



# CA Database Management Solutions for DB2 for z/OS Utilities Overview

Manuel Gómez Burriel – May, 31<sup>st</sup>, 2018 – 2.13

## Prague Technology Days

May 30 - June 1, 2018



# For Informational Purposes Only

This presentation was based on current information and resource allocations as of **May 2018** and is subject to change or withdrawal by CA at any time without notice. Notwithstanding anything in this presentation to the contrary, this presentation shall not serve to (i) affect the rights and/or obligations of CA or its licensees under any existing or future written license agreement or services agreement relating to any CA software product; or (ii) amend any product documentation or specifications for any CA software product. The development, release and timing of any features or functionality described in this presentation remain at CA's sole discretion. Notwithstanding anything in this presentation to the contrary, upon the general availability of any future CA product release referenced in this presentation, CA will make such release available (i) for sale to new licensees of such product; and (ii) to existing licensees of such product on a when and if-available basis as part of CA maintenance and support, and in the form of a regularly scheduled major product release. Such releases may be made available to current licensees of such product who are current subscribers to CA maintenance and support on a when and if-available basis. In the event of a conflict between the terms of this paragraph and any other information contained in this presentation, the terms of this paragraph shall govern.

Certain information in this presentation may outline CA's general product direction. All information in this presentation is for your informational purposes only and may not be incorporated into any contract. CA assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, CA provides this presentation "as is" without warranty of any kind, including without limitation, any implied warranties or merchantability, fitness for a particular purpose, or non-infringement. In no event will CA be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, lost investment, business interruption, goodwill, or lost data, even if CA is expressly advised in advance of the possibility of such damages. CA confidential and proprietary. No unauthorized copying or distribution permitted.

# Agenda

- Introduction
- Utility Review
  - CA Fast Load
  - CA Fast Unload
  - CA Database Analyzer
  - Rapid Reorg

# CA Database Management for DB2 is a set of products broken down into 5 suites.

Utilities	Administration	Recovery	Performance	PRF
Fast Load*	RC/Update	Fast Check	Bind Analyzer	Report Facility
Fast Unload*	RC/Compare	Fast Recover	Detector	Compile/PRF
Database Analyzer*	RC/Migrator	Log Analyzer	Insight	
Fast Index	RC/Query	Merge/Modify	Plan Analyzer	
Rapid Reorg*	RC/Secure	Quick Copy	SQL Ease	
	Partition Expert	Recovery Analyzer	Subsystem Analyzer	

# DB2 Products: Utilities

Category	Products	Purpose\Value
Utilities	CA Fast Index® for DB2 for z/OS	Can Build Indexes and Add to DB2 Tables
	CA Fast Load for DB2 for z/OS	Can quickly load data to minimize downtime
	CA Fast Unload® for DB2 for z/OS	Provides Multiple options to unload Data
	CA Database Analyzer for DB2 for z/OS	Automatically generates IBM or CA utility JCL
	CA Rapid Reorg™ for DB2 for z/OS	Offers not only speed but also effective control and management of the reorganization requirements in a DB2

# CA Fast Load

## Function

- Fast Load for DB2 for z/OS is a comprehensive high speed utility for loading data into DB2 tables.
- It offers superior performance and important functional advantages in the areas of flexible formatting, data conversions.
- CA Fast Load can automatically resize and reallocate underlying DB2 datasets.
- Sort avoidance techniques can significantly reduce Load time.
- Collects DB2 statistics during the load process and can concurrently update the DB2 catalog.

## Value

- Cost reduction: Significantly lower CPU and Elapsed Times.
- Business impact reduction: Reduced impact on DB2 resources.

