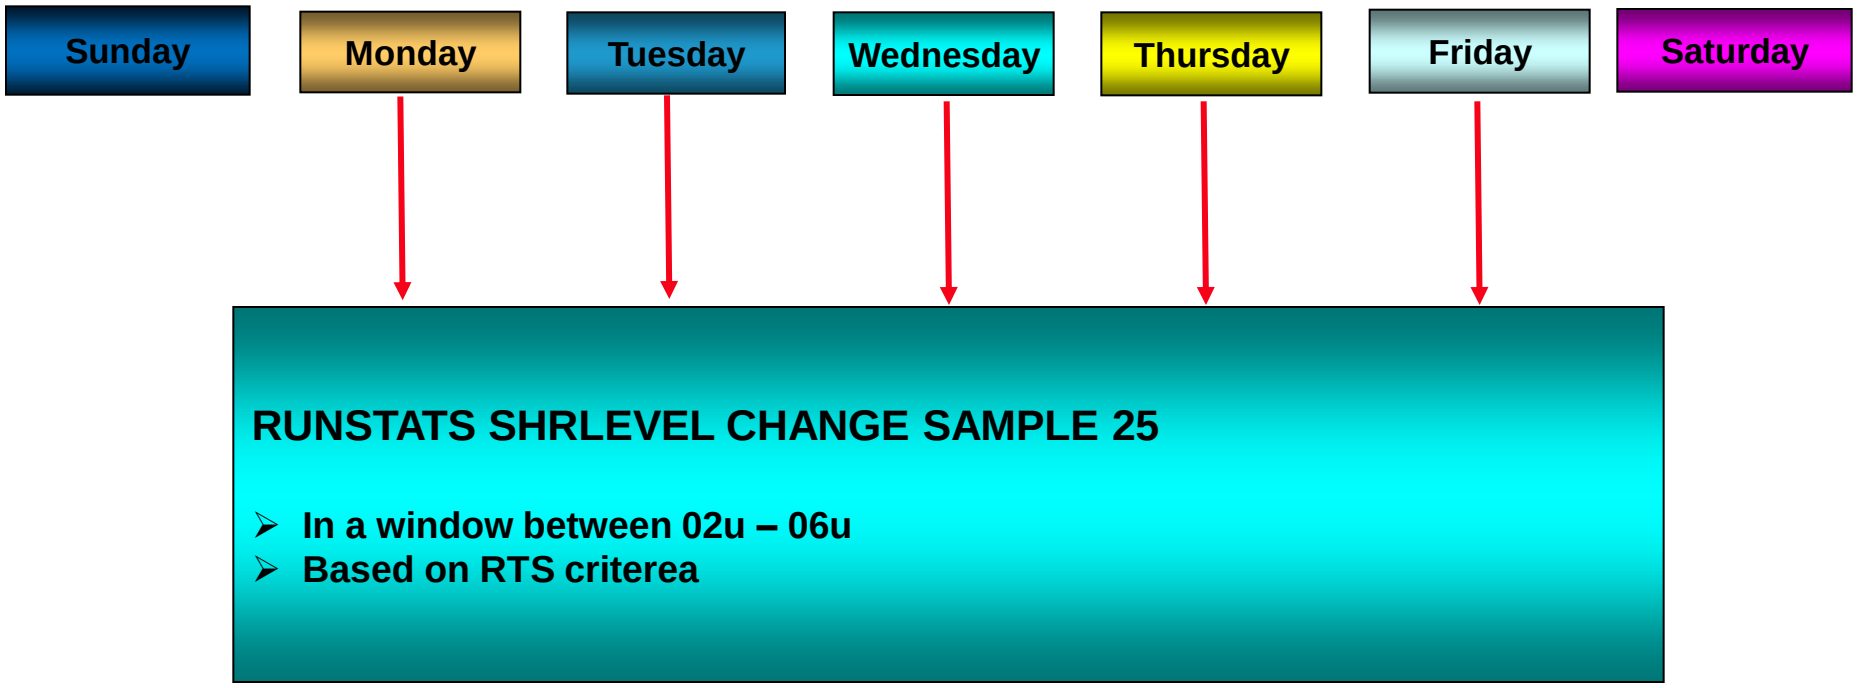

Runstats process at KBC

- RTS runstats criteria checked:
  - changed data > 20%
  - last runstats older then 2 months
- All TS an IX, except objects registered in STATS_EXCL table
- Inline stats (during reorg, load)
- Stats are manipulated for new empty tables (until data is available), because of some issues with -1 value.
- Utility jobs generated via KBC application and scheduled in batch window (run history is kept in a DB2 table)

