

CA Database Management Solutions for DB2 for z/OS

Release Notes

Version 17.0.00, Sixth Edition



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, 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 LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2013 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This document references the following CA Technologies products:

- CA ACF2™ for z/OS (CA ACF2)
- CA Bind Analyzer™ for DB2 for z/OS (CA Bind Analyzer)
- CA Compile/PRF for CICS and TSO (CA Compile/PRF)
- CA Application Performance Management (CA APM)
- CA Database Analyzer™ for DB2 for z/OS (CA Database Analyzer)
- CA Detector® for DB2 for z/OS (CA Detector)
- CA Endeavor® Software Change Manager (CA Endeavor SCM)
- CA Endeavor® Software Change Manager Interface for DB2 for z/OS (CA Endeavor SCM Interface for DB2)
- CA Fast Check® for DB2 for z/OS (CA Fast Check)
- CA Fast Index® for DB2 for z/OS (CA Fast Index)
- CA Fast Load for DB2 for z/OS (CA Fast Load)
- CA Fast Recover™ for DB2 for z/OS (CA Fast Recover)
- CA Fast Unload® for DB2 for z/OS (CA Fast Unload)
- CA Index Expert™ for DB2 for z/OS (CA Index Expert)
- CA InfoRefiner® (CA InfoRefiner)
- CA InfoTransport® (CA InfoTransport)
- CA Insight™ Database Performance Monitor for DB2 for z/OS (CA Insight DPM)
- CA Log Analyzer™ for DB2 for z/OS (CA Log Analyzer)
- CA LogCompress™ for DB2 for z/OS (CA LogCompress)
- CA Chorus™ (CA Chorus)
- CA Chorus™ Software Manager (CA CSM)
- CA Merge/Modify™ for DB2 for z/OS (CA Merge/Modify)
- CA NSM
- CA NSM Database Option for DB2 for z/OS
- CA OPS/MVS® Event Management and Automation (CA OPS/MVS)
- CA Partition Expert™ for DB2 for z/OS (CA Partition Expert)
- CA Plan Analyzer® for DB2 for z/OS (CA Plan Analyzer)
- CA Quick Copy for DB2 for z/OS (CA Quick Copy)

- CA Rapid Reorg® for DB2 for z/OS (CA Rapid Reorg)
- CA RC/Compare™ for DB2 for z/OS (CA RC/Compare)
- CA RC/Extract™ for DB2 for z/OS (CA RC/Extract)
- CA RC/Migrator™ for DB2 for z/OS (CA RC/Migrator)
- CA RC/Query® for DB2 for z/OS (CA RC/Query)
- CA RC/Secure™ for DB2 for z/OS (CA RC/Secure)
- CA RC/Update™ for DB2 for z/OS (CA RC/Update)
- CA Recovery Analyzer™ for DB2 for z/OS (CA Recovery Analyzer)
- CA Report Facility for CICS, IMS, and TSO (CA Report Facility)
- CA SQL-Ease® for DB2 for z/OS (CA SQL-Ease)
- CA Subsystem Analyzer for DB2 for z/OS (CA Subsystem Analyzer)
- CA Top Secret® for z/OS (CA Top Secret)

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at <http://ca.com/docs>.

Documentation Changes

The following updates have been made in the sixth edition of this documentation:

- [Support for Automated Switching to EXCP NO](#) (see page 68) in CA Quick Copy—Added this new topic.
- [DB2 10 MEMBER CLUSTER option support for universal tablespaces](#) (see page 78) in CA RC/Migrator—Added this new topic.

The following updates have been made in the fifth edition of this documentation:

- [Optimizer Recommended RUNSTATS Support](#) (see page 21) in CA Database Analyzer—Added this new topic.
- [Automatic Space Allocations with Models MJUTLAL and MJUTLAL1](#) (see page 23) in CA Database Analyzer—Added this new topic.
- [Temporal Table Support](#) (see page 72) in CA RC/Compare—Updated the description of this enhancement indicating that you can add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint when altering a table in a compare strategy.
- [Compare Rules Enhancement to Trigger Attributes](#) (see page 74)—Added this new topic.
- [Temporal Table Support](#) (see page 78) in CA RC/Migrator—Updated the description of this enhancement indicating that you can add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint when creating, altering, or templating a table in a compare or migrate strategy.
- [Temporal Table Support](#) (see page 87) in CA RC/Update—Updated the description of this enhancement indicating that you can add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint when creating, altering, or templating a table.

The following updates have been made in the fourth edition of this documentation:

- [DB2 11 Support](#) (see page 14)—Added this new topic that describes the DB2 11 support that has been added for the CA Database Management Solutions for DB2 for z/OS. For detailed information about the DB2 11 support that each product provides, see the product-specific enhancement sections.
- [IBM z/OS 2.1 Support](#) (see page 16)—Added this new topic.

In addition to DB2 11 support, the following product-specific enhancements have been made in the fourth edition:

- [Performance Improvements](#) (see page 25) in CA Database Analyzer—Added this new topic.
- Added the following new topics in CA Fast Unload:

- [DB2 10 Temporal Table Support](#) (see page 38)
- [New PART-SEPARATE Keyword](#) (see page 39)
- [Temporal Table Support](#) (see page 72) in CA RC/Compare—Updated the description of this enhancement indicating that you can now alter an index.
- Added or updated the following topics in CA RC/Query:
 - [Index Report Enhancements](#) (see page 86)
 - [Routine Report Enhancements](#) (see page 86)
 - [Table Report Enhancements](#) (see page 86)
 - [Wildcard Character Support](#) (see page 87)
 - [SPUFI Compatible HDDL Output Support](#) (see page 85)—Updated the description of this enhancement indicating that HDDL object selection is now possible.
 - [ALL Command Support](#) (see page 85)—Updated the description of this enhancement indicating that CA RC/Query can now automatically apply changes to more than 255 selected objects.
- Added or updated the following topics in CA RC/Migrator:
 - [AUTOREPAIR option in RC/Merger](#) (see page 80)
 - [Set SQL Terminator in the Batch Processor](#) (see page 82)
 - [PBG Tablespace Enhancements](#) (see page 81)
 - [Temporal Table Support](#) (see page 78) in CA RC/Migrator—Updated the description of this enhancement indicating that you can now alter an index.
- [Temporal Table Support](#) (see page 87) in CA RC/Update—Updated the description of this enhancement indicating that you can now alter an index.

The following updates have been made in the third edition of this documentation:

- [zIIP Support](#) (see page 16)—Added this new topic.
- [Reclaim Version Support](#) (see page 56) in CA Merge/Modify—Added this new topic.
- [Temporal Table Support](#) (see page 72) in CA RC/Compare—Updated the description of this enhancement indicating that you can now create or template an index.
- [Temporal Table Support](#) (see page 78) in CA RC/Migrator—Updated the description of this enhancement indicating that you can now create or template an index.
- [Temporal Table Support](#) (see page 87) in CA RC/Update—Updated the description of this enhancement indicating that you can now create or template an index.

The following updates have been made in the second edition of this documentation:

- [DB2 10 IFCID 363 Support](#) (see page 45)—Added this enhancement to CA Insight DPM.

Contents

Chapter 1: Enhancements	13
DB2 11 Support	14
DB2 10 Support	14
DB2 9 Support	15
IBM z/OS 2.1 Support	16
zIIP Support	16
Software Configuration Service Support	16
CA Bind Analyzer Enhancements	17
DB2 11 Support	17
z/OS XL C Compiler Support	17
CA Compile/PRF Enhancements	17
CA Database Analyzer Enhancements	17
DB2 11 Support	18
DB2 11 Archive Table Support	18
DB2 11 EXCLUDE NULL KEYS Support	19
DB2 11 Extended Logging Support	19
DB2 11 NOT LOGGED Declared Global Temporary Table Support	20
DB2 11 Optimizer Recommended RUNSTATS Support	21
DB2 10 and DB2 11 IBM RUNSTATS USE PROFILE Support	22
DB2 10 Temporal Table Support	22
Automatic Space Allocations with Models MJUTLAL and MJUTLAL1	23
New Action Conditions to Skip Migrated Objects	23
New Columns in Statistics Tables	24
Performance Improvements	25
REBIND PACKAGE Support for Stored Procedures	25
CA Detector Enhancements	25
DB2 11 Support	25
View the Collection Interval Size	26
CA Endeavor SCM Interface for DB2 Enhancements	26
DB2 11 Support	26
New ISPF Panel Interface Support	26
Associations Support	27
BACK BIND Support	27
Batch Utility Footprint Support	27
Element Action Backout Support	28
CA Fast Check Enhancements	28
DB2 11 Support	28

DB2 11 Extended Logging Support	28
DB2 10 Compress on INSERT Support	28
DB2 10 Temporal Table Support	29
CA Fast Index Enhancements	29
DB2 11 Support	29
DB2 11 EXCLUDE NULL KEYS Support.....	30
DB2 10 Compress on INSERT Support	30
DB2 10 Temporal Table Support	30
CA Fast Load Enhancements	31
DB2 11 Support	31
DB2 11 EXCLUDE NULL KEYS Support.....	31
DB2 11 Extended Logging Support.....	31
DB2 10 Inline LOB Support	32
DB2 10 MEMBER CLUSTER Support in Universal Tablespaces.....	32
DB2 10 Temporal Table Support	32
LOB 2 GB Column Support	33
New Partition Independence Keywords	33
zIIP Support	34
CA Fast Recover Enhancements	34
DB2 11 Support	34
DB2 11 Extended Logging Support.....	34
DB2 10 EAV Support.....	35
DB2 10 FlashCopy Support	35
DB2 10 Temporal Table Support	35
CA Fast Unload Enhancements	36
DB2 11 Support	36
DB2 11 Extended Logging Support.....	37
DB2 10 FlashCopy Support	37
DB2 10 Inline LOB Support	37
DB2 10 Temporal Table Support	38
Dynamic Allocation Support.....	38
Native LOB Support	39
Unload Objects in Descending Size Order	39
New PART-SEPARATE Keyword	39
New XMLDB2V9 Parameter for OUTPUT-FORMAT Keyword.....	40
zIIP Support	40
CA Insight DPM Enhancements	40
DB2 11 Support	41
DB2 11 Extended Logging Support.....	41
DB2 11 IFCID 378 AND 379 Support.....	42
DB2 10 IFCID 316 and 401 Support	42
DB2 10 IFCID 363 Support	45

DB2 10 IFCID 365 Support	45
DB2 10 Statement ID-Based Trace Record Support	45
Cumulative Data Sharing Statistics for CA Cross-Enterprise APM.....	46
Expanded Sort Metrics Support	47
Extended Print Format for Intervals.....	47
IBM DB2 Accelerator Support	48
Improved Lock Contention Display	48
Improved I/O Trace Record Processing.....	49
Improved Package Information on the Current and History Package List Panels	49
Improved System Condition Monitor Functionality.....	50
New IQL Record (1044) APM-APPL-SUM	50
New SORT Exit Routines.....	50
CA Log Analyzer Enhancements	51
DB2 11 Support	51
DB2 11 Extended Logging Support.....	52
DB2 10 Creation of Partition-by-Growth Tablespaces with NUMPARTS and MAXPARTITIONS	52
DB2 10 DEFINE NO Support for LOB and XML Tablespaces	52
DB2 10 Temporal Table Support	53
DB2 10 TIMESTAMP Support	53
DB2 10 Unique Index INCLUDE Column Support	54
New DDL Activity Object Report	54
CA Merge/Modify Enhancements	55
DB2 11 Support	55
DB2 11 Extended Logging Support.....	55
DB2 10 EAV Support.....	55
DB2 10 FlashCopy Support.....	56
DELETE Options for Image Copy Records.....	56
CA Merge/Modify Reclaim Version Support	56
CA NSM Database Option for DB2 for z/OS Enhancements.....	56
CA Plan Analyzer Enhancements.....	57
DB2 11 Support	57
DB2 10 Index Probing Support	57
DB2 10 Non-Inline SQL Scalar Function Package Support.....	58
DB2 10 Row and Column Access Control Support	59
DB2 10 Statement-Level Optimization Hint Support	60
DB2 10 XMLXSROBJECTID Support.....	61
DB2 9 XMLCAST Specification Clause Support	61
Conditionally Indent SQL Substatements.....	61
Identify Non-SQL Access Path Changes.....	62
Improved Cost Difference Filters for the Compare Versions Cost Compare Report.....	62
New DB2 Profile Services Facility	63
Create Explain Table Indexes Automatically During a Future Explain.....	64

Update all Explain Tables with Source SQL Catalog Data During a Future Explain	64
New Physical Rule 1083	64
New Reporting Services QUERY Reporting Option	65
CA Quick Copy Enhancements	66
DB2 11 Support	67
DB2 11 Extended Logging Support	67
DB2 10 EAV Support	67
DB2 10 FlashCopy Support	67
DB2 10 Temporal Table Support	67
Support for Automated Switching to EXCP NO	68
CA Rapid Reorg Enhancements	68
DB2 11 Support	68
DB2 11 EXCLUDE NULL KEYS Support	69
DB2 11 Extended Logging Support	69
DB2 10 MEMBER CLUSTER Support in Universal Tablespaces	69
DB2 10 Temporal Table Support	70
New Partition Independence Keywords	70
zIIP Support	71
CA RC/Compare Enhancements	71
DB2 11 Support	71
DB2 10 Temporal Table Support	72
DB2 10 TIMESTAMP Support	72
DB2 10 Unique Index INCLUDE Column Support	73
DB2 9 Implicitly Hidden Columns Support	73
DB2 9 ROW CHANGE TIMESTAMP Support	73
DB2 9 Clone Support	73
Extended DB2 9 Native SQL Stored Procedures Support	74
Compare Rules Enhancement to Trigger Attributes	74
CA RC/Extract Enhancements	74
DB2 11 Support	74
DB2 10 Temporal Table Support	74
DB2 10 TIMESTAMP Support	75
DB2 10 Unique Index INCLUDE Column Support	75
Enhanced Sort DYNALLOC Options	75
Hardware Compression Support	76
Isolation Level Support for SQL Access	76
LIMIT and WHERE Clause Support	76
Named Relationship Sets	76
Split Extract Unit Count Support	76
Alias Confirmation Removed	77
CA RC/Migrator Enhancements	77
DB2 11 Support	77