# CA Fast Load

## Benefits

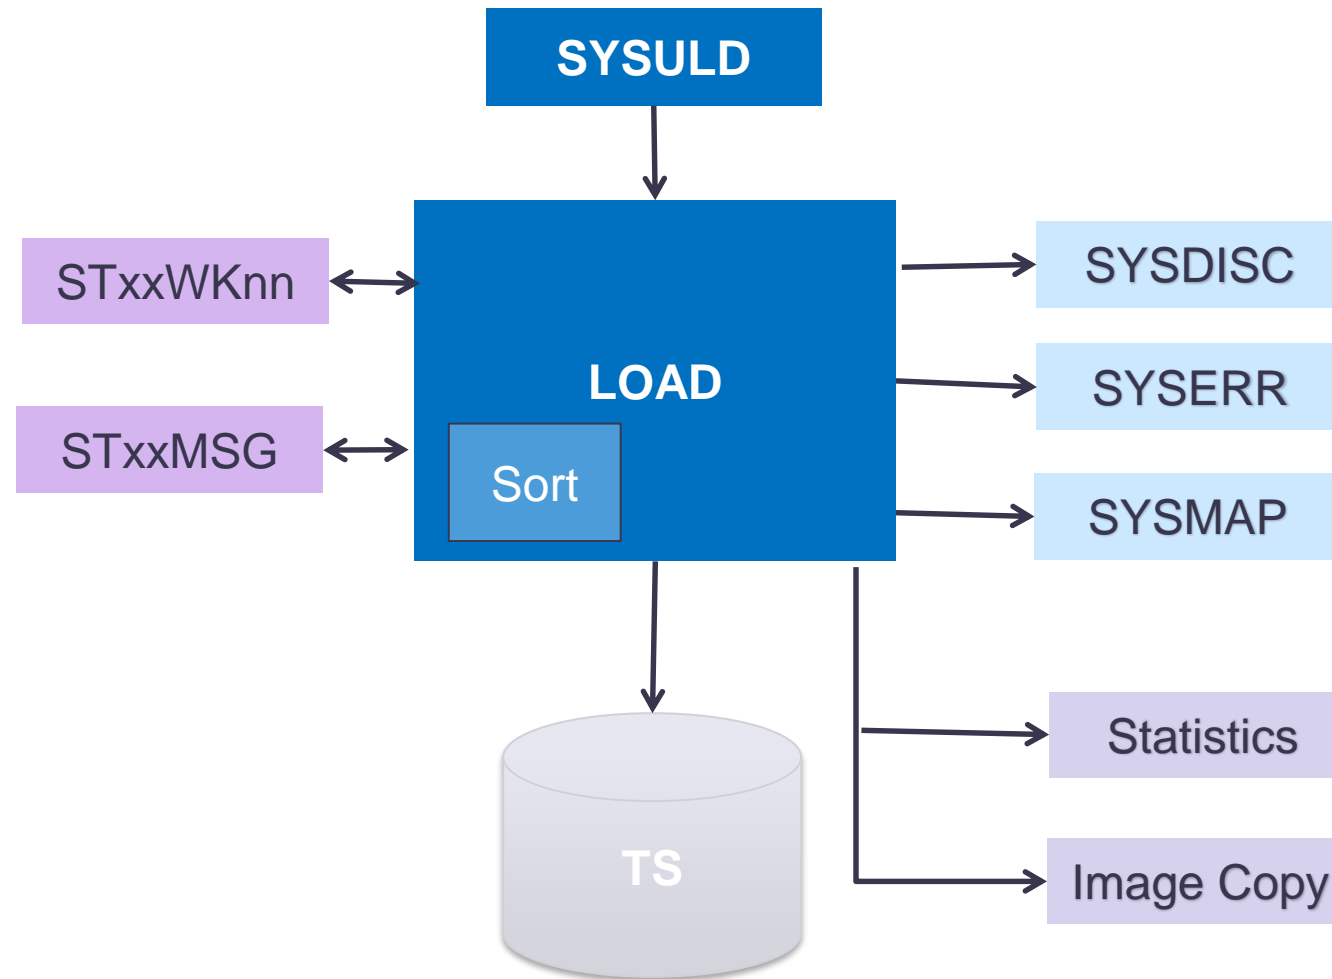
- Multiple Formats
  - Output-control ALL
  - Output-control Build
- Multi-task sort and build index when more than one index
- Optionally update PDA statistics and DB2 catalog
- Multi-task input datasets for partitioned tablespace
  - Partition Independence
- Multiple SYSREC datasets for output available, multi-task load phase
- Boolean logic supported for the DEFAULTIF and NULLIF parameters
- Dynamically allocate Sortwork datasets
- OBID conversion