The future at KBC

- Upgrade to DB2 V11:
  - Re-evaluate all RTS criteria…
    less need for reorg
  - … or make use of new DSNACCOX

KBC

BNPPBFortis Volvo
IS4F KBC

Statistics collection

Frequency	% DB2 data All = 100%	RTS usage	Statistics collection and •reorg •load , ..	DB2 Statistics manipu -lation	Reasons Policy
Daily, weekly, monthly, yearly					
Parameter driven	100%	Yes	RUNSTATS		
	2	Share	No	Online	DBA trigger
	All	Yes	Inline stats	No	daily
	95%	Yes	Yes, inline stats are taken	New empty tables (until data is available). Because of some issues with -1 value	Criteria: 1) Changed data > 20% 2) Last runstats older then 2 months

BNPPBFortis Volvo
IS4F KBC

Triggering / Scheduling Utilities

Trigger →	↓			Scheduled →	↓			
Fixed (Daily/Weekly/ monthly/ ...)			Daily	OEM tool specific for DB2			no	Every day/ based upon RTS criteria
Own RTS queries	YES	yes	-	Std job scheduler	YES		TWS	TWS controlled
RTS + DSNACCOX			RTS + DSNACCOX (PDA)	manual			On request	RTS are queried on a daily base. We use KBC utility criteria instead of DSNACCOX.
OEM DB2 Tool								No
Home grown applications	YES	YES		Java application upon Z/OS		40,000 by day,		YES

Statistics Collection – Discussion items

- Inline stats with LOAD/REORG vs RUNSTATS
- Only collection stats with REORG
- Use of STATS profiles ?
- Monitor stats changes between runstats collection (history stats) – do they vary enough to warrent their collection
- Not afraid that changed stats will cause access path regression at rebind (or prepare) ?
- RTS stats are used as input !
  - Via own pgm
  - Via DSNACCOX
- Std job scheduler used to run the jobs

Statistics Collection – V11 Enhancements

Statistics Enhancements

- More zIIP offload for RUNSTATS distribution statistics
  - Up to 80% zIIP-eligible
- zIIP offload for inline statistics
  - Additional 30% offload to zIIP
- Enhance inline statistics for RUNSTATS avoidance
  - Inline statistics collection on NPSIs during REORG with SORTNPSI
  - Inline histogram statistics
  - Inline DSTATS
- New RUNSTATS RESET option to reset existing statistics
- Improved PROFILE usability for LISTDEF processing
  - Gather default statistics if no profile exists for table

Statistics Enhancements

- Optimizer determination of missing statistics
  - Optimizer identifies missing statistics & writes information to new catalog table SYSIBM.SYSSTATFEEDBACK
  - EXPLAIN will populate userid.DSN_STAT_FEEDBACK table
  - OQWT modifies statistics profile
  - Automation Tool detects profile change & builds RUNSTATS job
- -ACCESS DATABASE … MODE(STATS) option to externalize RTS statistics
- RTS overhead reduction

BIND / REBIND Process

To Rebind or not to Rebind at KBC

- Binds are done after database/application changes
- Mass rebinds after each DB2 migration
  - Make use of new optimizer
    - New accesspaths are trusted (APCOMPARE warnings are logged)
    - plan_mgmt: basic (some APREUSE exceptions)
  - Removal of invalid SPROC's
  - But once a bind is done in a specific DB2 release our DBA's prefer stability of the accesspath.