DB2 11 Extended Logging Support	78
DB2 10 MEMBER CLUSTER Option Support for Universal Tablespaces	78
DB2 10 Temporal Table Support	78
DB2 10 TIMESTAMP Support	79
DB2 10 Unique Index INCLUDE Column Support	79
DB2 9 Clone Support	79
DB2 9 Histogram Statistics Support	79
DB2 9 Implicitly Hidden Columns Support	80
Extended DB2 9 Native SQL Stored Procedures Support	80
DB2 9 ROW CHANGE TIMESTAMP Support	80
AUTOREPAIR Option in RC/Merger	80
Parallel Unload and Load Support	81
PBG Tablespace Enhancements	81
Set SQL Terminator in the Batch Processor	82
CA RC/Query Enhancements	82
DB2 11 Support	82
DB2 11 Extended Logging Support	83
DB2 10 Temporal Table Support	83
DB2 9 Roles and Trusted context Support	84
DB2 9 ROW CHANGE TIMESTAMP Support	84
ALL Command Support	85
SPUFI-Compatible HDDL Output Support	85
Index Report Enhancements	86
Routine Report Enhancements	86
Table Report Enhancements	86
Wildcard Character Support	87
CA RC/Secure Enhancements	87
CA RC/Update Enhancements	87
DB2 11 Support	87
DB2 10 Temporal Table Support	87
DB2 10 TIMESTAMP Support	88
DB2 10 Unique Index INCLUDE Column Support	88
DB2 9 Clone Support	88
DB2 9 Roles and Trusted Context Support	89
DB2 9 ROW CHANGE TIMESTAMP Support	89
Extended DB2 9 Native SQL Stored Procedures Support	89
CA Recovery Analyzer Enhancements	89
DB2 11 Support	90
DB2 11 Extended Logging Support	90
DB2 10 Temporal Table Support	90
New Block Size Field for Disaster Recovery Strategies	91
CA Report Facility Enhancements	91

DB2 11 Support	91
RUN PROC Command within a Procedure	92
Define the Procedure Error Threshold	92
Specify the KEEP Column Name Suffix	92
Specify the Precedence of Sources for Variable Resolution	93
CA SQL-Ease Enhancements	94
DB2 11 Support	94
DB2 10 Index Probing Support	94
DB2 10 Row and Column Access Control Support	95
New Create Explain Table(s) Facility	95
Conditionally Indent SQL Substatements	96
New DB2 Profile Services Facility	96
Create Explain Table Indexes Automatically During a Future Explain	97
New Physical Rule 1083	98
CA Subsystem Analyzer Enhancements	98
DB2 11 Support	98
View the Collection Interval Size	98
Value Pack Enhancements	99
DB2 11 Support	99
DB2 10 DSNZPARM Support	99
DB2 10 TIMESTAMP Support	101
Increased Active Log Data Set Support	101
Data Sharing Support for -DISPLAY CA ACTIVE Command	101
Named Relationship Sets	102

Chapter 2: Documentation 103

CA HTML Bookshelf	103
Access the Bookshelf	103
Access the Guides	104
Search the Bookshelf	104
Common Guides	105
Product-Specific Documentation	105
Backup and Recovery	105
Database Administration	106
Database Performance Management	106
Reporting and Information Management	107
Edition Numbers	107

Chapter 1: Enhancements

Welcome to the *CA Database Management Solutions for DB2 for z/OS Release Notes* for Version 17.0. This guide describes new enhancements, known issues, published solutions, and the product documentation set.

Note: For more information about system requirements, and installation and upgrade considerations, see the *Installation Guide*.

This section contains the following topics:

- [DB2 11 Support](#) (see page 14)
- [DB2 10 Support](#) (see page 14)
- [DB2 9 Support](#) (see page 15)
- [IBM z/OS 2.1 Support](#) (see page 16)
- [zIIP Support](#) (see page 16)
- [Software Configuration Service Support](#) (see page 16)
- [CA Bind Analyzer Enhancements](#) (see page 17)
- [CA Compile/PRF Enhancements](#) (see page 17)
- [CA Database Analyzer Enhancements](#) (see page 17)
- [CA Detector Enhancements](#) (see page 25)
- [CA Endeavor SCM Interface for DB2 Enhancements](#) (see page 26)
- [CA Fast Check Enhancements](#) (see page 28)
- [CA Fast Index Enhancements](#) (see page 29)
- [CA Fast Load Enhancements](#) (see page 31)
- [CA Fast Recover Enhancements](#) (see page 34)
- [CA Fast Unload Enhancements](#) (see page 36)
- [CA Insight DPM Enhancements](#) (see page 40)
- [CA Log Analyzer Enhancements](#) (see page 51)
- [CA Merge/Modify Enhancements](#) (see page 55)
- [CA NSM Database Option for DB2 for z/OS Enhancements](#) (see page 56)
- [CA Plan Analyzer Enhancements](#) (see page 57)
- [CA Quick Copy Enhancements](#) (see page 66)
- [CA Rapid Reorg Enhancements](#) (see page 68)
- [CA RC/Compare Enhancements](#) (see page 71)
- [CA RC/Extract Enhancements](#) (see page 74)
- [CA RC/Migrator Enhancements](#) (see page 77)
- [CA RC/Query Enhancements](#) (see page 82)
- [CA RC/Secure Enhancements](#) (see page 87)
- [CA RC/Update Enhancements](#) (see page 87)
- [CA Recovery Analyzer Enhancements](#) (see page 89)
- [CA Report Facility Enhancements](#) (see page 91)
- [CA SQL-Ease Enhancements](#) (see page 94)
- [CA Subsystem Analyzer Enhancements](#) (see page 98)
- [Value Pack Enhancements](#) (see page 99)

DB2 11 Support

The products that comprise the CA Database Management Solutions for DB2 for z/OS have been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

Note: This support is not currently available for CA LogCompress and CA Partition Expert.

In addition to DB2 11 NFM and CM support, the following DB2 11 functionality is also provided:

- Archive table support in CA Database Analyzer
- EXCLUDE NULL KEYS support in CA Database Analyzer, CA Fast Index, CA Fast Load, and CA Rapid Reorg
- Extended logging support in CA Database Analyzer, CA Fast Check, CA Fast Load, CA Fast Recover, CA Fast Unload, CA Insight DPM, CA Log Analyzer, CA Merge/Modify, CA Quick Copy, CA Rapid Reorg, CA RC/Migrator, CA RC/Query, and CA Recovery Analyzer
- IFCID 378 and 379 support in CA Insight DPM
- NOT LOGGED Declared Global Temporary Table (DGTT) support in CA Database Analyzer
- Optimizer recommended RUNSTATS support in CA Database Analyzer
- USE PROFILE support in CA Database Analyzer

Note: For more information about these enhancements, see the product-specific enhancement sections.

DB2 10 Support

The following additional DB2 10 support has been added to the CA Database Management Solutions for DB2 for z/OS:

- Compress on INSERT support in CA Fast Check and CA Fast Index
- Creation of partition-by-growth tablespaces with Numparts and Maxpartitions support in CA Log Analyzer
- DEFINE NO support for LOB and XML tablespaces in CA Log Analyzer
- DSNZPARM support in Thread Termination\Dynamic DSNZPARM
- EAV support in CA Fast Recover, CA Merge/Modify, and CA Quick Copy
- FlashCopy support in CA Fast Recover, CA Fast Unload, CA Merge/Modify, and CA Quick Copy

- IFCID 316, 363, 365, and 401 support in CA Insight DPM
- INCLUDE column support in CA RC/Compare, CA RC/Migrator, and CA RC/Update
- Index probing support in CA Plan Analyzer and CA SQL-Ease
- Inline LOB support in CA Fast Load and CA Fast Unload
- MEMBER CLUSTER support in universal tablespaces in CA Fast Load and CA Rapid Reorg
- Non-inline SQL scalar function package support in CA Plan Analyzer
- Row and column access control support in CA Plan Analyzer and CA SQL-Ease
- Statement ID-based trace record support in CA Insight DPM
- Statement-level optimization hint support in CA Plan Analyzer
- Temporal table support in CA Fast Check, CA Fast Index, CA Fast Load, CA Fast Recover, CA Fast Unload, CA Log Analyzer, CA Quick Copy, CA Rapid Reorg, CA RC/Compare, CA RC/Extract, CA RC/Migrator, CA RC/Query, CA RC/Update, and CA Recovery Analyzer
- Additional temporal table support in CA Database Analyzer
- TIMESTAMP with greater precision support in CA Log Analyzer, CA RC/Compare, CA RC/Extract, CA RC/Migrator, CA RC/Update, and the Value Pack Batch Processor component
- TIMESTAMP WITH TIME ZONE support in CA Log Analyzer, CA RC/Compare, CA RC/Extract, CA RC/Migrator, CA RC/Update, and the Value Pack Batch Processor component
- Unique index INCLUDE column support in CA Log Analyzer, CA RC/Compare, CA RC/Extract, CA RC/Migrator, and CA RC/Update
- USE PROFILE support in CA Database Analyzer
- XML support in CA Plan Analyzer for XMLXSROBJECTID

Note: For more information about these enhancements, see the enhancements list for each product.

DB2 9 Support

The following additional DB2 9 support has been added to the CA Database Management Solutions for DB2 for z/OS:

- Histogram statistics support in CA RC/Migrator
- Implicitly hidden columns support in CA RC/Compare and CA RC/Migrator
- Extended clone support in CA RC/Compare, CA RC/Migrator, and CA RC/Update

- Full support for native stored procedures (NSPs) with multiple versions, including group ICL support, in CA RC/Compare, CA RC/Migrator, and CA RC/Update
- Roles and trusted context support in CA RC/Query and CA RC/Update
- ROW CHANGE TIMESTAMP column support in CA RC/Compare, CA RC/Migrator, CA RC/Query, and CA RC/Update
- XMLCAST specification clause support in CA Plan Analyzer
- XML PATTERN support in CA RC/Migrator and CA RC/Update

Note: For specific enhancements and limitations, see the product-specific enhancement sections.

IBM z/OS 2.1 Support

IBM z/OS 2.1 support is now provided. The products that comprise the CA Database Management Solutions for DB2 for z/OS have been enhanced to run under IBM z/OS 2.1.

zIIP Support

IBM System z Integrated Information Processor (zIIP) support is now provided for the following products:

- CA Fast Load
- CA Fast Unload
- CA Rapid Reorg

Additionally, CA Insight DPM provides zIIP CPU time package information, in tabular format, on the Current Package/DBRM List, Package/DBRM History List, and Package/DBRM History Summary List panels.

Software Configuration Service Support

You can now use the Software Configuration Service (SCS) component of CA Chorus Software Manager to configure products that you have already acquired, installed, and deployed through CA Chorus Software Manager. This functionality guides you through the post-installation configuration process and indicates which configuration steps to perform manually.

Note: CA Chorus Software Manager Release 4.1 or higher is required. For more information about using CA Chorus Software Manager to configure your products, see the *CA Database Management Solutions for DB2 for z/OS Installation Guide*.

CA Bind Analyzer Enhancements

The following enhancements have been made to CA Bind Analyzer with this release:

- DB2 11 NFM and CM support is now provided.
- z/OS XL C compiler support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

z/OS XL C Compiler Support

The CA Bind Analyzer postcompiler module now supports the z/OS XL C compiler with the METAL and SQL options enabled. The postcompiler is a subcomponent of the Binder function. The postcompiler can now use DBRM changes to determine whether DB2 binds are required for a Metal C source member.

This enhancement helps reduce overhead by limiting the number of BINDs executed for Metal C host language programs. Limiting the BINDs helps stabilize access paths and SQL performance.

Note: For more information about how the postcompiler works, see the *CA Bind Analyzer User Guide*.

CA Compile/PRF Enhancements

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

CA Database Analyzer Enhancements

The following enhancements have been made to CA Database Analyzer with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 archive table support is now provided.
- DB2 11 EXCLUDE NULL KEYS support is now provided.
- DB2 11 extended logging support is now provided.

- DB2 11 NOT LOGGED Declared Global Temporary Table (DGTT) support is now provided.
- DB2 11 optimizer recommended RUNSTATS support is now provided.
- DB2 10 and DB2 11 IBM RUNSTATS USE PROFILE support is now provided.
- Additional DB2 10 temporal table support is now provided.
- Automatic space allocation is now defined in models MJUTLAL and MJUTLAL1.
- New action conditions to skip migrated objects are now provided.
- New columns have been added to the statistics tables.
- Performance improvements have been added.
- REBIND PACKAGE support is now provided for stored procedures.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Archive Table Support

DB2 11 archive table support is now provided. Archive tables can automatically store rows that are deleted from archive-enabled tables, and you can choose whether to include archive tables in your SQL queries. Archive tables are useful for managing historical data.