# CA Fast Load: OUTPUT-CONTROL **BUILD**

- Fastest option for Fast Load
- Builds the tablespace directly, eliminating the need for large amounts of secondary storage for SYSUT1 and SYSRECs.
- Use the BUILD type load when any of the following conditions apply:
  - None of the indexes are unique.
  - Unique indexes exist, but you are sure that there are no duplicate rows.
  - You are not concerned with restarting jobs.
  - You are sure of the data and you want to load it directly into an empty tablespace.



# CA Fast Load : OUTPUT-CONTROL BUILD

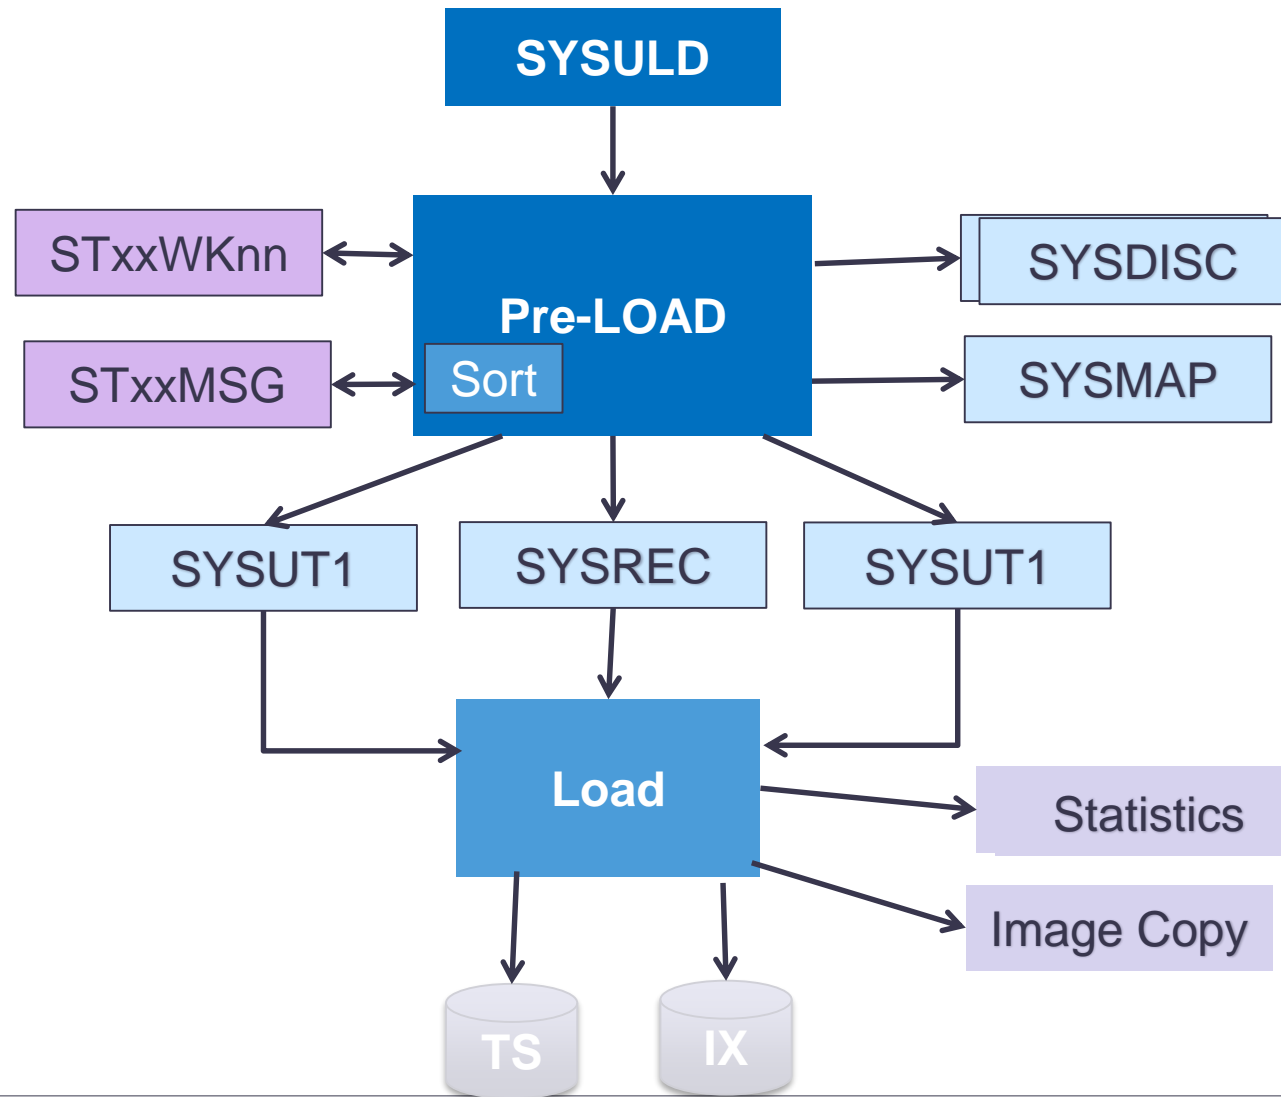


# CA Fast Load: OUTPUT-CONTROL **ALL**

- The ALL format writes data to secondary storage data sets (**classic mode**)
- Does not stop processing if it encounters duplicate keys or out-of-sequence index entries.
- Fully restartable
- Use the ALL type load if any of the following conditions apply:
  - Duplicate keys or out-of-sequence index entries may be present.
  - You want the ability to restart in the LOAD phase.

# CA Fast Load

## OUTPUT-CONTROL ALL



# CA Fast Load: Enhancements

- DB2 11 EXCLUDE NULL KEYS Support
  - Index defined with EXCLUDE NULL KEYS is supported
- Specify the Column Order in a Delimited Input File
  - Field specifications are now honored when INPUT-FORMAT DELIMITED is specified
- DB2 10 Inline LOB Support
  - Inline large object (LOB) support in a base table
- New ESTIMATED-ROWS Option
