Db2 13 for z/OS Technical Overview

TRIDEX

September 8, 2022

Robert Catterall, IBM
Principal Db2 for z/OS Technical Specialist

Agenda

Background

Installation and migration

Availability and scalability

Performance

SQL and application management

Security

Insight and oversight

Utilities

Summary

Db2 13 – Setting the Context for New Innovation

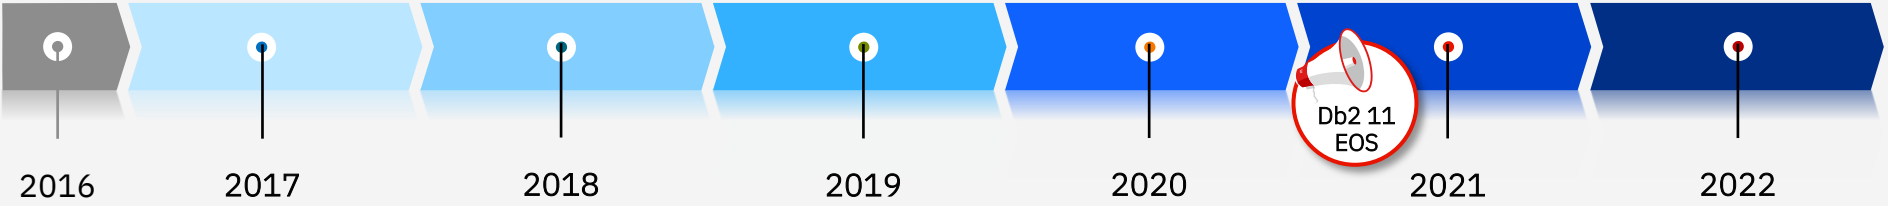

2016

Db2 12 for z/OS GA
Oct 2016

+ Over 100 new Db2 12 features delivered since GA that are not tied to specific Function Level

2017

FL 501 – 1st post-GA delivery
- LISTAGG

2018

FL 502 – Apr
APAR PI95511
- Transparent Dataset Encryption: Db2 DBA controls
- Casting numeric to GRAPHIC / VARGRAPHIC

FL 503 – Sep
APAR PH00506
- Db2 AI for z/OS (Db2ZAI) – prereqs FL500
- Migration support on DATA CHANGE OPERATION for temporal auditing
- Enable replication of system-period temporal tables and generated expression columns

2019

FL 504 – Mar
APAR PH07672
- Huffman data compression
- New SQL syntax alternatives
- Prevent new deprecated objects
- Passthru of Built-In Functions (OLAP to Db2 Analytics Accelerator)

FL 505 – Jun
APAR PH09191
- Rebind Phase-in
- New RUNSTATS sampling
- BIF for Transparent Column Encryption (TCE)
- Temporal and archive on WHEN clause of triggers
- Indexes for Decfloat

FL 506 – Nov
APAR PH16829
- DROP TABLE automatic drop of explicit table spaces
- Alternative spellings for existing SQL built-in scalar functions to improve compatibilit

2020

FL 507 – Jun
APAR PH24371
- Application-level locking control
- Improved statistics management
- CREATE OR REPLACE
- Db2 Analytics Accelerator passthrough expressions

FL 508 – Oct
APAR PH29392
- Multi-table table space conversion to Universal Table Space (UTS)

2021

FL 509 – Feb
APAR PH33015
- Support for tamper-proof audit policies
- HA for AOTs
- Specify a compression algorithm at the table, table space, or partition level
- Temporal RI allows UPDATE or DELETE on the parent table

FL 510 – Apr
APAR PH33727
- Automation and groundwork to set up ease of migration from Db2 12 to Db2 VNext

2022

Db2 13 - May
- Industry-first AI-infused insights into business data
- Enhanced operational flexibility with no application outages
- Improved interoperability with cloud applications
- Streamlined version upgrade with no application outages

Function Level cadence resumes with 1 in 2H22 and 3 planned for 2023

Announced April 5, 2022

Available May 31, 2022

IBM z16™

+

IBM Db2 13 for z/OS

Better together!