CA Database Analyzer can now process archive tables. For example, CA Database Analyzer can gather statistics for these objects, include them in action processing, and audit these objects for physical integrity.

The existing model members MJUTLICS and MJUTLRCS demonstrate how you can use the new %ARCRLATD symbolic. This symbolic represents the relationship the currently processed object has to archive tables. The symbolic value determines whether LISTDEF INCLUDE ARCHIVE control statements are added to the generated JCL.

Note: For more information about this symbolic, see the online help. Select **T** (Tutorial) from the CA Database Analyzer main menu, then **7** (Symbolic Variable Descriptions). Place your cursor on a symbolic and press Enter to view its description.

DB2 11 EXCLUDE NULL KEYS Support

DB2 11 EXCLUDE NULL KEYS support is now provided. When an index is defined with EXCLUDE NULL KEYS, DB2 creates an index entry only when a key column contains a value other than NULL. If every key column contains the NULL value, no index entry is created. This option reduces the size of the index, which can improve the performance of your index scans.

CA Database Analyzer can now process indexes that are defined with EXCLUDE NULL KEYS and adjust the related index statistics, such as HIGHKEY, LOWKEY, and COLCARD. CA Database Analyzer can also adjust the multi-column group frequency and cardinality statistics.

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Database Analyzer can now process DB2 objects that have extended RBA and LRSN values. For example, CA Database Analyzer can gather statistics for these objects, include them in action processing, and audit these objects for physical integrity.

CA Database Analyzer provides the following new action conditions (ACs). These conditions identify whether objects are in basic RBA format or DB2 11 extended RBA format. Use them to include or exclude these objects when generating action JCL:

- (IX) RTS include extended RBA
- (IX) RTS include basic RBA objects
- (TS) RTS include extended RBA objects
- (TS) RTS include basic RBA objects
- (IXP) RTS include extended RBA objects
- (IXP) RTS include basic RBA objects
- (TSP) RTS include extended RBA objects
- (TSP) RTS include basic RBA objects

You can also use the following new extract object conditions (EOCs) to identify the object format and include or exclude those objects during processing:

- (IX) RTS include extended RBA
- (IX) RTS include basic RBA objects
- (TS) RTS include extended RBA objects
- (TS) RTS include basic RBA objects

You can also extend the following existing ACs or EOCs with new #INCLUDE model members MCEAC60 (TS) and MCEAC61 (IX). These models extend the SQL predicates of the existing conditions to include objects with basic RBAs during processing:

- (IX) RTS all: use excludes/cond preds
- (IX) RTS (any of the other IX conds)
- (TS) RTS all: use excludes/cond preds
- (TS) RTS (any of the other TS conds)

Note: For more information about these ACs and EOCs, see the online help. On the Action Conditions or Extract Conditions panel, type **H** (Help) next to the condition and press Enter. For more information about #INCLUDE processing, see the online help. Select **T** (Tutorial) from the CA Database Analyzer main menu, then **10** (Glossary Terms), then follow the #INCLUDE hyperlink in A1 Action Conditions or E4 Extract Object Conditions.

When using the existing model member MJUTLROO to create mapping tables for reorganizations, you can also use the following new symbolics. These symbolics identify the RBA format of a whole tablespace or an individual partition:

- %RBAFMTPT (for individual partitions)
- %RBAFMPTS (for whole tablespaces)

Note: For more information about these symbolics, see the online help. Select **T** (Tutorial) from the CA Database Analyzer main menu, then **7** (Symbolic Variable Descriptions). Place your cursor on a symbolic and press Enter to view its description.

DB2 11 NOT LOGGED Declared Global Temporary Table Support

DB2 11 NOT LOGGED Declared Global Temporary Table (DGTT) support is now provided. DGTTs are temporary tables that are defined for the current DB2 connection. Starting with DB2 11, you can now specify a NOT LOGGED clause on a DGTT table so that changes to this table are not logged. Avoiding unnecessary logging can improve system performance.

CA Database Analyzer uses DGTTs in Real Time Object Selection (RTOS) Extracts and now takes advantage of the NOT LOGGED clause on DGTTs.

DB2 11 Optimizer Recommended RUNSTATS Support

DB2 11 optimizer recommended RUNSTATS statistics are now supported. When the optimizer selects access paths, it recognizes when relevant statistics are missing or conflicting. This information is written to the SYSIBM.SYSSTATFEEDBACK table and broken down by statistics type and reason. Statistics collection utilities can then collect the missing statistics so that the optimizer can select efficient access paths.

CA Database Analyzer now provides two new extract object conditions (EOCs) and action conditions (ACs) that evaluate the information in the SYSSTATFEEDBACK table. These conditions select objects at the tablespace or partition level when they have a REASON value of BASIC, KEYCARD, or CONFLICT in the SYSSTATFEEDBACK table.

Note: For more information about the SYSSTATFEEDBACK table and its column values, see the IBM documentation.

The following new conditions are now provided:

- (TS) RTS RUNSTATS SYSSTATFEEDBACK TS lvl (tablespace level)
- (TS) RTS RUNSTATS SYSSTATFEEDBACK pt lvl (partition level)

These conditions can execute RUNSTATS TABLE ALL USE PROFILE or INDEX ALL syntax to collect the statistics on the affected tablespaces and indexes. RUNSTATS also removes the applicable entries from the SYSSTATFEEDBACK table.

Note: For more information about these EOCs and ACs, see the online help. On the Extract Conditions or Action Conditions panel, type **H** (Help) next to the condition and press Enter. For more information about USE PROFILE, see the IBM documentation.

DB2 10 and DB2 11 IBM RUNSTATS USE PROFILE Support

DB2 10 and DB2 11 IBM RUNSTATS USE PROFILE syntax is now supported. USE PROFILE lets RUNSTATS use an existing statistics profile to gather statistics. The use of RUNSTATS profiles permits autonomic statistics maintenance. These profiles also let you quickly invoke the RUNSTATS utility with a predefined set of options at the individual table level.

CA Database Analyzer supports the USE PROFILE function through the new %PRFRWRS symbolic (PRoFile RoW RunStats Auto). This symbolic can be used in the existing model MJUTLRS and in the new model MJUTLRSS. This symbolic represents whether the USE PROFILE syntax can be used. On DB2 11, USE PROFILE can be used regardless of whether a RUNSTATS profile row exists. On DB2 10, USE PROFILE can be used only when the row exists. DB2 10 users can use the provided INSERT SQL sample to create SYSTABLES_PROFILES table level rows with default RUNSTATS options or to use the RUNSTATS SET PROFILE syntax.

Note: For more information about USE PROFILE, see the IBM documentation. For more information about the %PRFRWRS symbolic and the MJUTLRS and MJUTLRSS models, see the online help panels and the *CA Database Analyzer Reference Guide*.

DB2 10 Temporal Table Support

Additional DB2 10 temporal table support is now provided. CA Database Analyzer can now generate JCL to process the SYSTEM_TIME temporal tables (both the base table and the history table) in a tablespace set.

This functionality is provided through the following new JCL models. These models are accessed through the corresponding IC, QU, and RC utility codes on the Build Action Procedure panel:

MJUTLICS

Contains default JCL and DB2 utility statements that execute the IBM DB2 COPY utility against a tablespace set. This model uses the DB2 LISTDEF and TEMPLATE statements and the %STTRLATD variable.

MJUTLQUS

Contains default DB2 utility statements that execute the IBM QUIESCE (QU) utility against a tablespace set. This model uses the DB2 TABLESPACESET keyword.

MJUTLRCS

Contains default DB2 utility statements that execute the IBM RECOVER (RC) utility against a tablespace set. This model uses the DB2 LISTDEF statement and the %STTRLATD variable.

Note: These models also support COPY YES indexes, XML objects, and LOB objects as part of the tablespace set. For more information about these models, see the comments within the models, which are found at *hlq.CDBAMDL*. For more information about the utility codes, see the online help panels and the *CA Database Analyzer Reference Guide*.

Automatic Space Allocations with Models MJUTLAL and MJUTLAL1

The MJUTLAL and MJUTLAL1 models have been updated to let DB2 calculate the optimal primary and secondary data set quantities for you. These models now conditionally set PRIQTY -1, and always set SECQTY -1. This SECQTY value causes DB2 to use its sliding scale algorithm to calculate the secondary extents. Models MJUTLAL and MJUTLAL1 generate DB2 ALTER statements that resize a tablespace or index. Letting DB2 manage the quantities during an ALTER increases the chances of the object reaching its maximum data set size before reaching maximum extents.

Notes:

- For more information about how to use the MJUTLAL and MJUTLAL1 models, see the comments within the models.
- You can edit the models to specify the primary and secondary quantities yourself. For more information about the automatic symbolics that are available for this purpose, see the online help. Select the tutorial from the CA Database Analyzer main menu, and then select Symbolic Variable Descriptions. Place your cursor on a symbolic and press Enter to view its description.

New Action Conditions to Skip Migrated Objects

After the version 16.0 GA release of the CA Database Management Solutions for DB2 for z/OS, two new action conditions were added to identify migrated DB2 objects. You can use these conditions to skip the migrated objects when generating action JCL. Skipping the objects means that their underlying data sets are not recalled, which can save time and resources.

Note: Use these action conditions when CA Database Analyzer statistics are not being collected (that is, when relying on IBM Real-Time Statistics, or RTS). When CA Database Analyzer statistics are collected, migrated objects are skipped automatically.

The following action conditions are now available:

(TS) REJECT migrated DB2 objects

Identifies and skips migrated tablespaces during action condition processing.

(IX) REJECT migrated DB2 objects

Identifies and skips migrated indexes during action condition processing.

Note: For more information about these action conditions, see the online help. On the Action Conditions panel, type **H** (Help) next to the action condition and press Enter.

New Columns in Statistics Tables

Several columns have been added to the CA Database Analyzer statistics tables. These columns record various DB2 object attributes. For example, the RAIX_HASH column indicates whether an index is a hash overflow index. You can add attribute columns to the SQL predicates in action conditions and to extended queries (EQFs) in online query and report procedures. Adding these columns lets you include or exclude certain object types during processing.

The following columns have been added to the RATS_STATS_XXXX table:

- RATS_INSTANCE
- RATS_IPREFIX
- RATS_MAXPARTITIONS
- RATS_ORGANIZATIONTYPE

The following columns have been added to the RAIX_STATS_XXXX table:

- RAIX_EXTENSION_TYPE
- RAIX_HASH
- RAIX_INSTANCE
- RAIX_IPREFIX
- RAIX_SPARSE

An RAXX_UPDATE_DB2CAT column has also been added to every statistics table. This column is for internal use only.

Note: For more information about these columns, see the *CA Database Analyzer Reference Guide*.

Performance Improvements

CA Database Analyzer has been enhanced to include SYSIBM.SYSTABLES DBNAME in its WHERE clauses when using Real Time Object Selection (RTOS). Including DBNAME lets the DB2 optimizer use the DB2 10 index SYSIBM.DSNDTX05 that is on DBNAME and TSNAME.

Note: For a similar performance improvement in DB2 9 systems, add the index that is defined in *hlq.CDBASRC(CATDTX08)*.

This enhancement has the potential to lower your CPU and elapsed times and perform object maintenance more quickly. These improvements can be experienced when using tablespace extract object conditions (EOCs) after a tablespace is selected for processing.

Note: This support was added after the initial GA announcement of the previous release.

REBIND PACKAGE Support for Stored Procedures

REBIND PACKAGE support is now provided for native SQL stored procedures. CA Database Analyzer can generate the action JCL to rebind packages for native SQL stored procedures, which are listed in the SYSIBM.SYSPACKAGE table with TYPE=N. Rebinding these packages is useful for picking up new access paths. In previous releases, these packages were rebound manually. Rebinding them through CA Database Analyzer simplifies system maintenance.

This functionality is provided through the new MJUTLRPN JCL model. This model is accessed through the existing US utility code on the Build Action Procedure panel.

Note: For more information about the utility codes and models, see the online help panels and the *CA Database Analyzer Reference Guide*.

CA Detector Enhancements

The following enhancements have been made to CA Detector with this release:

- DB2 11 NFM and CM support is now provided.
- You can now view the amount of space that each collection interval uses.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

View the Collection Interval Size

You can now view the amount of space that each collection interval uses. This information helps prevent storage issues by showing how much space an interval uses, especially when collection options are changed.

This functionality is provided in the new SIZE (MB) column on the Datastore Interval Display.

CA Endeavor SCM Interface for DB2 Enhancements

The following enhancements have been made to CA Endeavor SCM Interface for DB2 with this release:

- DB2 11 NFM and CM support is now provided.
- A new ISPF panel interface is now provided.
- Associations support is now provided.
- BACK BIND support is now provided.
- Batch utility footprint support is now provided.
- Element action backout support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

New ISPF Panel Interface Support

The ISPF panel interface has been updated to provide a new main menu and many new sub panels to support new features. The new interface is much more intuitive and easy to use.

Associations Support

CA Endeavor SCM Interface for DB2 now lets you view, define, update, and delete associations in a table (EDB2EEAT) that the backout functionality in CA Endeavor SCM uses. Earlier releases only supported the package backout exit from CA Endeavor SCM.

The association table relates certain objects with particular CA Endeavor SCM inventory locations. This table can be used for services that need specific information for a system. The table owner can reference the association and can permit other users to manage the association. By default, the associations are restricted to use by only the owner only.