  - Real-Time Statistics (RTS) to specify the existing number of rows
  - Efficiently allocate sort work data sets and estimate the sort file size
- Additional Real-Time Statistics Support
  - Update the Real-Time Statistics (RTS) tables during processing

# CA Fast Load: Enhancements

- New SORTNUM Parameter
  - SORTWRK file dynamic allocation.
  - SORTNUM AUTO provides dynamic allocation per load, based on #rows and row-size
- MAX TASKS AUTO
  - Maximum number of concurrent sort tasks
  - Based on NPI and Clustering INDEXES
- DISCARD processing
  - Allow for more control by the user
  - RECOVER PENDING has been eliminated
  - Do NOT include WHERE processing discards
- Enhanced COPY-BUFFERS Support
  - Allocate up to 360 I/O buffers when calling CA Quick Copy for DB2 for z/OS

# CA Fast Unload

## Function

- High speed utility for offloading DB2 table data to various sequential file formats.
- It offers superior performance and important functional advantages in the areas of flexible formatting of data and selection of DB2 table data.
- Can process data directly from DB2 tables or can unload from a DB2 image copy.
- Can unload multiple DB2 tables concurrently.
- Unload partitions concurrently
- Can unload data using a sampling technique to aid in building subsets of data.

## Value

- Lower CPU and Elapse Times. Down time is reduced saving money.
- Reduced impact on DB2 resources.