KBC

BNPPBFortis Volvo
IS4F KBC

Bind / Rebind -1

Frequency	New application version	After reorg Statistics collection	After a new DB2 version	Policy - Bind - Rebind …
Explain	YES	YES	YES	
	Yes	No	Yes	rebind
	Yes	Yes	Yes	Rebind
	yes	no	yes	
Using PLANMGMT	YES	YES	YES	
	Yes (Extended)	Yes (Extended)	Yes (Extended)	Rebind
	Yes (basic)	No	Yes (basic)	
Using APREUSE	NO	NO	YES	
	No	No	No	
	no	No	No (Some exceptions)	After DB2 release migration we let the optimizer take new access paths.
Using APCOMPARE	NO	NO	NO	
	No	No	No	
	no	No	warning	DBA's follow up the warnings.
Access path comparison in same/another environment via OEM tool or home grown applications	Home grown	Home grown	Home grown	
	yes			
	No	No	No	
	In an acceptance environment.	No	In an acceptance environment.	

BNPPBFortis Volvo
IS4F KBC

Bind / Rebind -2

Frequency	New application version	After reorg Statistics collection	After a new DB2 version	Policy - Bind - Rebind …
Follow-up/decision	Auto	Auto	Auto	Manual validation for those that are considered worse
	yes			
	No	No	No	
	DBA's	No	DBA's + System Engineers	
Strengths / weaknesses		Ongoing discussion: to rebind or not		No large issues anymore after using this setup
		Once a bind is done in a specific DB2 release, we prefer stability of the access path.		
	yes			

BIND / REBIND - Discussion items

- Everybody bind/rebind in new version – very good !
- PLANMGMT usage
  - Not after REORG/RUNSTATS at KBC (but used in other cases) – why ?
- APREUSE not used a lot
  - Trust the optimizer and/or use plan mgmt to fall back ?
- What about dynamic SQL ?
  - No plan mgmt to 'freeze' acces path
  - Worries , concerns ?
- APCOMPARE - (almost) not used
- Rebind after REORG/RUNTATS
  - Why ? – not afraid of regression ? – mitigated by using plan mgt ?
  - High degree of confidence in the optimizer
    - Not that many problems after rebinding

BIND / REBIND - V11 Enhancements

- APREUSE(WARN)
  - New in DB2 11 (DB2 10 delivered APREUSE(ERROR))
  - Upon failure of reuse, Optimizer will generate a new access path choice
    - Thus failure of 1 SQL will not fail the entire package
  - PLAN_TABLE output will represent a valid plan
    - For both ERROR or WARN
- Statistics feedback
  - Give optimizer the best opportunity to find good access path by providing all 'useful' stats for a query
  - See statistics topic
- DB2 11 Selectivity Overrides (FF hints)
  - Process of supplying more robust selectivity (Filter Factor) input
  - Rather than a whole OPTHINT – just FF hints

Schema Changes

Schema Changes at KBC

- During application release changes
  - 8 times a year, max outage 4u
  - Pressure to do more often (Agile development)
    -> small changes = once a month
- Changes are made via alter (60%) and drop/create (40%)
- After change integrity is checked (data + index)
- After alter change reorg is run (we want al objects in RW status)

KBC

BNPPBFortis Volvo
IS4F KBC

Schema changes

Frequency	Max Outage Online Batch	Type % drop/create % alter	Reorg delay after alter	Consistency Integrity check	Reasons Policy	
Daily/weekly						
Monthly	During major releases: 4u	40% drop/create 60% alter	Yes, we want al TS in a RW status after schema change.	Check data en check index are run.	Pressure to do schema changes more often (agile development) .	
Quarterly		Alter when possible	Yes		P weekend	
On demand	100	50/50	No	Yes	performance	
	depends	No stats	Daily reorg (check on AREOR status)	If needed check utility is run		

Schema Changes – Discussion items

- People exploiting online schema changes ?
  - Manually or only because a ‘tool’ exploits them ?
  - Which is the most important one that is missing ?
- Object that are not in RW state
  - If not DB2 does a DBET check may result in LC12 contention if the object is heavily used