An association can have multiple attribute sets as follows:

- DB2 (a collection of DB2 subsystem IDs)
- JCL (used to submit jobs on a particular system)
- HOST (system identifier for submitted jobs)

This new functionality can be accessed from ISPF and in batch using the new DEFINE and DELETE control statements. The new BACK BIND control statement generates DB2 BIND statements for all CA Endeavor SCM locations automatically.

BACK BIND Support

Back bind is a process where an element is bound to all DB2 systems where it exists as it moves through the CA Endeavor SCM environment map. This process keeps the DB2 systems up to date with the most current version of a program.

A new control statement, BACK BIND, has been added to support this functionality.

Note: For more information about this processing, see the *CA Endeavor SCM Interface for DB2 User Guide*.

Batch Utility Footprint Support

The footprint function has been added to the batch utility. You can now use the batch utility to create and delete data for packages and plans in the footprint tables. Multiple BIND statement read is supported.

This function replaces the existing footprint function (ENCRFTP).

Element Action Backout Support

Element Action Backout support is now available. The existing package backout exit program now supports the element backout.

This functionality lets you:

- Backout individual elements from a package and display the generated statements.
- Backout a complete package.

CA Fast Check Enhancements

The following enhancements have been made to CA Fast Check with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 compress on INSERT is now supported.
- DB2 10 temporal table support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Fast Check can now process objects that have extended RBA and LRSN values.

DB2 10 Compress on INSERT Support

DB2 10 compress on INSERT is now supported. CA Fast Check can now expand and format compressed data rows using the compression dictionary that is built during INSERT processing.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Fast Check now supports all three temporal table types. You can check the integrity of the data in those tables and their indexes, as you can with nontemporal tables.

CA Fast Index Enhancements

The following enhancements have been made to CA Fast Index with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 EXCLUDE NULL KEYS support is now provided.
- DB2 10 compress on INSERT is now supported.
- DB2 10 temporal table support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 EXCLUDE NULL KEYS Support

DB2 11 EXCLUDE NULL KEYS support is now provided. When an index is defined with EXCLUDE NULL KEYS, DB2 creates an index entry only when a key column contains a value other than NULL. If every key column contains the NULL value, no index entry is created. This option reduces the size of the index, which can improve the performance of your index scans.

CA Fast Index can now process indexes that are defined with EXCLUDE NULL KEYS.

DB2 10 Compress on INSERT Support

DB2 10 compress on INSERT is now supported. CA Fast Index can now expand and format compressed data rows using the compression dictionary that is built during INSERT processing.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined SYSTEM_TIME period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined BUSINESS_TIME period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Fast Index now supports all three temporal table types. You can create and drop indexes for those tables, as you can with nontemporal tables.

CA Fast Load Enhancements

The following enhancements have been made to CA Fast Load with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 EXCLUDE NULL KEYS support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 inline large object (LOB) support is now provided.
- DB2 10 MEMBER CLUSTER support in universal tablespaces is now provided.
- DB2 10 temporal table support is now provided.
- LOB columns up to 2 GB are now supported.
- New partition independence keywords are now provided.
- zIIP support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 EXCLUDE NULL KEYS Support

DB2 11 EXCLUDE NULL KEYS support is now provided. When an index is defined with EXCLUDE NULL KEYS, DB2 creates an index entry only when a key column contains a value other than NULL. If every key column contains the NULL value, no index entry is created. This option reduces the size of the index, which can improve the performance of your index scans.

CA Fast Load can now process indexes that are defined with EXCLUDE NULL KEYS.

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Fast Load can now load tables that have extended RBA and LRSN values.

DB2 10 Inline LOB Support

DB2 10 inline large object (LOB) support is now provided. DB2 10 allows adding an inline LOB column to a base table where part of the LOB data exists in the base table LOB column. Placing part of the LOB data in the base table can improve performance for smaller LOBs by avoiding the need to access the LOB tablespace.

CA Fast Load can now process tables with inline LOB columns.

Note: For more information about using CA Fast Load to process LOB data, see the *CA Fast Load User Guide*.

DB2 10 MEMBER CLUSTER Support in Universal Tablespaces

DB2 10 MEMBER CLUSTER support in universal tablespaces (UTS) is now provided. When the MEMBER CLUSTER option is used, INSERTs do not cluster the data by the implicit or explicit clustering indexes. Instead, the data is placed in the tablespace according to available space. This option can reduce contention during heavy INSERT activity in a data sharing environment.

CA Fast Load can now process tables in a UTS that are set to use the MEMBER CLUSTER option.

Note: For more information about MEMBER CLUSTER support, see the *CA Fast Load User Guide*.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined SYSTEM_TIME period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined BUSINESS_TIME period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.

- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Fast Load now supports all three temporal table types. You can load data into these tables, as you can with nontemporal tables.

LOB 2 GB Column Support

CA Fast Load now supports LOB (CLOB, BLOB, and DBCLOB) columns up to 2 GB. This support lets you process larger LOB objects than previous releases, which were limited to the maximum region availability below the line. This support uses storage that is above the line, which permits CA Fast Load to process the larger LOB sizes. This support also permits CA Fast Load to adjust the buffer sizes for the input data sets automatically. You no longer need to use the SPANNED-RECSIZE keyword to adjust these buffer sizes in your syntax.

Note: You do not need to remove this keyword from your existing jobs. The utility ignores any SPANNED-RECSIZE value that is specified.

New Partition Independence Keywords

New keywords for greater control of partition independence (PI) processing are now provided. PI processing consumes DB2 resources like locks and buffers. Utilities can hold uncommitted updates for an extended period when multiple PI tasks are running concurrently, causing problems with resource availability and contention. The new keywords help you control the resource usage for better performance and efficiency.

The following keywords are now available in the *hlq.CDBAPARM(UTIL)* parmlib member. You can also specify them in CA Fast Load and CA Rapid Reorg SYSIN syntax:

PI-KEY-COUNT

Specifies the maximum number of keys to delete or insert before a checkpoint is made.

PI-MAXTASKS

Specifies the maximum number of concurrent partition independence (PI) tasks in a DB2 data sharing group that can access the same DB2 buffer pool. If data sharing is not used, this keyword specifies the maximum number of tasks in a DB2 instance.

PI-RETRY-COUNT

Specifies the number of times to retry deleting or inserting index keys before stopping the task.

Note: For more information about these keywords, see the parmlib help panels or the product user guides.

zIIP Support

System z Integrated Information Processor (zIIP) support is now provided. Offloading zIIP-eligible workloads to zIIP processors can reduce the processing load on your CPUs. zIIP processing can also decrease overall processing time by implementing tasks in parallel.

CA Fast Load can now use zIIP processors when updating indexes. No syntax changes are required.

CA Fast Recover Enhancements

The following enhancements have been made to CA Fast Recover with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 Extended Address Volume (EAV) support is now provided.
- DB2 10 FlashCopy support is now provided.
- DB2 10 temporal table support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

The extended RBA and LRSN conversion is enabled through the new RBALRSN_CONVERSION keyword and the new Set RBALRSN_CONVERSION format for REBUILD INDEX (RBALRSN_CONVERSION) parameter in the *hlq.CDBAPARM*(PFR) member.

CA Fast Recover now records 10-bytes RBA and LRSN values in PFR messages.

DB2 10 EAV Support

DB2 10 Extended Address Volume (EAV) support is now provided. EAV support provides more than 65,520 cylinders, which lets you store more DB2 data on a single volume.

DB2 10 FlashCopy Support

DB2 10 FlashCopy support is now provided.

The FlashCopy feature creates a point-in-time copy that is written to a VSAM data set. Starting with DB2 10, this image copy is also recorded in SYSIBM.SYSCOPY. The image copy is immediately available for read and write access. This copy method greatly reduces the amount of time that tablespaces are unavailable.

CA Fast Recover can now recover data from FlashCopy image copies, including the copies that are recorded in SYSCOPY. Previously, CA Fast Recover was unable to use these image copies as input unless they were first converted to sequential image copies. The ability to process them directly eliminates the conversion step, simplifying the process and reducing the chance of user error.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.

- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Fast Recover now supports all three table types using the new VERIFYSET keyword. The VERIFYSET keyword specifies whether to verify that all related objects are included in the RECOVER statement when performing a point-in-time recovery. For a point-in-time recovery, temporal and history tables must be recovered in the same statement.

Note: For more information about VERIFYSET and temporal table support, see the *CA Fast Recover User Guide*.

CA Fast Unload Enhancements

The following enhancements have been made to CA Fast Unload with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 FlashCopy support is now provided.
- DB2 10 inline large object (LOB) support is now provided.
- DB2 10 temporal table support is now provided.
- Dynamic allocation support is now provided.
- Native LOB support is now provided.
- You can now unload objects in descending size order.
- A new PART-SEPARATE keyword is now provided.
- A new XMLDB2V9 parameter has been added to OUTPUT-FORMAT.
- zIIP support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Fast Unload can now process objects that have extended RBA and LRSN values.

DB2 10 FlashCopy Support

DB2 10 FlashCopy support is now provided. The FlashCopy feature creates a point-in-time copy that is written to a VSAM data set. Starting with DB2 10, this image copy is also recorded in SYSIBM.SYSCOPY. The image copy is immediately available for read and write access. This copy method greatly reduces the amount of time that tablespaces are unavailable.

CA Fast Unload can now unload data directly from FlashCopy image copies, including the copies that are recorded in SYSCOPY. Previously, CA Fast Unload was unable to use these image copies as input unless they were first converted to sequential image copies. The ability to process them directly eliminates the conversion step, simplifying the process and reducing the chance of user error.

This functionality is enabled through the new FLASHCOPY parameter in the INPUT-FORMAT keyword. A new INDDN-COUNT keyword also lets you specify the number of FlashCopy data sets to process.

Notes:

- For more information about these keywords, see the *CA Fast Unload User Guide*.
- CA Fast Unload also supports the FlashCopies that CA Quick Copy generates.

DB2 10 Inline LOB Support

DB2 10 inline large object (LOB) support is now provided. DB2 10 allows adding an inline LOB column to a base table where part of the LOB data exists in the base table LOB column. Placing part of the LOB data in the base table can improve performance for smaller LOBs by avoiding the need to access the LOB tablespace.

CA Fast Unload can now process tables with inline LOB columns.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Fast Unload now supports all three temporal table types. You can unload the data from those tables, as you can with nontemporal tables.

Note: For more information about processing temporal tables, see the *CA Fast Unload User Guide*.

Dynamic Allocation Support

Dynamic allocation support is now provided. You can now allocate the `SYSREC`, `SYSDDL`, and `SYSCTL` data sets dynamically instead of manually. These data sets contain the unloaded rows, `CREATE TABLE` DDL statements, and load control statements that CA Fast Unload generates.

This method eliminates the need to code the output data set requirements in the JCL manually. Allocating these output data sets dynamically lets CA Fast Unload calculate the required amount of space for the data sets. Dynamic allocation also supports the use of symbolics, which let you generate unique names for the data sets.

Dynamic allocation is implemented by using a `SYSTEMPL DD` statement in the JCL or a `TEMPLATE` control statement within the `SYSIN` syntax.

Note: For more information about dynamic allocation and the `SYSTEMPL` and `TEMPLATE` statements, see the *CA Fast Unload User Guide*.

Native LOB Support

Native large object (LOB) support is now provided. In previous releases, CA Fast Unload provided LOB support only with SQL-ACCESS ONLY, which uses DB2 to process the SQL statements. You can now process LOBs with SQL-ACCESS NONE and EXTENSION. These SQL-ACCESS options use CA Fast Unload to process SQL statements natively.

Processing LOBs natively can decrease processing time, because CA Fast Unload can multitask and unload multiple tables concurrently. Native processing can also improve the overall performance of a DB2 subsystem by decreasing DB2 overhead, which reduces the subsystem workload.

Note: For more information about LOB support, see the *CA Fast Unload User Guide*.

Unload Objects in Descending Size Order

You can now unload objects in descending order by size. This enhancement upgrades the internal logic that CA Fast Unload uses to ATTACH unload subtasks. Unloading the objects in descending order can potentially reduce the overall elapsed time.

This functionality is enabled through the new OBJ-ORDER keyword.

Note: For more information about OBJ-ORDER, see the *CA Fast Unload User Guide*.

New PART-SEPARATE Keyword

A new PART-SEPARATE keyword is now provided. This keyword specifies whether to separate each partition into its own subtask so that the partitions can be processed in parallel. Parallel processing can decrease processing time.

This keyword can be specified as a global option, where it applies to all partitions specified in all SELECT statements. This keyword can also be specified as a SELECT option, where it applies only to the partitions specified within that SELECT statement.

PART-SEPARATE {NO|YES}

NO

Does not separate the partitions into individual subtasks. This value is the default.

YES

Separates each partition into its own subtask and processes the partitions in parallel.

Note: For more information about this keyword, see the *CA Fast Unload User Guide*.

New XMLDB2V9 Parameter for OUTPUT-FORMAT Keyword

A new XMLDB2V9 parameter is now provided for the OUTPUT-FORMAT keyword. This parameter is used when unloading XML data for use by the DB2 9 version of the IBM LOAD utility. DB2 9 requires a RECFM=VB (variable-block) data set. The DB2 9 IBM LOAD utility also requires that the XML data be stored in a variable-length field with a 2-byte length component. The new XMLDB2V9 parameter lets you generate output format with this criteria so that IBM LOAD can process the output.

Note: This parameter is not required in DB2 10, which supports RECFM=VBS (variable-blocked-spanned) data sets.