# CA Fast Unload: Benefits

- Much faster and more flexible than DSNTIAUL
- Processes DB2 data directly from TS's or IC's
  - Can unload data directly from DB2 data pages or from a full image copy, DSN1COPY, or a DFSMS concurrent copy
  - Can also unload data from an image copy that no longer matches the current table definition by reading the table's source DDL.
  - Flashcopy support
- Multiple output formats (
  - Supports standard SQL against image copies
  - Supports any valid SELECT statement; you can unload table rows selectively using a SQL WHERE clause, a total unloaded row count, or a sample of  $n$  rows every  $x$  rows
- Unload data by partition, or selectively by using a WHERE clause specifying the data to unload
- Dynamic allocation with LISTDEF
- OBID conversion

# CA Fast Unload: Best Practices

- For recommended settings that apply to the entire installation, set them in the PFU member of *high-level.CDBAPARM*
- Enable sort to dynamically allocate sort work data sets for processing ORDER BY clauses in SELECT statements.
- Specify COPY-BUFFERS, IO-BUFFERS, and VSAM-BUFFERS if the default values are not appropriate for your use.
- Use EXCP MM as the default for maximum speed and flexibility.
- Use CURRENT-DEGREE ANY and MULTI-ROW-FETCH to improve performance with SQL-ACCESS ONLY. MULTI-ROW-FETCH is enabled by default.
- **Use OUTPUT-FORMAT LOAD** for best performance. However, when unloading long variable length columns, OUTPUT-FORMAT VARIABLE is usually more appropriate.



# CA Fast Unload: Enhancements

- LOAD Statement Support with Comma-Delimited Output
  - generate load control statements when specifying OUTPUT-FORMAT COMMA-DELIMITED
- INTO Statement Support with Comma-Delimited Output
  - include INTO statements in your syntax when specifying OUTPUT-FORMAT COMMA-DELIMITED
- Enhanced Literals Support in Comma-Delimited Output
  - NULL-FIELD - Specifies the character string that will represent a null value in the output file.
  - For the following keywords, literal support has been expanded from 1 byte to 4 bytes:
  - COMMA - Specifies which character to use as the field separator.
  - ROW-DELIMITER - Specifies which character to use as the end-of-record delimiter.

# CA Fast Unload: Enhancements

- New FILL Keyword
  - FILL keyword is now provided to specify whether to add leading zeros to a field when converting numeric data types
- New Data Types Support
  - BIGINT EXTERNAL and SMALLINT EXTERNAL
- Enhanced Null Indicator Support
  - specify which null indicator to use in the output and where to place it
- Timestamp Pattern
  - specify a TIMESTAMP pattern for non-timestamp data types
- Improved Performance with Large WHERE Clauses (GA)
  - column IN (value-list)

# CA Fast Unload: Enhancements

- [Use SYSREC01 and SYSCTL01 with ONE-SYSREC](#)
- [New QUOTE-EMBED-SINGLE Keyword](#)
- [DB2 12 Support \(Updated\)](#)
- [Flexible INTO Statement Syntax](#)
- [New Default LRECL Value for Comma-Delimited Output](#)
- [Enhanced Timestamp Formatting with Comma-Delimited Output](#)
- [LOAD Statement Support with Comma-Delimited Output](#)
- [INTO Statement Support with Comma-Delimited Output](#)
- [Enhanced Literals Support in Comma-Delimited Output](#)
- [New FILL Keyword](#)
- [New Data Types Support](#)
- [Enhanced Null Indicator Support](#)
- [Improved Performance with Large WHERE Clauses](#)
- [New EXITS Parameters](#)
- [zIIP Support](#)
- [Enhanced Support for Load Control Statements](#)
- [DB2 10 Temporal Table Support](#)

# CA Database Analyzer

## Function

- Database Analyzer has been referred to as a “DBA in a box”.
- Performs analysis on thousands of DB2 tablespaces and indexes.
- Selectively chooses objects in need of maintenance based on either Database Analyzer collected statistics, or IBM Real Time Statistics.
- Automatically generates IBM or CA utility JCL for common IBM or CA utilities based on user selected thresholds.
  - REORG
  - Image copy
  - RUNSTATS

## Value

- Simplify the DBA day to day responsibilities
- Optimize which objects are acted upon.