Simpler installation and migration

- Fewer ZPARMs and improved defaults or allowable ranges
  - 16 removed, 12 updated
- Migration only from V12R1M510 with validation of migration readiness
- V13R1M100 introduces no new external function
  - No catalog structural changes in V13 until after V13R1M500 to minimize risk and complexity
  - But catalog must be V13R1M100 “level,” so trivial CATMAINT required at M100
- V13R1M500 supports new features without catalog changes
- V13R1M501 supports all the initial new features of Db2 13

Availability and scalability (1 of 8)

- ALTER TABLE DATA CAPTURE CHANGES will no longer serialize with application threads
  - Before: had to quiesce static packages, quiesce and invalidate dynamic statements
  - Now: no quiesce of static packages or quiesce or invalidation of cached dynamic statements
  - This is an initial step in improving concurrency from an ALTER perspective

ALTER	base run 1	base run 2	base run 3	new run 1	new run 2	new run 3
Executed	10	10	10	10	10	10
Failed	9	10	9	0	0	0

Availability and scalability (2 of 8)

- Online conversion from PBG to PBR
  - New ALTER TABLE syntax for pending alter to convert to PBR RPN
  - All related artifacts preserved (e.g., indexes) – table ready for use after online REORG
  - Currently not supported for tables with LOB or XML columns

```
ALTER TABLE SCR001.TB01 ALTER PARTITIONING TO
                          PARTITION BY RANGE (ACCT_NUM)
                        ( PARTITION 1 ENDING AT (199),
                          PARTITION 2 ENDING AT (299),
                          PARTITION 3 ENDING AT (399),
                          PARTITION 4 ENDING AT (MAXVALUE));
```

Availability and scalability (3 of 8)

- Online removal of active log datasets allows complete online replacement of logs
  - New –SET LOG REMOVELOG option
    - Remove from BSDS or mark removal pending for later removal
  - Enhanced –DISPLAY LOG command (new DETAIL option)
- Quadruple size of critical SPT01 and SYSLGRNX table spaces to 256 GB
  - Catalog/directory PBGs are defined with MAXPARTITIONS 1 DSSIZE 64G
  - SPT01 defaults to inline LOB length of the maximum (32138)
  - V13: SYSLGRNX, SPT01 changed to DSSIZE 256G on REORG after M500
    - Optional and not part of any CATMAINT

Availability and scalability (4 of 8)

- Support more open datasets
  - Actual open dataset limit is a function of the amount of memory consumed per open dataset below 2GB bar and amount of 31-bit private storage
  - Db2 APARs PH09189 (threshold processing) and PH27493 (utility data sets) improved open dataset management
  - z/OS 2.5 reduces below-the-bar memory consumption per dataset by ~35%
  - Optionally move SWB blocks above 2 GB bar
    - Update ALLOCxx member of PARMLIB, or…
    - Issue command SETALLOC SYSTEM,SWBSTORAGE=ATB
  - Db2 13 exploits new z/OS API
  - Double DSMAX limit from 200,000 to 400,000

Availability and scalability (5 of 8)

- Virtual storage constraint relief
  - Reduce ECSA consumption by Db2 IFCID tracing
    - Cut by up to 80% (from around 30 MB to 4-8 MB)
  - Reduced ECSA footprint for distributed threads
  - Move dynamic SQL statement text above 2 GB bar
  - BIND/REBIND/PREPARE storage reduction
- Improved storage management
  - Improved storage contraction for thread and ECSA pool storage
  - Prevent flood of DISCARDDATA requests after DDF workload spikes (improved “right-sizing” of DBAT pool following workload spike)
  - Remove REALSTORAGE_MANAGEMENT ZPARM

Availability and scalability (6 of 8)

- RTS improvements to support larger workloads
  - Change various “counter” columns from INTEGER to BIGINT (up to about 9 quadrillion)
  - Change to LOCKMAX 0 to avoid lock escalation
- PAGESET_PAGENUM ZPARM default changes to RELATIVE (relative page numbering for PBR table spaces)
  - Note: PBR RPN requires EA/EF datasets
- Simplify IRLM memory management
  - Value of ZPARM parameter MAX STORAGE FOR LOCKS now passed by Db2 when IRLM is started
  - Can no longer be overridden in SMFLIMxx PARMLIB member
    - Effect: cut down on “surprise” changes in IRLM storage
  - Can still be changed via MODIFY IRLM command

Availability and scalability (7 of 8)