This new CA Fast Unload parameter can now be specified using the following syntax:

```
OUTPUT-FORMAT {LOAD|FIXED|DSNTIAUL|COMMA-DELIMITED|  
VARIABLE,[DB2LOAD|FASTLOAD|XMLDB2V9]|EXTERNAL}
```

XMLDB2V9

Unloads the data in a format that the DB2 9 IBM LOAD utility can use.

Note: For more information about this parameter, see the *CA Fast Unload User Guide*.

zIIP Support

System z Integrated Information Processor (zIIP) support is now provided. Offloading zIIP-eligible workloads to zIIP processors can reduce the processing load on your CPUs. zIIP processing can also decrease overall processing time by implementing tasks in parallel.

CA Fast Unload can now use zIIP processors when processing LOBs (including inline LOBs) natively. No syntax changes are required.

CA Insight DPM Enhancements

The following enhancements have been made to CA Insight DPM with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 11 IFCID 378 and 379 support is now provided.
- DB2 10 IFCID 316 and 401 support is now provided.
- DB2 10 IFCID 363 support is now provided.
- DB2 10 IFCID 365 support is now provided.

- DB2 10 statement ID-based retrieval of dynamic and static SQL statement statistics, SQL text, or both is now provided.
- Cumulative data sharing statistics are now provided to CA Cross-Enterprise APM.
- Expanded sort metrics support is now provided.
- An extended print format for interval timestamps is now provided.
- IBM DB2 Accelerator support is now provided.
- Current lock contention information for each lock holder and waiter is now provided on a single line.
- All database I/O type trace records are now processed in a single record.
- Improved package information is now provided on the current and history package list panels.
- Improved System Condition Monitor (SCM) functionality is now provided.
- A new IQL record, (1044) APM-APPL-SUM, is now provided.
- New SORT exit routines for sorting DB2 trace records from GTF or SMF are now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Insight DPM now provides extended RBA and LRSN support in all online and batch reports that display RBA values.

DB2 11 IFCID 378 AND 379 Support

DB2 11 IFCID 378 and 379 support is now provided. CA Insight DPM now reports on query entry and exit to the DB2 Accelerator. This data is provided online and in batch SQL and record trace reports. This functionality improves your ability to monitor data processing performance on the accelerator.

Note: For more information about IFCID 378 and 379, see the *CA Insight DPM Writing Requests Reference Guide*.

More information:

[IBM DB2 Accelerator Support](#) (see page 48)

DB2 10 IFCID 316 and 401 Support

DB2 10 IFCID 316 and 401 support is now provided. CA Insight DPM now lets you monitor static and dynamic SQL execution statistics in real-time, online history, and in batch. CA Insight DPM provides this functionality by using the typically lower-overhead, lower-cost SQL statistics data that IFCID 316 and IFCID 401 collect.

This enhancement helps eliminate the need to run expensive SQL traces. Now you can generate reports that show how SQL statistics have changed over time. You can also obtain a current or historical snapshot of SQL statistics and identify the top resource users.

CA Insight DPM provides the following information for real-time analysis:

- A snapshot of dynamic SQL activity for statements in the dynamic statement cache
- A snapshot of static SQL in the EDM pool

Use the Focus command to limit displayed data to selected programs, SQL statements, or SQL statements that exceed a specified threshold value. You can also view the data for the current history collection interval.

For historical reporting, CA Insight DPM collects data from synchronous and asynchronous IFCID 316 and IFCID 401 records, and builds its own internal records. These records reflect dynamic and static SQL activity for a collection interval. View the data in an interval by interval display, or view a summary of all selected intervals. In addition, you can optionally collect the SQL statement text (IFCID 317 records) for the dynamic SQL statements with the worst performance. Use history selection criteria to limit the data to selected programs, SQL statements, or SQL statements that exceed a specified threshold value.

In addition, batch reports are provided for reporting on the new, internal, interval-based trace records that are unloaded from the online history files.

The following list describes the new #IFI variables, SYSPARMS, CA Insight DPM main menu options, and batch reports that enable this functionality:

#IF VARIABLES

IQL-INCLUDE-IFCID318

Controls whether to start the DYNSTATS IQL request, which starts IFCID 318. When IFCID 318 is enabled, DB2 collects statistics for dynamic SQL and it stores the statistics in the dynamic statement cache.

IQL-INCLUDE-IFCID400

Controls whether to start the STASTATS IQL request, which enables IFCID 400. When IFCID 400 is enabled, DB2 collects statistics for static SQL and it stores the statistics in the EDM pool.

SYSPARMS

HIST-IFC317-TOT-ELAP-TOP=0

Captures (317) DYNAMIC-SQL-STMT records for the top xx resource users by total elapsed time during the collection interval.

HIST-IFC317-AVG-ELAP-TOP=0

Captures (317) DYNAMIC-SQL-STMT records for the top xx resource users by average elapsed times per execution during the collection interval.

HIST-IFC317-TOT-CPU-TOP=0

Captures (317) DYNAMIC-SQL-STMT records for the top xx resource users by total CPU time during the collection interval.

HIST-IFC317-AVG-CPU-TOP=0

Captures (317) DYNAMIC-SQL-STMT records for the top xx resource users by average CPU time during the collection interval.

IQL-INCLUDE-IFCID318

Controls the setting of the IQL-INCLUDE-IFCID318 #IF IQL variable, which controls the starting of the DYNSTATS IQL request, enabling IFCID 318. If IFCID 318 is enabled, DB2 collects statistics for dynamic SQL in the dynamic statement cache.

IQL-INCLUDE-IFCID400

Controls the setting of the IQL-INCLUDE-IFCID400 #IF IQL variable, which controls the starting of the STASTATS IQL request, enabling IFCID 400. If IFCID 400 is enabled, DB2 collects statistics for static SQL in the EDM pool.

Main Menu Options

SQL Exec Current Stats

Contains a current snapshot of dynamic and static SQL statistics.

SQL Exec History

Contains dynamic and static SQL statistics for a specified interval.

Batch Reports

The trace report shows how the SQL statement execution changed for each interval. The short and long summary reports identify the most resource intensive and poorly performing SQL statements.

BTDSQLTX Dynamic Cache SQL Statement Text

Displays a complete trace report of all of the dynamic SQL statement text that was captured for each interval.

BTDSQSML Long Summary of Dynamic SQL Executed

Displays a complete breakdown summary of executed dynamic SQL, grouped by program and statement. This report shows the SQL statements with the worst performance, ordered by a summary of accumulated elapsed time for each SQL statements.

BTDSQSMS Short Summary of Dynamic SQL Executed

Displays a line-by-line summary of executed dynamic SQL, grouped by program and statement. This report is ordered according to a summary of accumulated elapsed time for each SQL statement.

BTDSQTRL Long Report of Dynamic SQL Executed

Displays a complete trace report of all of the dynamic SQL statements that were executed for each interval.

BTSSQSML Long Summary of Static SQL Executed

Displays a complete breakdown summary of executed static SQL, grouped by program and statement. This report shows the SQL statements with the worst performance, ordered by a summary of accumulated elapsed time for each SQL statements.

BTSSQSMS Short Summary of Static SQL Executed

Displays a line-by-line summary of executed static SQL, grouped by program and statement. This report is ordered according to a summary of accumulated elapsed time for each SQL statement.

BTSSQTRL Long Report of Static SQL Executed

Displays a complete trace report of all of the static SQL statements that were executed for each interval.

Note: For more information about how to collect these statistics, see the *CA Insight DPM* documentation.

DB2 10 IFCID 363 Support

DB2 10 IFCID 363 support is now provided. CA Insight DPM now reports on how DB2 uses the straw model distribution method when processing a parallel group. This information can help you understand how parallelism impacts DB2 performance. This record is written when performance trace class 8 is on. This record is used in the Batch Record Trace Report (BTRECTRC).

Note: For more information about this record, see the *CA Insight DPM Writing Requests Reference Guide*.

DB2 10 IFCID 365 Support

DB2 10 IFCID 365 support is now provided. CA Insight DPM can now report on which remote locations are connecting to the DB2 subsystem. CA Insight DPM can also report on the type and volume of activity that each remote location generates. For each remote location, CA Insight DPM provides remote location statistics for each interval or a summary of activity for all intervals.

Use this information to help assess your distributed activity and to understand how users access and use DB2 across your network. For example, you can identify and investigate excessive activity that is occurring in unexpected locations.

To obtain the DB2 activity by remote location, specify IFCID 365 in the HIST-STATS-RECS data collector SYSPARM parameter. Data is provided online, in list and detail formats, for real-time and historical activity. You can also obtain remote activity data through the following new batch reports: Distributed Activity by Remote Location (BTSTATR3 IQL request) and Summary of Remote Activity (BTSTASM3 IQL request).

DB2 10 Statement ID-Based Trace Record Support

DB2 10 statement ID-based retrieval of dynamic and static SQL statement statistics, SQL text, or both is now provided. CA Insight DPM provides this support for online IQL requests that support trace data containing statement IDs. Being able to retrieve SQL statement data based on the statement ID lets you more easily monitor and resolve problems with dynamic and static SQL. For example, you can quickly determine which SQL statements were involved in a timeout or deadlock.

Functionality is provided by using the T line command from the following panels:

- Detail Trace of DDF Activity (DISTRIB IQL request)
- SQL Statements that Scan Many Pages (EXPSQL IQL request)
- Timeouts and Deadlocks History (HISTREQ (HSDLKTO, HSUDLKTO) IQL requests)
- Thread Accounting Summary (HPRACCT (HPROBE) IQL request)
- SQL Statement Detail (HPRSTDTL (HPROBE) IQL request)
- Lock Timeouts and Deadlocks (LOCKCONT IQL request)
- Detail Trace of DB2 Activity (RECTRACE IQL request)
- Trace All SQL Statements (SQLTRACE IQL request)
- Probe Thread Summary (THRDTRAC (PROBE) IQL request)

For dynamic SQL statements that are available in the cache, this command displays performance data, including the SQL text, for the selected SQL statement. For static SQL statements, this command retrieves and displays the SQL statement text from the DB2 catalog.

Important! To avoid a tablespace scan of the SYSPACKSTMT table while retrieving static SQL, create a user index that contains a STMT_ID column.

Note: For more information about this functionality, see the *CA Insight DPM User Guide* or the online help.

Cumulative Data Sharing Statistics for CA Cross-Enterprise APM

CA Insight DPM can now provide cumulative data sharing statistics from all members of a data sharing group for consumption by CA Cross-Enterprise APM. This data facilitates more effective monitoring and management of DB2 subsystems.

This functionality is enabled by including the DATASHAREABLE keyword in a SUMMARIZE IQL request that reports on real-time (snapshot) data.

Note: For more information about this keyword, see the *CA Insight DPM Writing Requests Reference Guide*.

Expanded Sort Metrics Support

Expanded sort metrics support is now provided. CA Insight DPM now provides the total, average, and peak (high-water mark) values for the following types of real-time or historical sort activity:

- Number of sorts, sort elapsed time, and number of sorted records
- Sort space and sort record, data area, and key length sizes
- Work files used, initial work files, and partitioned work files
- Sort merge passes, columns, and keys
- Sorts with return code 0 and sorts with non-zero return codes

You can also identify the SQL statements with the highest sort values.

These metrics help determine sort space requirements, monitor sort use, and tune queries.

An expanded (151) IDB2ACCT record and new sort detail displays enable this functionality.

Note: For more information about how to collect and display sort metrics, see the *CA Insight DPM User Guide*.

Extended Print Format for Intervals

CA Insight DPM now supports extended print formats for interval times.

You can now specify an hourly interval print format up to eight digits when specifying a date-time literal mask in your field-attribute syntax. For example:

HHHHHHH:MM:SS.TTTTTT

This enhancement lets you format your interval fields, such as (EVENT-DUR), with greater precision.

Note: For more information about field-attribute syntax, see the *CA Insight DPM Writing Requests Reference Guide*.

IBM DB2 Accelerator Support

CA Insight DPM now collects and reports the statistics that the IBM DB2 Accelerator generates. The IBM DB2 Accelerator quickly filters and processes enormous amounts of data. Using CA Insight DPM to collect and report on IBM DB2 Accelerator-related trace data helps monitor and manage your IBM DB2 Accelerator use more effectively. For example, you can compare the IBM DB2 Accelerator SQL execution time against the locally executed SQL and can tune your IBM DB2 Accelerator use.

CA Insight DPM provides the following IBM DB2 Accelerator-related information:

- Current IBM DB2 Accelerator status
- Resource use
- Failures or delays
- Number of queries using the IBM DB2 Accelerator and average execution time per query
- Volume of data that is transferred between DB2 and the IBM DB2 Accelerator

The new Accelerator Services displays, provided at the subsystem and thread history levels, enable this functionality. In addition, accelerator-related statistics are now provided in the following CA Insight DPM batch reports:

- Detailed Summary of Accounting Information
- Detailed Trace of Accounting Information
- Summary of DB2 Database Address Space Statistics
- Statistics Data Trace
- System Parameters

Note: For more information about using these displays and reports, see the *CA Insight DPM* documentation.

More information:

[DB2 11 IFCID 378 AND 379 Support](#) (see page 42)

Improved Lock Contention Display

CA Insight DPM now displays lock contention information for each lock holder and waiter on a single line. Previously this information was provided on multiple lines, so scrolling was required to view multiple contentions. The updated display makes it easier to monitor multiple contentions simultaneously.

This functionality is provided on the Current Lock Contentions panel.

Note: For more information about accessing the Current Lock Contentions panel, see the *CA Insight DPM User Guide*.