# CA Database Analyzer: Benefits

- Sophisticated analysis
  - Gathers and uses comprehensive statistics to provide graphs, trending and forecasts.
- Flexible scripting language
  - Provides customizable predefined rules, conditions and job control language (JCL) templates.
- Efficient processing
  - Utilizes zIIP technology, processes objects concurrently, allocates multiple buffers, exploits extended memory and utilizes fast I/Os.

# CA Database Analyzer: Best Practices

- Retrieving Values Directly from the PDA, RTS, or DB2 catalog tables for CA Database Analyzer Action Procedures
- Monitor the IBM RTS Tables for Missing Information
- Customize your installation
  - Use INCLUDE Statements to Specify Longer Predicate Criteria
  - Use INCLUDE Statements to Create New Action Conditions
- Using Real-Time Statistics with CA Database Analyzer to reduce statistics collections
- Using a Combination of IBM RUNSTATS and CA Database Analyzer Extract Procedures
- Using RTOS (Real Time Object Selection) to Dynamically Select Objects for Maintenance
- Using RTOS to Reorganize Hash PBGs

# CA Database Analyzer: Best Practices

## Action procedure: Utilities

- Sequence of up to 10 actions (utilities) to happen when predefined conditions are true.
  - Standard DB2 utilities (such as an image copy or REORG)
  - CA-defined utilities
  - Command Processor access
  - User-defined applications
- Requires to be tied to an “extract” procedure .

## Extract procedure: Objects

- specifies which objects to gather statistics (RUNSTATS)
- Specifies which objects to process. An extract procedure can access any combination of databases, tablespaces, indexspaces, and tables residing on the same DB2 subsystem
- Allows “filtering” criteria based on profiles
  - Object **Selection** Profile (OSP)
  - **Prioritize** Object Maintenance (POM)

# CA Database Analyzer: Enhancements

- Improved Control Over Job Name Generation
  - When generating job names, you can now control where numeric sequencing occurs within the mask. Characters that are not part of the sequencing digits are preserved. This enhancement makes it easier to generate unique job names and comply to site-specific job naming standards. Masks up to eight characters are supported.
- Simplified Object Selection for RTOS-based Extract Procedures (ACTION oriented)
  - An Object Selection Profile (OSP) simplifies the process of selecting objects that you want to include in or exclude from an RTOS extract procedure.
  - Previously, users could limit the object selection only by adding a #INCLUDE on each Extract Object Condition, specifying an EQF, or both.
- Prioritized Object Maintenance (POM)
  - POM will prioritize (processing order) objects in action JCL so that the most critical objects will be processed based on selected POCs (Prioritize Object Conditions).
  - POCs are different for every action (utility). Use standard recommendations based on stats values.
  - Object Sets (Tablespaces and Indexes) are ranked based on the most critical criteria



# CA Database Analyzer: Last minute enhancements...

- **New Action Conditions for DSNACCOX**
  - The following tablespace conditions are now available and can be used at the partition level:
    - (TSP) RTS DSNACCOX full image copy
    - (TSP) RTS DSNACCOX incremental image copy
    - (TSP) RTS DSNACCOX RUNSTATS recommended
  - The following indexspace conditions are now available and can be used at the partition level:
    - (IXP) RTS DSNACCOX full image copy
    - (IXP) RTS DSNACCOX REORG recommended
    - (IXP) RTS DSNACCOX RUNSTATS recommended