Schema Changes – V11 Enhancements

- Alter limit key has become a ‘pending’ change
  - See before
- DROP column online schema change support
- Some PIT recovery restrictions after online schema change have been removed
- DDL/BIND break-in
  - Including persistent (local) thread support

Ad Hoc Recoveries

Ad hoc DB2 recovery at KBC

- Each year only a few cases
  => risk of loosing the skills!
- Tactics to avoid errors during recoveries
  - Application error(most cases): avoid recovery
    - Standard recovery to “independent outspace” , errors corrected via BMP
    - Generate undo/redo SQL
  - Dry run recoveries before executing them
  - Several template JCL’s available
    Commands, syntax: if not in the template: check the manual!
  - 4 eyes principle!

KBC

BNPPBFortis Volvo
IS4F KBC

Ad hoc DB2 recovery

Frequency	Nbr of objects	FIC + log vs. 'backout' from log	Logical recovery (under a different name)	Max outage •Batch •Online	Reasons Policy	Practice/test
Several per day						
Daily						
Weekly						
Quarterly	80%	YES			Refresh	
On Demand			YES		Verification	
	depends	FIC + log	no	As long as it takes	Only TS recovery + rebuild indexes	For Financial critical systems yearly test
Yearly	< 5	Fic + log	No		Logical problem	(daily… all zero)
Almost none	Very few	Depending on the case	The standard recovery procedure is to recover the requested data into a shadow TS.	As long as it takes.	We only recover TS. Indexes are always rebuild.	Very little.

AD Hoc Recovery - Discussion items

- Very few ad hoc recoveries
  - That is good news I guess
  - Better developers/operations staff ?
  - May result in knowledge of how to do this getting lost (Practice ?)
- Are there SLAs for those recoveries ? Or always 'best effort'
- What kind of recoveries
  - To quiesce point
  - One objects / application / entire system ?
  - DB2 + other data (IMS, CICS, MQ, flat files....) ?
- When using 'shadow' to recover into
  - Why ?
  - Using Tooling to assist with that ?

Ad Hoc Recovery – V11 Enhancements

Backup & Recovery Enhancements

- Faster catalog/directory recovery
  - Enhanced SYSLGRNX recording
- New VCAT name translation for RESTORE SYSTEM for system cloning
  - Support logapply when RESTORE SYSTEM used for cloning purposes
- Improved recoverability with COPY-REORG concurrency
  - Permit COPY to run concurrent with long-running REORGs
- Avoid allocating empty image copy datasets for incremental or CHANGELIMIT copies

Backup & Recovery Enhancements

- Lifted many restrictions on point-in-time recovery prior to materializing REORG
  - PIT recovery restrictions lifted for
    - LOB table spaces
    - XML table spaces
    - PBR table spaces
    - Including when immediate alters have occurred since materializing REORG
  - PIT recovery restrictions still in place
    - Table space conversion
    - PBG table spaces
    - PBG partition pruning
    - Online DROP COLUMN
- Faster index recovery with FLA support for index log records
  - APAR PI07694 for V9 and above

Other DB2 11 Utilities Enhancements

Compression Dictionaries

- Avoid decompression failures for IFI 306 readers when new compression dictionary built by REORG/LOAD
- Old compression dictionary stored on log
- New SYSCOPY record written pointing to old compression dictionary for CDC tables
- IFI 306 read automatically retrieves old compression dictionary if necessary
- Avoid need for replication target refresh when dictionary changes

General Enhancements

- Lots of changes to all utilities to support extended RBA/LSRN format
  - Largest effort of all changes in DB2 11 to support this
- Greater parallelism for faster utilities
  - 11% elapsed time reduction measured for REORG, LOAD, REBUILD INDEX
- PARALLEL *n* option for parallelism control for LOAD, REORG, REBUILD INDEX, UNLOAD, CHECK INDEX
- -DISPLAY UTILITY enhancements
  - Remove serialization between –DIS UTIL and –TERM UTIL
  - Jobname, start timestamp
  - Late addition: SWITCHTIME and NEWMAXRO
- Utility impact reduction on bufferpools
  - Extend MRU for UNLOAD, REORG TABLESPACE, RUNSTATS TABLESPACE, RUNSTATS INDEX, REBUILD INDEX, CHECK INDEX, CHECK DATA
- Improved dataset cleanup in utility stored procedures
  - Previously, datasets remained allocated on utility failure, preventing cleanup