Improved I/O Trace Record Processing

CA Insight DPM now processes all database I/O type trace records in a single record. This enhancement helps eliminate the problems that occurred when combining different I/O type records in the same IQL request.

A new IQL record, (1012) IO, provides this functionality. This record consolidates the information from the READ-IO and WRITE-IO trace records and is included in the IOBYPSET and APPLIO IQL requests.

Note: For more information about (1012) IO, see the *CA Insight DPM Writing Requests Reference Guide*.

Improved Package Information on the Current and History Package List Panels

CA Insight DPM now provides the following package information, in tabular format, on the Current Package/DBRM List, Package/DBRM History List, and Package/DBRM History Summary List panels:

- Number of SQL statements in the package or DBRM
- CPU and zIIP CPU time
- I/O suspension time
- Lock, latch, claim, and drain suspension time
- Suspension time unrelated to I/O or lock/latch

In addition, this data is now sortable.

Use this information to evaluate packages effectively, without accessing another panel. For example, you can quickly assess which packages or DBRMs are consuming the most resources.

Notes:

- Detailed column descriptions are provided in the CA Insight DPM online help (F1).
- For more information about accessing these panels, see the *CA Insight DPM User Guide*.

Improved System Condition Monitor Functionality

The CA Insight DPM System Condition Monitor (SCM) now supports the following functionality:

- Displaying an active exception count for each DB2 subsystem. The exception count is provided through the new Excpt Count field in the System Condition Monitor. This count reflects all active exceptions that have the IN-ALERT status.
- Sorting the displayed data by column number. Using the primary command SORT eliminates the need to scroll through a long list of subsystems. For example, you can sort the data so the subsystems experiencing the most exceptions display at the top of the subsystem list in descending order.
- Invoking a list of exceptions (X), active threads (T), or subsystem statistics (A) for the selected subsystem.

Use this information to help monitor the health of multiple DB2 subsystems and quickly determine which DB2 subsystems are experiencing problems.

Note: For more information about the System Condition Monitor, see the *CA Insight DPM System Reference Guide* or *CA Insight DPM User Guide*.

New IQL Record (1044) APM-APPL-SUM

CA Insight DPM now supports a new IQL record, (1044) APM-APPL-SUM. This record provides a summary of application activity for an interval following the activation of the IFCID 1044 record in the data collector. The application activity that this record collects is provided to CA Cross-Enterprise APM.

This record collects only the data that CA Cross-Enterprise APM uses and is designed to provide improved application activity reporting to CA Cross-Enterprise APM users.

A new IQL request, DSQAPMAS, is provided to support this record.

Note: For more information about (1044) APM-APPL-SUM, see the *CA Insight DPM Writing Requests Reference Guide*.

New SORT Exit Routines

CA Insight DPM now supports two new SORT exit routines:

NSIGHT15

Processes each GTF or SMF record before the record is sorted.

NSIGHT35

Processes each GTF or SMF record after the record is sorted.

These routines facilitate the sorting of DB2 trace records based on the more precise time value that is provided in QWHSSTCK. As a result, concatenated GTF and SMF files are accurately sorted in ascending sequence. CA Insight DPM Batch Report Writer IQL requests that specify the INTERVAL parameter process successfully.

Sample JCL is provided in the SMFSORT member in the Insight common SOURCE data set.

Note: For more information about NSIGHT15 and NSIGHT35, see the *CA Insight DPM Batch Report Reference Guide*.

CA Log Analyzer Enhancements

The following enhancements have been made to CA Log Analyzer with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 creation of partition-by-growth tablespaces with NUMPARTS and MAXPARTITIONS is now supported.
- DB2 10 DEFINE NO support is now provided.
- DB2 10 temporal table support is now provided.
- DB2 10 TIMESTAMP support is now provided.
- DB2 10 unique index INCLUDE column support is now provided.
- A new DDL Activity Object report is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Log Analyzer has added extended logging support to DML activity reporting. This functionality is enabled through the ENABLE parmlib member.

Notes:

- For more information about extended logging support, see informational solution RI64114 on the CA Support web site.
- The log report forms have been updated to accommodate the longer RBA and LRSN values. Before you use extended log values, update any customized report forms to accommodate these values. A new Extended Values field is provided on the Report Form Definition panel to specify the use of longer values. For more information about supporting a subsystem after converting to extended logging, see the *CA Log Analyzer User Guide*.

DB2 10 Creation of Partition-by-Growth Tablespaces with NUMPARTS and MAXPARTITIONS

DB2 10 creation of partition-by-growth (PBG) tablespaces with NUMPARTS and MAXPARTITIONS is now supported. Specifying both NUMPARTS and MAXPARTITIONS in the same CREATE TABLESPACE statement lets you create a PBG tablespace with multiple partitions. This capability gives you greater control over your DASD allocations.

CA Log Analyzer can now generate CREATE TABLESPACE DDL statements with NUMPARTS and MAXPARTITIONS.

DB2 10 DEFINE NO Support for LOB and XML Tablespaces

DB2 10 DEFINE NO support for LOB and XML tablespaces is now provided. Defining a tablespace as DEFINE NO lets you defer the creation of this tablespace and its indexes. DB2 logs the tablespace in the DB2 catalog, but does not create it until a LOAD or INSERT is performed. Using DEFINE NO can help you improve space management.

CA Log Analyzer can now generate DDL statements with DEFINE NO syntax.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Log Analyzer now supports all three temporal table types. You can report on changes to those tables, and can undo or redo changes to those tables, as you can with nontemporal tables.

CA Log Analyzer also supports the system-maintained history tables that are associated with `SYSTEM_TIME` tables. When generating REDO SQL or Log Apply statements, you can specify whether to include or exclude activity against history tables. For example, you can exclude the history activity to allow data propagation to targets that are also defined with a `SYSTEM_TIME` period history table relationship. This functionality is enabled through the new Redo History Table Updates field on the DML Related Updates Options panel.

DB2 10 TIMESTAMP Support

DB2 10 `TIMESTAMP` support is now provided. CA Log Analyzer now supports greater `TIMESTAMP` precision in DML activity reports, data queries, load files, Log Apply processing, and DDL file mapping. When you specify `TIMESTAMP EXTERNAL (length)` in an `INTO TABLE` statement, you can use the length value to specify the number of microseconds.

CA Log Analyzer also supports `TIMESTAMP WITH TIME ZONE` in DML activity reports, data queries, load files, Log Apply processing, DDL file mapping, and DDL generation.

DB2 10 Unique Index INCLUDE Column Support

DB2 10 New Function Mode (NFM) includes a new INCLUDE clause on the CREATE INDEX and ALTER INDEX statements. The INCLUDE clause lets you define nonkey columns in a unique index. These nonkey columns let queries use the unique index for index-only access, so you can eliminate indexes that were created solely to enable index-only access. Eliminating unnecessary indexes can help improve your system performance, simplify index maintenance, and decrease physical storage requirements.

CA Log Analyzer now supports unique indexes with nonkey columns. When using an index to select key columns for a table, you can specify to use only unique index columns as key columns. This functionality is provided through the new U (Unique) option on the Index List panel.

Note: For more information about selecting key columns, see the *CA Log Analyzer User Guide*.

New DDL Activity Object Report

You can now generate DDL Activity Object reports. These reports show the objects that were created, altered, and dropped for each Unit of Recovery (UR) within a specified log range.

This report can help you perform auditing tasks quickly. For example, you can generate an object report that shows all objects that a specific user created. Using an object report to gather this information instead of a detail report or DDL statements reduces the amount of data that you must view.

This functionality is enabled through the new O (Object) option for the Level of Detail field on the DDL Activity Report Options panel.

Note: For more information about this new option, see the online help panels in the product. For more information about generating DDL activity reports, see the *CA Log Analyzer User Guide*.

CA Merge/Modify Enhancements

The following enhancements have been made to CA Merge/Modify with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 Extended Address Volume (EAV) support is now provided.
- DB2 10 FlashCopy support is now provided.
- A new option DELETE KEEP GDGLIMIT is provided for managing SYSCOPY image copy records.
- Reclaim version support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Merge/Modify now supports 10-byte RBA and LRSN values in the change accumulation recovery tables and the surge table. You can convert rows of change accumulation recovery tables from previous releases to the current release using the CONVERT_ALOGRANGE keyword. The initial conversion is performed during post-installation processing by using a DB2 catalog customization task.

Note: The change accumulation recovery tables and the stage table have been updated for Version 17.0. For more information about the changes to the objects, see the *Implementation Guide*. For more information about how to convert the tables using the DB2 catalog customization task, see the *Installation Guide*. For more information about the CONVERT_ALOGRANGE keyword, see the *CA Merge/Modify User Guide*.

DB2 10 EAV Support

DB2 10 Extended Address Volume (EAV) support is now provided. EAV support provides more than 65,520 cylinders, which lets you store more DB2 data on a single volume.

DB2 10 FlashCopy Support

DB2 10 FlashCopy support is now provided. The FlashCopy feature creates a point-in-time copy that is written to a VSAM data set. Starting with DB2 10, this image copy is also recorded in SYSIBM.SYSCOPY. The image copy is immediately available for read and write access. This copy method greatly reduces the amount of time that tablespaces are unavailable.

CA Merge/Modify can now process data directly from FlashCopy image copies, including the copies that are recorded in SYSCOPY. Previously, CA Merge/Modify was unable to use these image copies as input unless they were first converted to sequential image copies. The ability to process them directly eliminates the conversion step, simplifying the process and reducing the chance of user error.

DELETE Options for Image Copy Records

After the version 16.0 GA release of CA Database Management Solutions for DB2 for z/OS, a new option KEEP GDGLIMIT was added for the MODIFY RECOVERY DELETE keyword. Use this option to retain the SYSCOPY records of local primary (LP) full image copies based on the GDG limit of the most recent LP full image copy.

CA Merge/Modify Reclaim Version Support

CA Merge/Modify now supports reclaiming object version numbers during the MODIFY RECOVERY processing. Reclaiming version information prevents you from running out of available version numbers and helps you keep objects available for future changes (ALTER).

This functionality is enabled through the new Use reclaim version functionality (RECLAIM-VERSION) parameter in the PMM parmlib member (*hlq.CDBAPARM*). The functionality is disabled by default.

Note: This support was added following the Version 16.0 GA release of the CA Database Management Solutions. For more information about updating the version information, see the *CA Merge/Modify User Guide*. For more information about editing the PMM parmlib member, see the *Implementation Guide*.

CA NSM Database Option for DB2 for z/OS Enhancements

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

CA Plan Analyzer Enhancements

The following enhancements have been made to CA Plan Analyzer with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 10 index probing is now supported.
- DB2 10 non-inline SQL scalar function packages are now supported.
- DB2 10 row and column access control is now supported.
- DB2 10 statement-level optimization hints are now supported.
- DB2 10 XMLXSROBJECTID support is now provided.
- DB2 9 XMLCAST specification clauses are now supported.
- Conditionally indented SQL substatements are now supported.
- Compare Explain Versions and Compare Versions Cost Compare reports can now report on historical versions of explained SQL sources where the access path has changed and the SQL is unchanged.
- Improved cost difference filters for the CA Plan Analyzer Compare Versions Cost Compare Report are now supported.
- A new DB2 Profiles Services facility is now provided to facilitate the management of DB2 profiles.
- Explain table indexes can now be automatically created during a future explain.
- Installed explain tables are now updated during a future explain with the plan and package data for catalog explain sources.
- Physical Rule 1083 is now supported.
- A new QUERY Reporting option is now provided for requesting catalog based reports about data that is contained in the access path repository.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 10 Index Probing Support

DB2 10 index probing is now supported. CA Plan Analyzer Future Explain now reports on statistics that the DB2 optimizer uses during index probing.

DB2 uses index probing to obtain more accurate filter factor estimates for matching predicates when conditions prevent the DB2 optimizer from obtaining an accurate estimate. Improved filter factor estimates can help stabilize the access path.

Note: For more information about how DB2 uses index probing, see the *IBM DB2 for z/OS Performance Topics* guide.

Index probing is supported through the following new or updated features:

- Enhanced explain Index Probe Statistics Report
- DSN_COLDIST_TABLE and DSN_KEYTGTDIST_TABLE explain tables
- IXPROBSTAT parmlib option
- REPORT enhanced explain card

A new enhanced explain report, Index Probe Statistics, shows how index probing influenced the access path that the DB2 optimizer selected during SQL execution. For example, you can determine which runtime statistics the DB2 optimizer applied to the query.

The Index Probe Statistics Report consists of two subreports. The first subreport focuses on column distribution statistics. This report is generated only when statistics are present in the DSN_COLDIST_TABLE explain table. The second subreport focuses on key target distribution statistics. This report is generated only when statistics are present in the DSN_KEYTGTDIST_TABLE explain table. Each report is generated on a per statement basis.

Two new explain tables, DSN_COLDIST_TABLE and DSN_KEYTGTDIST_TABLE facilitate the generation of this report. Create these tables using the Create Explain/Input Table(s) facility.

The following new PPA parmlib option supports the new Index Probe Statistics Report:

Index Probe Stats Report (IXPROBSTAT)

Indicates whether to produce the Index Probe Statistics Report during an explain. Valid values are Y or N. The default is N.

Note: For more information about editing product-specific parmlib members, see the *Implementation Guide*.

DB2 10 Non-Inline SQL Scalar Function Package Support

DB2 10 non-inline SQL scalar function packages are now supported. This package type lets you use the enhanced features of the CREATE FUNCTION statement.