# CA Rapid Reorg

## Function

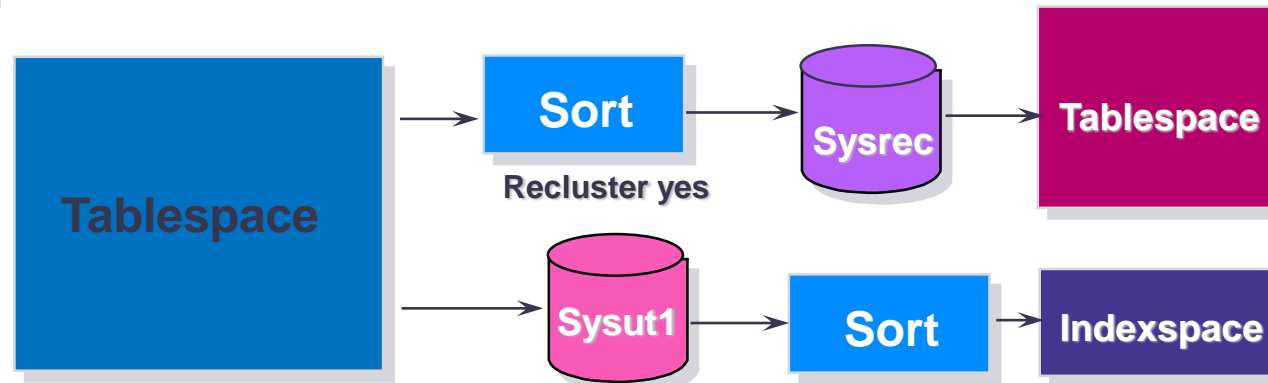
- Perform quick and effective DB2 data reorganizations
- Online reorg capability enables read and update operations to tables during the reorganization process (DATA-AVAILABLE NORMAL/HIGH/**CRITICAL**....OFF/RO/**RW**)
- Automatic space allocation facilities make it easy to redefine data sets during the reorganization
- Produce up to eight image copies and collect and update DB2 catalog statistics during the reorg

## Value

- Increase data availability and performance and save resources.
- Application access to data is left unaffected during online reorgs.
- Reduce CPU time, I/O activity and the costs associated with downtime due to reorgs