- Dynamic alteration of lock structure to avoid structure full conditions
  - More responsive than ALLOWAUTOALT in CFRM policy

```
DXR189I  <irlmname> ALTERING LOCK STRUCTURE SIZE
DXR190I  <irlmname> ALTER LOCK STUCTURE COMPLETED
```

```
IXC530I START ALTER REQUEST FOR STRUCTURE DSNCAT_LOCK1 ACCEPTED
  TARGET SIZE:                    8 M
IXC534I REQUEST TO ALTER STRUCTURE DSNCAT_LOCK1
COMPLETED.  TARGET ATTAINED.
CURRENT SIZE:                  8 M  TARGET:          8 M
CURRENT ENTRY COUNT:        4427    TARGET:       4427
CURRENT ELEMENT COUNT:         0    TARGET:          0
CURRENT EMC COUNT:             0    TARGET:          0
```

Availability and scalability (8 of 8)

- Positioning for statement-level dependency in future (think about ALTERs that invalidate entire packages)
  - Disabled in V13, but some catalog objects are created
    - New table SYSPACKSTMTCOPY
    - New table SYSPACKSTMTDEP
    - New column SYSPACKSTMT.VALID
    - Indexes on new tables
    - All new page sets are DEFINE NO

Performance (1 of 4)

- Improved PBG insert processing – reduce incidence of invalid “no space for INSERT” situations
  - Retry conditional partition locks
  - Smarter space search across partitions
- More efficient handling of DBAT threads
  - Improved termination of pooled and keep dynamic DBATs caused by workload spikes
- Intelligent use of internal block fetch
  - Track statement fetch history for better internal block fetch decisions

Performance (2 of 4)

- Intelligent index lookaside support
  - Adjusts usage based on random vs. sequential index access (usage based on learning, versus just catalog stats)
    - Index lookaside benefits sequential access (FTBs primarily benefit random access)
- Greater index FTB eligibility
  - Unique indexes: Increase max key length from 64 to 128 bytes
  - Non-unique indexes: Increase max key length from 56 bytes to 120 bytes
- Drop DSNKDX02 on SYSIBM.SYSPACKDEP
  - Redundant since creation of DSNKDX03

Performance (3 of 4)

- More efficient IFCID306 log read - better support for partition-level reads
  - New partition range log read support in IFCID 306 interface
  - One result: ability to drive parallel replication streams for one partitioned table space
- Reduced false lock contention for PBR RPN in data sharing
  - Reduced page P-lock contention for row-level locking
- Authorization cache improvements to reduce RACF contention
  - Cache successful external plan auth checks
  - Smarter caching to cache more auth IDs per plan

Performance (4 of 4)

- More efficient data sharing castout processing
  - Lower GBPOOLT checking interval from 10s to 1s
  - Reduce write failure retry interval for GBP full conditions
- SQL sort performance improvements
  - Use query history (learning) to expand SORTL use
    - Db2 support for z15 SORTL instruction originally in PH31684
- System recovery performance boost
  - z16 automatic support for boost during restart of Db2 and IRLM (V12 and V13)

SQL and application management (1 of 9)

- SQL Data Insights
  - Infuse AI to deliver SQL semantic query support
  - New built-in functions
    - AI_ANALOGY
    - AI_SIMILARITY
    - AI_SEMANTIC_CLUSTER
  - Exploits zIIP engines

Semantic SQL functions

First set of AI built-in functions available in Db2 13

Cognitive intelligence query	Functional description	Db2 functions
Semantic similarity and dissimilarities	• Matching rows/entities based on overall meaning (similarity/dissimilarity) • Suggest choices for incorrect or missing entities	AI_SIMILARITY
Semantic clustering	• **Find** entities/rows based on relationships between attributes in a given set • Example: Find animals similar to (lion, tiger, panther)	AI_SEMANTIC_CLUSTER
Reasoning analogy	• **Find** entities/rows based on relationships between attributes • Example: Moon : Satellite :: Earth; ?	AI_ANALOGY