This support is provided through the following reports and commands:

- Line commands DPLY, R, RO, ROA, F, FA, FO, and QE
- Primary commands REBINDPACK and FREEPACK
- Package List Report
- Package Detail Report
- Package Statement Report
- Package Object Dependency Report
- Bind Package Report
- Rebind Package Report
- Package User Authorizations Report
- Package Explain Report

Note: For more information about the CREATE FUNCTION statement, see the IBM *DB2 for z/OS Application Programming and SQL Guide*.

DB2 10 Row and Column Access Control Support

DB2 10 row and column access control is now supported.

This feature lets you manage access to a table at the row and column level by using row permissions and column masks. Controlling access in this way allows you to enforce stricter security policies that prevent users from accessing rows and columns that contain sensitive information.

The following CA Plan Analyzer functions now support row permissions and column masks:

Explain Services Catalog/Autobuild Object Explain Source

Column Mask and Row Permission are now supported in the catalog object and autobuild object explain source, available when using Explain Services in expert mode. This support expands your ability to assess the SQL statements included in the packages that row permissions and column masks depend on.

SQL Dependency Analysis Report

The SQL Dependency Analysis report now includes object support for column masks and row permissions and is accessible from the following product functions:

- Explain Services
- Object Reporting
- Statement Reporting

- Query Explain Database
- Identify Problem SQL
- Compare Explain Versions
- Source Database Maintenance

Object Reporting

Column Mask Dependency and Row Permission Dependency reports are now provided. Verify whether your SQL DML statements adhere to your row and column access security policies by using these reports.

DB2 10 Statement-Level Optimization Hint Support

DB2 10 statement-level optimization hints are now supported. You can now create and maintain these hints using the CA Plan Analyzer Optimization Hints facility. These hints help influence the access path that the DB2 optimizer selects for processing SQL statements. Creating and managing statement-level optimization hints can help control SQL statement costs.

The PLAN_TABLE Maintenance facility has been redesigned as the Optimization Hints facility to help automate the process of creating statement-level optimization hints. This redesigned facility continues to support PLAN_TABLE maintenance.

New commands and a new report are also provided to support the creation and maintenance of these hints:

BINDQRYOPT (BQO)

Generates the control cards to bind a query and inserts the statement-level hint into the access path repository.

FREEQRYOPT (FQO)

Frees queries from the access path repository. Query removal lets you reclaim disk space, capitalize on DB2 optimizer improvements, and remove inadvertently bound queries.

PTH

Provides detailed query information for the selected QUERYNO. You can then use line commands to generate reports to compare the access paths that DB2 uses to create the statement-level access path hint.

DSC

Generates the new Dynamic Statement Cache Report. This report provides detailed information about the SQL statements in your dynamic cache. Use this data to help determine which SQL statements to use when creating your statement-level optimization hints.

CA Plan Analyzer now lets you create the following tables to support the creation of statement-level optimization hints:

- DSN_USERQUERY_TABLE
- DSN_STATEMENT_CACHE_TABLE

The Create User PLAN_TABLE(s) facility has been redesigned as the Create Explain/Input Table(s) facility to support these tables.

Note: For more information about creating statement-level optimization hints and the tables to support them, see the *CA Plan Analyzer User Guide*.

DB2 10 XMLXSROBJECTID Support

DB2 10 XMLXSROBJECTID support is now provided. XMLXSROBJECTID is a scalar function that is used to indicate whether an XML document has been validated, and by whom. Validated XML documents are less prone to errors. You can now specify XMLXSROBJECTID as a search condition for CA Plan Analyzer Statement Reporting to expand the depth of your SQL statement reporting.

DB2 9 XMLCAST Specification Clause Support

DB2 9 XMLCAST specification clauses are now supported.

This clause supports casts between XML and non-XML data types. This flexibility is important when incorporating XML into your existing DB2 environment. You can specify XMLCAST specification clauses as search conditions for Statement Reporting and Explain options to expand the depth of your SQL statement reporting.

Conditionally Indent SQL Substatements

You can now control how your SQL statements display. SQL substatements can now be conditionally indented, making complex output easier to read.

This new functionality is enabled through the new PPINDENT parameter in the PPA parmlib member.

Note: For more information about editing product-specific parmlib members, see the *Implementation Guide*.

Identify Non-SQL Access Path Changes

Compare Explain Versions and Compare Versions Cost Compare reports can now report on historical versions of explained SQL sources where the access path has changed and the SQL is unchanged. When upgrading to a new DB2 release, this option can help you identify access path changes that occurred due to something other than SQL changes.

Enable this functionality using the following options:

- Specify **N** (SQL UNCHANGED, ACCESS CHANGED) in the PPCEVOR parameter in the PPA parmlib member.
- Specify **N** as the Change Level when setting your Compare Versions Report options.

The COMPOPTS enhanced explain and report batch cards have also been updated to support this functionality.

Note: For more information about editing product-specific parmlib members, see the *Implementation Guide*. For more information about setting the Compare Versions Report options, see the *CA Plan Analyzer User Guide*.

Improved Cost Difference Filters for the Compare Versions Cost Compare Report

Improved cost difference filters for the CA Plan Analyzer Compare Versions Cost Compare Report are now supported. The following cost difference filters are now filtered by a percentage or an actual value, or by a percentage and an actual value:

- Milliseconds
- Service Units
- Total Cost

The ability to filter cost differences on either a percentage or actual value, or both, makes it easier to review and assess the report output.

This functionality is provided through the PPA Compare Versions Cost Report Options panel.

The CVCCOPTS enhanced explain and report batch cards have also been updated to support this functionality.

Note: For more information about specifying cost report options, see the *CA Plan Analyzer User Guide*.

New DB2 Profile Services Facility

A new DB2 Profile Services facility is now provided. This facility simplifies DB2 profile management by facilitating the process of managing profiles and profile attributes in the DSN_PROFILE_TABLE and DSN_PROFILE_ATTRIBUTES tables. For example, this facility helps verify that only valid profiles and attributes are added to the profile tables.

This facility supports the following management tasks for profiles that are based on optimization parameter or modeling functions:

- Create, enable, disable, update, browse, delete, or obtain the status of a profile
- Start or stop profiles and determine whether profiles are running

Note: You cannot create or update profiles that are based on the MONITOR or ACCEL functions. However, you can delete, enable, start, stop, display, or view status reports for these profiles.

This functionality is accessed through the new DB2 Profile Services option that is provided in the CA Plan Analyzer Administrative Functions.

The following new panel-specific and line commands also support this functionality:

STARTPROF

Starts all enabled profiles and displays all DB2 command messages. The status reports for the enabled profiles follow these messages.

STOPPROF

Stops all active profiles.

DISPROF

Shows whether profiling is active or inactive.

S

Switches the profile status from enabled to disabled, or from disabled to enabled.

ST

Displays a status report for the selected profile.

CMA

Creates a MODELING profile using parameters that were captured from the modeled DB2 subsystem.

Note: This command requires access to Thread Termination\Dynamic DSNZPARM, which is a Value Pack component. The Value Pack components are provided at no additional charge with the CA Database Management Solutions for DB2 for z/OS.

If a profile is started, the Access Path report identifies the profile ID that was applied to the explain statement. This information is provided in all of the Access Path report formats.

Note: For more information about how to use this facility, see the *CA Plan Analyzer User Guide* and the online help.

Create Explain Table Indexes Automatically During a Future Explain

Explain table indexes can now be automatically created during a future explain.

Indexes can help retrieve table data more quickly. Creating the indexes automatically reduces the possibility that the index is set up incorrectly.

This new functionality is enabled through the new AUTOCRIX parameter in the PPA parmlib member.

Note: For more information about editing product-specific parmlib members, see the *Implementation Guide*.

Update all Explain Tables with Source SQL Catalog Data During a Future Explain

CA Plan Analyzer can now update each installed explain table with the plan and package data for catalog explain sources during a future explain. This enhancement expands your ability to query the explain tables outside of CA Plan Analyzer using the plan or package that is associated with the explained SQL. As a result, you can more easily access the catalog explain data that the future explain captures for every installed explain table that DB2 populated.

Functionality is enabled by specifying **A** (Commit All) in the PLAN_TABLE Option field when setting your explain options.

Note: For more information about explain tables, see the IBM DB2 for z/OS documentation.

New Physical Rule 1083

You can now use new Physical Rule 1083 to identify opportunities for index optimization, which can help reduce index maintenance and overhead. This rule is part of the default rule set and displays in the PPA Explain Physical Rules Report. This rule applies to systems running only on DB2 10 or higher.

Note: For more information about the rule-based Expert System application, see the *CA Plan Analyzer User Guide* or *CA SQL-Ease User Guide*.

New Reporting Services QUERY Reporting Option

CA Plan Analyzer now supports a new QUERY Reporting option.

Note: This option is applicable in DB2 10 NFM and up.

Use the QUERY Reporting option for requesting catalog based reports about data that is contained in the access path repository. The access path repository is a group of DB2 catalog tables that contain data about queries that are tied to statement-level optimization hints. Statement-level optimization hints help influence the access path that the DB2 optimizer selects for processing SQL statements.

Identify the status, validity, and parameters of your statement-level optimization hints using QUERY Reporting. Optionally free queries from within the QUERY reports. Free the queries from the access path repository to remove invalid hints, address performance degradation, or capitalize on DB2 optimizer improvements. You can also free queries when the contents of the DSN_USERQUERY_TABLE are bound inadvertently.

Important! DB2 allows a FREE QUERY against only a local location. Changing the location value from the local default causes the FREE QUERY to fail.

To access this option, select QR on the CA Plan Analyzer Main Menu.

QUERY Reporting supports the following new reports:

Query List Report

Lists detailed information for each query and indicates whether the statement-level optimization hints are runtime or access path hints.

Query Explain Report

Provides detailed DB2 performance and optimization information in short access path analysis format. This information focuses on the query statement that a statement-level access path hint uses.

Query Optimization Parameters Report

Lists the queries that are defined as statement-level runtime hints and provides details about the optimization parameters.

Query Statement Report

Displays the SQL statement text that is contained within the query.

The following new line commands can be used from within the QUERY reports:

L

Generates a Query List report.

E

Generates a Query Explain report.

OP

Generates a Query Optimization report.

ST

Generates a Query Statement report.

F

Frees an individual query.

Issue the new primary command, FREEQUERY, to free a list of queries. Use this command on any panel that contains a query list.

The following new batch report cards are generated:

QUERYID

Contains the unique identifier for a query.

USERFLTR

Groups a set of queries under a specified name.

The OPTION and REPORT batch report cards have also been updated to support query reporting.

The batch action command F has been modified to support queries. The batch action command FREEQUERY is now supported.

Note: For more information about QUERY reporting, see the *CA Plan Analyzer Reference Guide* or *CA Plan Analyzer User Guide*.

CA Quick Copy Enhancements

The following enhancements have been made to CA Quick Copy with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 Extended Address Volume (EAV) support is now provided.
- DB2 10 FlashCopy support is now provided.
- DB2 10 temporal table support is now provided.
- Support for automated switching to EXCP NO is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

DB2 10 EAV Support

DB2 10 Extended Address Volume (EAV) support is now provided. EAV support provides more than 65,520 cylinders, which lets you store more DB2 data on a single volume.

DB2 10 FlashCopy Support

DB2 10 FlashCopy support is now provided.

The FlashCopy feature creates a point-in-time copy that is written to a VSAM data set. Starting with DB2 10, this image copy is also recorded in SYSIBM.SYSCOPY. The image copy is immediately available for read and write access. This copy method greatly reduces the amount of time that tablespaces are unavailable.

The FlashCopy support was added to CA Quick Copy in Version 16.0. With this release, the FlashCopy support was enhanced to enable other utilities to use snapshot point-in-time image copies that CA Quick Copy created.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.

- User-maintained temporal tables have a defined BUSINESS_TIME period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Quick Copy now supports all three table types.

Support for Automated Switching to EXCP NO

Using the Execute Channel Program (EXCP) processing optimizes performance of CA Quick Copy. However, some objects are not supported by EXCP YES. To let you still benefit from EXCP YES processing, CA Quick Copy now automatically switches to EXCP NO for the objects that are not supported for EXCP processing. EXCP processing is still used for the supported objects.

Note: This enhancement was added in Version 17.0 as a post-GA PTF. For more information about EXCP processing, see the *CA Quick Copy User Guide*.

CA Rapid Reorg Enhancements

The following enhancements have been made to CA Fast Load with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 EXCLUDE NULL KEYS support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 MEMBER CLUSTER support in universal tablespaces is now provided.
- DB2 10 temporal table support is now provided.
- New partition independence keywords are now provided.
- zIIP support is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 EXCLUDE NULL KEYS Support

DB2 11 EXCLUDE NULL KEYS support is now provided. When an index is defined with EXCLUDE NULL KEYS, DB2 creates an index entry only when a key column contains a value other than NULL. If every key column contains the NULL value, no index entry is created. This option reduces the size of the index, which can improve the performance of your index scans.

CA Rapid Reorg can now process indexes that are defined with EXCLUDE NULL KEYS.

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Rapid Reorg can now reorganize objects that have extended RBA and LRSN values.

CA Rapid Reorg can also convert the RBA or LRSN formats for the selected objects during the reorganization. The log values can be converted from 6 bytes to 10 bytes, or from 10 bytes to 6 bytes. This functionality is provided through the new RBALRSN_CONVERSION keyword.

Note: To use this keyword, specify SWITCH-TO-IBM YES in *hlq.CDBAPARM(UTIL)*. For more information about this keyword, see the *CA Rapid Reorg User Guide*.