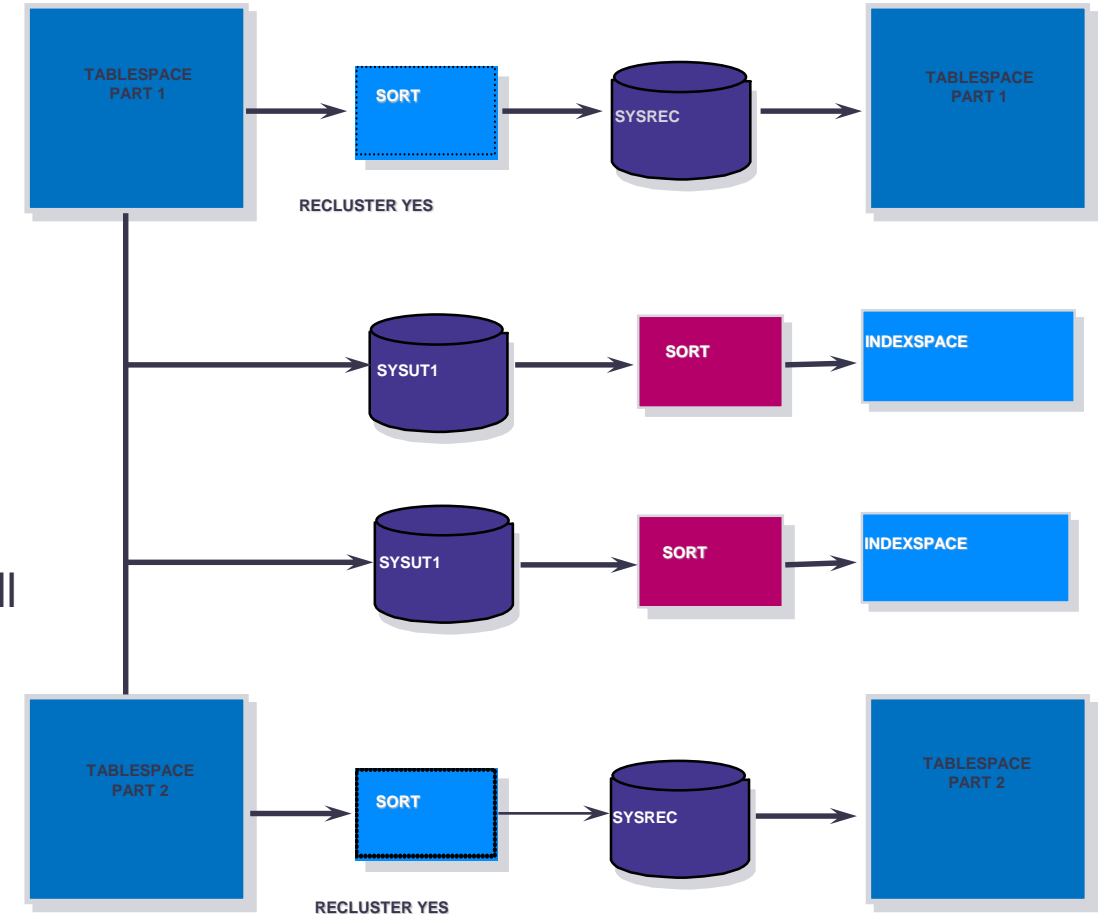
# CA Rapid Reorg

- Performs full reorganization including reclaiming of PCTFREE and PAGEFREE space, reclustering of rows and full data and index loads
- IBM REORG syntax is compatible with Rapid Reorg syntax
- Automatic space calculation and redefinition of both STOGROUP and VSAM datasets. No PAUSE is required to adjust space allocations.
- Optionally update the DB2 catalog or PDA statistics without running RUNSTATS
- Multi-tasking capabilities available in each phase of reorg
  - Partition Independence



# CA Rapid Reorg

- Select or delete rows from the tablespace during reorganization. Deleted rows can be saved in a dataset.
- Integrated image copy allows for up to eight image copies
- Dynamically allocates work and sortwork datasets
- Check pending and copy pending flags set
- Control over the number of concurrent tasks during all phases of reorg
- Interleave option will sort data in cluster key order
- Integrated with CA Fastload for LOAD RESUME and LOAD REPLACE processing



# CA Rapid Reorg: Online Reorg

## Function

- Read/Write accessibility during TS/IX reorg (DATA-AVAILABLE CRITICAL)
- Reduces both CPU usage and elapsed time
- Less Scanning of log records, No I/Os to the log dataset, Faster Renames
- Fully supports user-managed datasets by automating the naming and deletion processes
- Utilizes new IBM API to remove QUIESCE 911

## Value

- No application downtime is encountered – Fast Switch
- Reorgs can happen anytime of the day

# CA Rapid Reorg: Online reorg. Output data sets dynamic allocation

Function: usability

- **SYSREC**
  - Stores tablespace unload data.
- **SYSUT1**
  - Stores indexspace unload data.
- **SYSDISC**
  - Stores discard data.
- **RID**
  - Stores data in RIDNNNN work datasets for RID entries during log apply phase.
- **hlq.CDBAPARM(PRR):** standard/default naming conventions
- **PRR-SUFFIX (xxxxx):** alternate/specific for different business/applications.
- **Symbolics** allowed: %DATE, %DBNAME, %DSNUM, %INCR, %TIME, %SSID

# Questions?

...

# Thanks, Merci, Gracias, Dekuji,....



## Manuel Gómez Burriel

DB2 Utilities Product Owner

[ManuelAnibal.GomezBurriel@ca.com](mailto:ManuelAnibal.GomezBurriel@ca.com)



in



# Call for Speakers Now Open



Share your experience. Be the teacher. Enhance your resume.

**Register your session today!**



Learn more: [ca.com/caworld](https://ca.com/caworld)

© 2018 CA. ALL RIGHTS RESERVED.