SQL Data Insights – high-level overview

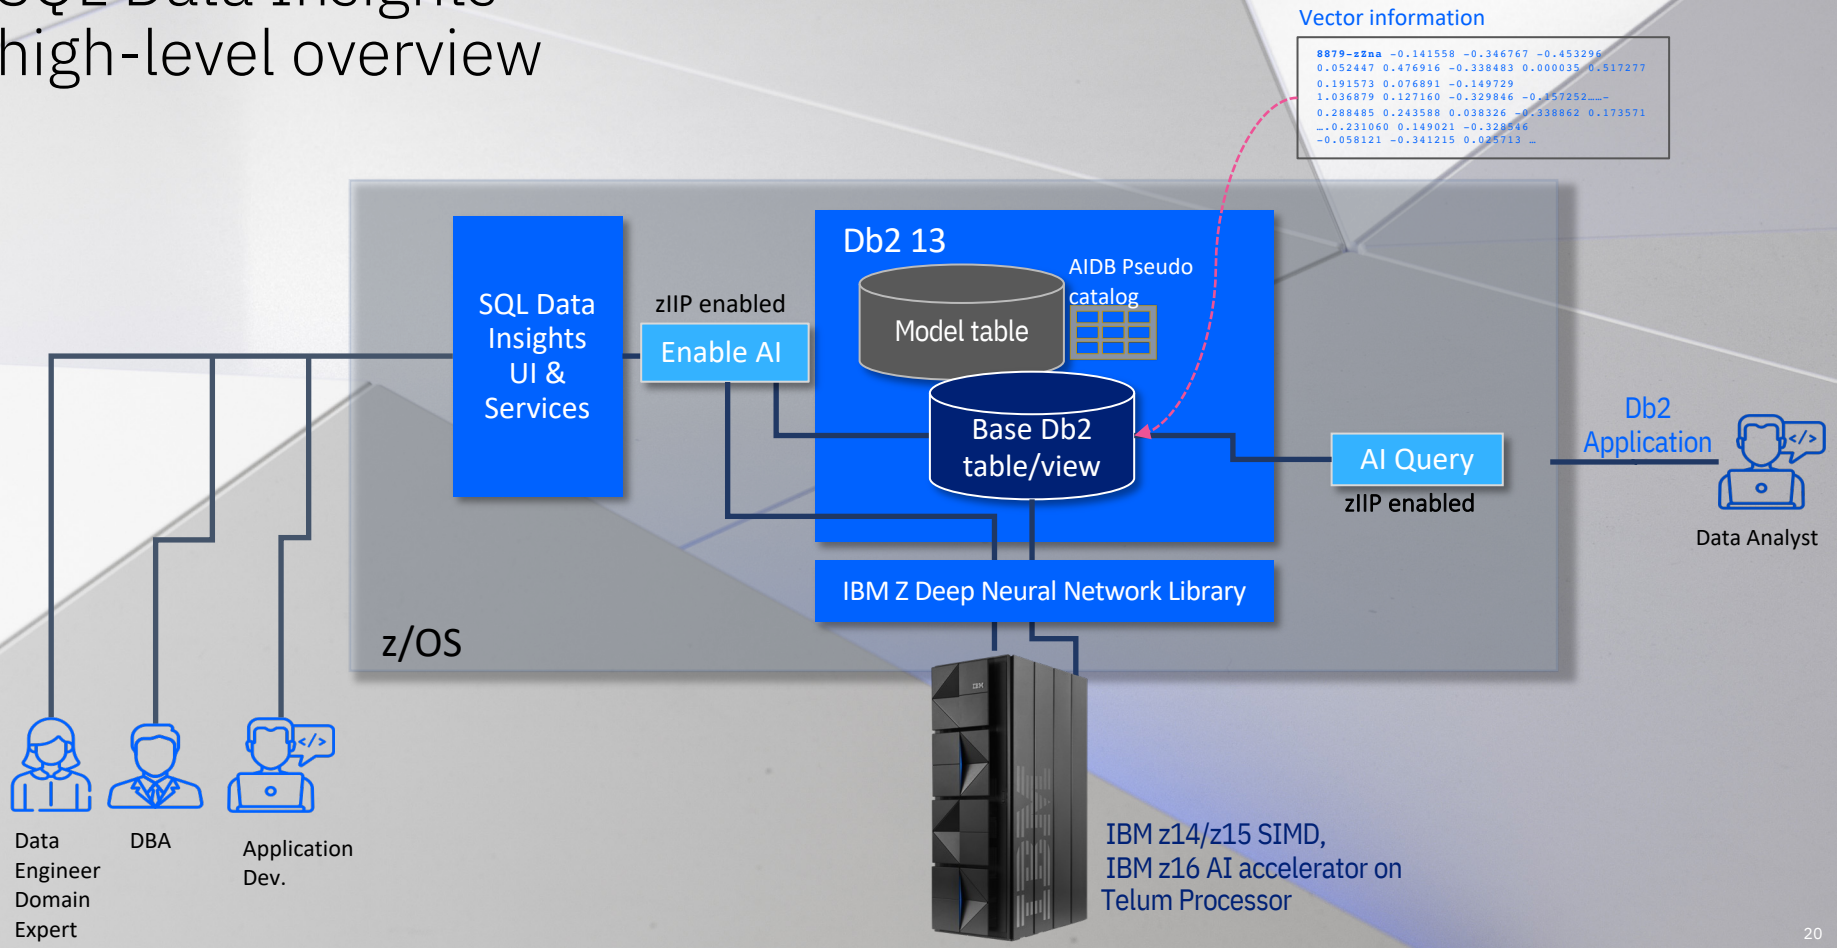

SQL and Application Management (4 of 9)

- New CURRENT LOCK TIMEOUT special register
  - Override IRLMRWT for claims and transaction L-locks
    - IRLMRWT also made online changeable!
  - NULL is default, equates to IRLMRWT
  - Range: -1 - 32767
  - ZPARM SPREG_LOCK_TIMEOUT_MAX provides system-level control
  - DSNT376I enhanced to indicate timeout value and source of timeout value
    - When special register set to 0 (application does not wait for lock), no lock holder information returned in message
  - Instrumentation enhanced
    - Accounting, stats trace information, IFCID 106 ZPARM information, IFCID 196 timeout information, new IFCID 437

SQL and application management (5 of 9)

- Note on setting current lock timeout at application level through new special register:
  - Data sharing processing overhead may make it difficult to achieve low lock timeout intervals (e.g. timeout 1 sec with deadlock cycle 1 sec)
    - Documented here:
      - https://www.ibm.com/docs/en/db2-for-zos/12?topic=processing-elapsed-time-until-timeout-data-sharing
    - Consider lowering deadlock detection cycle time to minimize delay
      - IRLM PH43770 and Db2 PH45103 (V12) reduce timeout delays

```
MIN GLOBAL TIMEOUT = timeout period + DEADLOCK TIME value
MAX GLOBAL TIMEOUT = timeout period + 4 * DEADLOCK TIME value
AVERAGE GLOBAL TIMEOUT = timeout period + 2 * DEADLOCK TIME value
```