General Enhancements

- Improved TEMPLATE support for large / EF datasets and local time values
  - DSNTYPE LARGE, EXTREQ, EXTPREF
  - New EATTR option on TEMPLATE to request extended attributes
  - New TIME LOCAL|UTC option
- Enforce NUMTCB=1 for stored procedures
- DSNACCOX performance

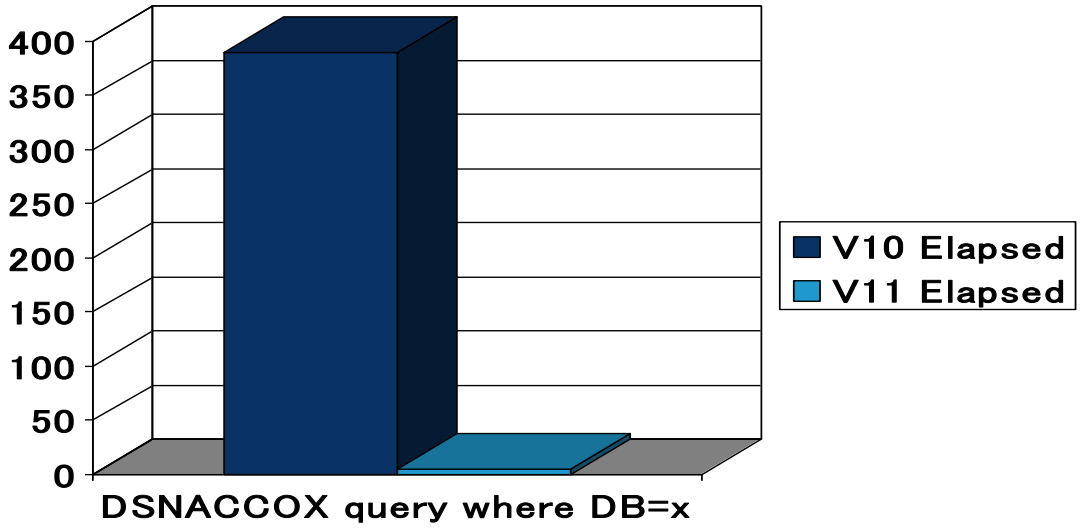

Deprecation

- REORG SHRLEVEL NONE for LOBs changed to RC8 from 11 CM onwards
  - Not supported in 10 NFM, but returns RC0 with MSGDSNU126I
- Still supported in 11, but no longer documented:
  - REORG
    - PARALLEL YES|NO

      Superseded by LISTPARTS
    - INDREFLIMIT
    - OFFPOSLIMIT
    - LEAFDISTLIMIT
    - UNLOAD ONLY
    - UNLOAD PAUSE
    - UNLOAD EXTERNAL
  - COPY
    - CHANGELIMIT

Recent News

- Part-level REORG NPSI insert performance improvement
  - PM87403 (V9)
  - 100m row table, 6 indexes
  - LOAD RESUME – 66% CPU reduction, 30% ET reduction
  - REORG PART 9 – 45% CPU reduction, 26% ET reduction
- Fast log apply for faster index recovery
  - PI07694 (V9)
- Retrofit REORG SWITCH phase performance to V10
  - PI09303 (V10)
  - Cut SWITCH phase duration by 90%
- Retrofit REORG REBALANCE SHRLEVEL CHANGE to V10
  - PI11839 (V10)
- LOAD REPLACE SHRLEVEL REFERENCE, LOAD RESUME SHRLEVEL REFERENCE, LOAD prevalidation
  - Delivered in Utilities Enhancement Tool / Utilities Solution Pack
  - PI04864
- LOB REORG performance improvement
  - PI17945 (V10)

Recent News

- Easy-tier storage support in REORG
  - PI35321
  - REORG shadows inherit temperature of original pagesets
  - No externals
  - OA46482 required for DFSMS support

LOAD & UNLOAD Enhancements

- Crossloader support for XML data
- Exploit FETCH CONTINUE for processing large LOBs & XML data in Crossloader
  - Reduce vstor requirement
  - Avoid DSNU1178i errors
  - 28% CPU reduction
    - Load of 1Mb LOBs
- zIIP offload for LOAD REPLACE PART clearing of NPSIs
  - 100% offload to zIIP for LOAD REPLACE with dummy input

LOAD & UNLOAD Enhancements

- LOAD SHRLEVEL NONE PARALLEL with single input dataset
  - Parallel data conversion
  - Not supported for PBGs
  - 50% ET reduction possible on single SYSREC load

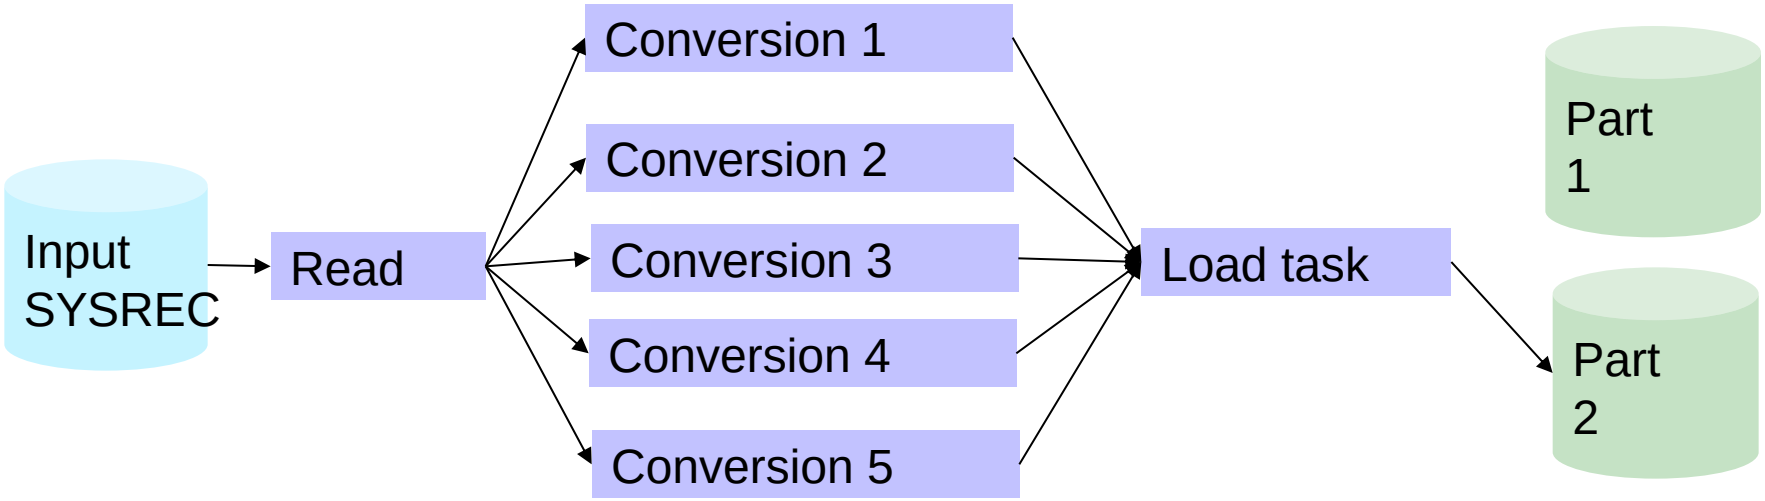

LOAD & UNLOAD Enhancements

- LOAD SHRLEVEL CHANGE PARALLEL
  - Supports non-partitioned as well as partitioned
  - Single input dataset
  - Not supported for PBGs
  - >80% ET reduction