SQL and application management (6 of 9)

- SYSIBMADM.DEADLOCK_RESOLUTION_PRIORITY built-in global variable
  - Influence deadlock victim decision
  - Range 0 – 255
  - IFCID 172 deadlock information enhanced
    - QW0172WAS
      - G = Global variable
      - O = Other
    - QW0172WA
      - Assigned worth value

SQL and application management (7 of 9)

- Profile table support for both SYSIBMADM.DEADLOCK_RESOLUTION_PRIORITY global variable and CURRENT LOCK TIMEOUT special register
- Profile table support for RELEASE_PACKAGE
  - Control package RELEASE option for both distributed and local threads
- Profile table support for local-to-Db2 applications (for application-level lock timeout and deadlock priority, and for package RELEASE setting)
- IFI306 support for decoding EDITPROCs
  - New WQALLOPT flag option
    - Provides standard mechanism for using EDITPROC-decode support for vendors of Db2 for z/OS data replication tools (vendors formerly required to work with Db2 for z/OS development to enable support)

IBM Data and AI on IBMZ / © 2022 IBM Corporation

SQL and application management (8 of 9)

- Improved consistency for creation of PBG table spaces
  - Background: for create of a PBG with explicit MAXPARTITIONS value, DSSIZE defaults based upon the following table

Page size	MAXPARTITIONS	DSSIZE default
Any	1-254	4G
4K	255-4096	4G
8K	255-4096	8G
16K	255-4096	16G
32K	255-4096	32G

IBM Data and AI on IBMZ / © 2022 IBM Corporation

SQL and application management (9 of 9)

- Improved consistency for creation of PBG table spaces (continued)
  - From V12R1M504, CREATE TABLESPACE will default to PBG with MAXPARTITIONS 256, instead of traditional segmented
  - This should have followed table (on preceding slide) *for default DSSIZE based upon page size*, but instead DSSIZE 4G is always given – we violated our own rule
  - With APPLCOMPAT set to V13R1M500 or higher, CREATE TABLESPACE without explicit MAXPARTITIONS will default to MAXPARTITIONS 254 – result: we're within the rules
    - Change applies to implicit PBGs also

IBM Data and AI on IBMZ / © 2022 IBM Corporation

Security

- Improved management of package ownership
  - Allow a DBA with role-based security to create packages using authid security
  - New SQL and BIND syntax to indicate whether owner is a role or a user
    - For CREATE PROCEDURE: PACKAGE OWNER *auth-name* AS ROLE | USER
    - For BIND/REBIND PACKAGE: OWNERTYPE(ROLE | USER)
- Db2 support for IBM compliance controls
  - https://www.ibm.com/products/z-security-and-compliance-center

Insight and oversight (1 of 2)

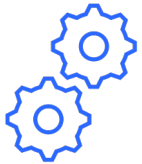

- New tracking of page GBP residency time for better GBP tuning tracking
- Track index page splits
  - New IFCID 396 (stats class 3 – active by default) when split takes >1 second
  - New SYSINDEXSPACESTATS columns
    - REORGTOTALSPLITS, REORGSPLITTIME, REORGEXCSPLITS

Insight and oversight (2 of 2)

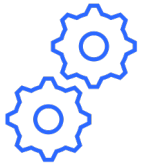

- New tracking of longest lock or latch wait by thread in accounting IFCID 3
  - Also longest wait for sync or async I/O, drain lock, service task
- STATIME_MAIN default changed from 60s to 10s to better diagnose workload peaks
- Include IFCID 369 (aggregate accounting information by connection type) in STATS CLASS 1 and 2
- New utility history table in Db2 catalog (more info to come)
- Support names > 16 bytes in deadlock and timeout messages DSNT375I, DSNT376I
  - And in IFCIDs 172 and 196

Utilities (1 of 4)

- Significant utilities investment supports core Db2 13 features such as online conversion from PBG to PBR
- Many major utility enhancements delivered via APAR in the years following Db2 12 general availability, including:
  - Redirected recovery
  - LOAD PRESORT
  - REORG inline copy improvements
  - LOAD REPLACE with SHRLEVEL REFERENCE
  - ...
- Following slides describes Db2 13-specific utility enhancements

Utilities (2 of 4)

- Central oversight over utility executions with new utility history catalog table
  - New ZPARM parameter: UTILITY_HISTORY NONE | UTILITY
  - Collect and store utility-level information in new SYSIBM.SYSUTILITIES catalog table
  - Who, what, when, where, etc.
  - New EVENTID column added to SYSCOPY to allow correlation
  - Object- and phase-level detail to follow later

Utilities (3 of 4)

- REORG INDEX performance
  - REORG INDEX NOSYSUT1 delivered for Db2 12 via PH25217
    - Keys not unloaded to data set – passed in memory to index build
    - Up to 80% ET and 90% CPU savings for REORG INDEX
    - Enabled via NOSYSUT1 parm of REORG, or REORG_INDEX_NOSYSUT1 parameter in ZPARM (Db2 12 default for that ZPARM is NO)
  - REORG_INDEX_NOSYSUT1 parameter in ZPARM will no longer have any effect as of V13R1M500
    - REORG INDEX will always avoid use of SYSUT1
    - REORG option NOSYSUT1 is ignored from M500 onwards
- Easier space-level RECOVER from part-level image copies
  - Supported for universal table spaces, associated partitioned indexes, associated LOB table spaces
  - Previously required use of LISTDEF PARTLEVEL keyword

Utilities (4 of 4)

- Page sampling support for inline stats
  - Support TABLESAMPLE for page sampling in inline stats
  - Also support STATPGSAMP in ZPARM for inline stats
    - Default of SYSTEM will mean that page sampling will be the default
- New REPAIR WRITELOG option to allow for writing of decompression dictionaries to Db2 log when dictionary created by other vendor products
  - Required for decompression of log records by IFI306 processing when new compression dictionaries are built to replace old dictionaries
  - Available in V13R1M100
  - Vendors responsible for writing correct information and inserting correct SYSCOPY entry

Thanks for your time

Robert Catterall
rfcatter@us.ibm.com