DB2 10 MEMBER CLUSTER Support in Universal Tablespaces

DB2 10 MEMBER CLUSTER support in universal tablespaces (UTS) is now provided. When the MEMBER CLUSTER option is used, INSERTs do not cluster the data by the implicit or explicit clustering indexes. Instead, the data is placed in the tablespace according to available space. This option can reduce contention during heavy INSERT activity in a data sharing environment.

CA Rapid Reorg can now process tables in a UTS that are set to use the MEMBER CLUSTER option.

Note: For more information about MEMBER CLUSTER support, see the *CA Rapid Reorg User Guide*.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Rapid Reorg now supports all three types of temporal tables. You can reorganize the data in those tables, as you can with nontemporal tables.

Note: For more information about temporal table support, see the *CA Rapid Reorg User Guide*.

New Partition Independence Keywords

New keywords for greater control of partition independence (PI) processing are now provided. PI processing consumes DB2 resources like locks and buffers. Utilities can hold uncommitted updates for an extended period when multiple PI tasks are running concurrently, causing problems with resource availability and contention. The new keywords help you control the resource usage for better performance and efficiency.

The following keywords are now available in the `hlq.CDBAPARM(UTIL)` parmlib member. You can also specify them in CA Fast Load and CA Rapid Reorg SYSIN syntax:

PI-KEY-COUNT

Specifies the maximum number of keys to delete or insert before a checkpoint is made.

PI-MAXTASKS

Specifies the maximum number of concurrent partition independence (PI) tasks in a DB2 data sharing group that can access the same DB2 buffer pool. If data sharing is not used, this keyword specifies the maximum number of tasks in a DB2 instance.

PI-RETRY-COUNT

Specifies the number of times to retry deleting or inserting index keys before stopping the task.

Note: For more information about these keywords, see the parmlib help panels or the product user guides.

zIIP Support

System z Integrated Information Processor (zIIP) support is now provided. Offloading zIIP-eligible workloads to zIIP processors can reduce the processing load on your CPUs. zIIP processing can also decrease overall processing time by implementing tasks in parallel.

CA Rapid Reorg can now use zIIP processors during the log apply phase and when performing online reorganizations (that is, when DATA-AVAILABLE CRITICAL is specified). No syntax changes are required.

CA RC/Compare Enhancements

The following enhancements have been made to CA RC/Compare in this release:

- DB2 11 NFM and CM support is now provided.
- DB2 10 temporal table support is now provided.
- DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) support is now provided.
- DB2 10 unique index INCLUDE column support is now provided.
- DB2 9 implicitly hidden columns support is now provided.
- DB2 9 ROW CHANGE TIMESTAMP column support is now provided.
- Extended DB2 9 clone support is now provided.
- Full support for DB2 9 native stored procedures (NSPs) with multiple versions is now provided.
- Compare rules enhancement to trigger attributes

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided.

The following new table types are now supported:

- Temporal
- History (for system-period or bi-temporal tables)

When altering a table, you can edit a temporal table in a compare strategy. This processing includes:

- Adding or changing a default attribute of a column defined AS ROW BEGIN, AS ROW END, or AS TRANSACTION START ID.
- Adding a PERIOD definition.

You can also use the Alternate Catalog Mapping (ACM) analysis option in compare strategies.

When creating, altering, or templating a table in a compare strategy, you can now add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint. To enable this support, a new option, B, is provided on the Unique Constraints Management panel.

When creating, altering, or templating a table, you can now add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint. To enable this support, a new option, B, is provided on the Unique Constraints Management panel.

When creating, altering, or templating an index, you can now specifically include or exclude the BUSINESS_TIME WITHOUT OVERLAPS clause in the index key. To enable this support, a new field, BTWOO, is provided on the Index Create, Alter, and Template panels.

Note: For more information about using temporal tables, see the *CA RC/Migrator User Guide*.

DB2 10 TIMESTAMP Support

DB2 10 TIMESTAMP support is now provided. CA RC/Compare can now process tables containing DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) data types.

DB2 10 Unique Index INCLUDE Column Support

DB2 10 New Function Mode (NFM) includes a new INCLUDE clause on the CREATE INDEX and ALTER INDEX statements. The INCLUDE clause lets you define nonkey columns in a unique index. These nonkey columns let queries use the unique index for index-only access, so you can eliminate indexes that were created solely to enable index-only access. Eliminating unnecessary indexes can help improve your system performance, simplify index maintenance, and decrease physical storage requirements.

Note: For more information about adding INCLUDE column support, see the *CA RC/Migrator User Guide*.

DB2 9 Implicitly Hidden Columns Support

DB2 9 implicitly hidden columns are now supported. Support is now provided for the IMPLICITLY HIDDEN clause in a table. The DB2 result set of a SELECT * statement includes the column only when it is explicitly referenced. This setting can be beneficial for sites that use SELECT *. Database administrators can add new columns to tables without having to change all the programs that use the modified tables.

DB2 9 ROW CHANGE TIMESTAMP Support

The DB2 9 ROW CHANGE TIMESTAMP column type is now supported. For each row change (insert or update), DB2 generates a timestamp value for the column that corresponds to the time of the insert or update. When GENERATED BY DEFAULT is specified, the timestamp column values can be specified manually. Knowing how often updates occur during specific time ranges helps Database Administrators plan for data replication and auditing tasks. This enhancement also helps to perform the following tasks:

- Ensure data integrity and limit the time that locks are held when optimistic locking is implemented.
- Reduce application development costs when DB2 handles the timestamp logic automatically.

DB2 9 Clone Support

DB2 9 cloned objects are now supported. Clones can provide maximum data availability. For example, you can replace tables quickly by applying changes to a clone table and then exchanging the clone table for the base table.

Note: For more information about clone support, see the *CA RC/Migrator User Guide*.

Extended DB2 9 Native SQL Stored Procedures Support

Full support for DB2 9 native stored procedures (NSPs) with multiple versions is now available.

Compare Rules Enhancement to Trigger Attributes

You can now use compare rules for the object creator and text within the action text. This enhancement lets you use CA RC/Compare on triggers between environments where the schema or object names are different.

CA RC/Extract Enhancements

The following enhancements have been made to CA RC/Extract in this release:

- DB2 11 NFM and CM support is now provided.
- DB2 10 temporal table support is now provided.
- DB2 10 TIMESTAMP support is now provided.
- DB2 10 unique index INCLUDE column support is now provided.
- Enhanced dynamic sort options are now provided.
- Hardware compression support is now provided.
- Isolation level customization support for SQL access is now provided.
- LIMIT and WHERE clause support during an extract is now provided.
- Named relationship sets support is now provided.
- Split extract unit count support is now provided.
- Alias confirmation has been removed.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA RC/Extract now supports `SYSTEM_TIME` and `BUSINESS_TIME` temporal tables. You can extract and load data in those tables and their indexes, as you can with nontemporal tables. You can also load to a normal table as well as a temporal table.

DB2 10 TIMESTAMP Support

DB2 10 TIMESTAMP support is now provided. CA RC/Extract can now extract and load tables containing DB2 10 TIMESTAMP WITH TIME ZONE and `TIMESTAMP(precision)` data types.

DB2 10 Unique Index INCLUDE Column Support

DB2 10 New Function Mode (NFM) includes a new `INCLUDE` clause on the `CREATE INDEX` and `ALTER INDEX` statements. The `INCLUDE` clause lets you define nonkey columns in a unique index. These nonkey columns let queries use the unique index for index-only access, so you can eliminate indexes that were created solely to enable index-only access. Eliminating unnecessary indexes can help improve your system performance, simplify index maintenance, and decrease physical storage requirements.

Enhanced Sort DYNALLOC Options

The enhanced sort options enable you to control the dynamic allocation of the sort data sets that are used to perform the sort. You can specify the device type and number of data sets to be used.

The sort values are used in extract processing to populate the `DYNALLOC(device_type,dataset_num)` sort option statement.

Hardware Compression Support

DFSMS hardware compression for the extract data set is now supported. CA RC/Extract supports batch SQL extraction with data class supporting compression.

Isolation Level Support for SQL Access

You can now specify the isolation level for SQL access to extract. This enhancement helps reduce the contention on DB2 resources.

LIMIT and WHERE Clause Support

You can now change the LIMIT and other filtering criteria during extract processing. You can also specify a WHERE clause during the extraction step. This enhancement lets you use the same SRCDEF to extract different sets of data by varying the row filters.

If you unload extracted data in CA RC/Extract using SQL, you can now use the Total limit option to specify the number of rows to extract and retrieve an intact set of RI data.

Named Relationship Sets

CA RC/Extract lets you extract and load tables with multiple relationships with other tables using named relationship sets in RI Manager. You can use RI Manager to name and save a set of referential integrity (RI) relationships for a table, including user-defined relationships. For example, in addition to the DB2 primary key, there can be several application managed relationships. You can define several user RI relationships for the same table. This functionality lets you create test data for a table using different foreign table relationships.

Split Extract Unit Count Support

If you unload extracted data in CA RC/Extract using CA Fast Unload, you can now add a unit count value to each split extract file and the extract object. This enhancement helps to protect against space allocation errors and allows a data set to span multiple volumes when more than one is needed. For example, if the extract object is large.

Alias Confirmation Removed

When you are creating a source definition and there is only one alias, the alias is now auto-selected. This enhancement simplifies the process of creating a source definition.

Previously, a selection list was displayed from which you could select the alias. When there is only one alias, this selection list is not needed.

CA RC/Migrator Enhancements

The following enhancements have been made to CA RC/Migrator in this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided in RC/Merger.
- DB2 10 MEMBER CLUSTER option support for universal tablespaces
- DB2 10 temporal table support is now provided.
- DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) support is now provided.
- DB2 10 unique index INCLUDE column support is now provided.
- Extended DB2 9 clone support is now provided.
- DB2 9 histogram statistics support is now provided.
- DB2 9 implicitly hidden columns support is now provided.
- Full support for DB2 9 native stored procedures (NSPs) with multiple versions is now provided.
- DB2 9 ROW CHANGE TIMESTAMP column support is now provided.
- An AUTOREPAIR option is now provided in RC/Merger.
- Parallel unload and load support is now provided.
- Partition by growth (PBG) support has been enhanced.
- Overriding the default SQL terminator is now supported in the Batch Processor.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

Extended RBA and LRSN support is now provided in the RC/Merger component. RC/Merger can now automatically detect RBA format differences between the source and target tablespaces and update the target catalog accordingly.

DB2 10 MEMBER CLUSTER Option Support for Universal Tablespaces

The MEMBER CLUSTER clause is no longer omitted when migrating a universal tablespace (type R and G) to another DB2 10 subsystem. When migrating to a DB2 9 subsystem, an informational message indicates that the clause is dropped.

When an analysis is performed on a migration strategy in CA RC/Migrator and that strategy contains a tablespace with MEMBER CLUSTER defined, the tablespace DDL contains the MEMBER CLUSTER clause.

Note: This enhancement was added in Version 17.0 as a post-GA PTF.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided.

The following new table types are now supported:

- Temporal
- History (for system-period or bi-temporal tables)

When altering a table, you can now edit a temporal table in a migrate or compare strategy. This processing includes:

- Adding or changing a default attribute of a column defined AS ROW BEGIN, AS ROW END, or AS TRANSACTION START ID.
- Adding a PERIOD definition.

You can also use the Alternate Catalog Mapping (ACM) analysis option in migrate, alter, and compare strategies.

When creating, altering, or templating a table, you can now add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint. To enable this support, a new option, B, is provided on the Unique Constraints Management panel.

When creating, altering, or templating an index, you can now specifically include or exclude the `BUSINESS_TIME WITHOUT OVERLAPS` clause in the index key. To enable this support, a new field, `BTWOO`, is provided on the Index Create, Alter, and Template panels.

Note: For more information about using temporal tables, see the *CA RC/Migrator User Guide*.

DB2 10 TIMESTAMP Support

DB2 10 TIMESTAMP support is now provided. CA RC/Migrator can now process tables containing DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) data types.

DB2 10 Unique Index INCLUDE Column Support

DB2 10 New Function Mode (NFM) includes a new INCLUDE clause on the CREATE INDEX and ALTER INDEX statements. The INCLUDE clause lets you define nonkey columns in a unique index. These nonkey columns let queries use the unique index for index-only access, so you can eliminate indexes that were created solely to enable index-only access. Eliminating unnecessary indexes can help improve your system performance, simplify index maintenance, and decrease physical storage requirements.

Note: For more information about adding INCLUDE column support, see the *CA RC/Migrator User Guide*.

DB2 9 Clone Support

DB2 9 cloned objects are now supported. Clones can provide maximum data availability. For example, you can replace tables quickly by applying changes to a clone table and then exchanging the clone table for the base table. If a base table has a clone, the base and clone and all dependent objects are migrated.

Note: For more information about clone support, see the *CA RC/Migrator User Guide*.

DB2 9 Histogram Statistics Support

DB2 9 histogram statistics are now supported. CA RC/Migrator can update the DB2 9 histogram statistics in SYSIBM.SYSCOLDIST so that the optimizer chooses the desired SQL access path on migrated DB2 objects (without having to run RUNSTATS).

DB2 9 Implicitly Hidden Columns Support

DB2 9 implicitly hidden columns are now supported. Support is now provided for the IMPLICITLY HIDDEN clause in a table. The DB2 result set of a SELECT * statement includes the column only when it is explicitly referenced. This setting can be beneficial for sites that use SELECT *. Database administrators can add new columns to tables without having to change all the programs that use the modified tables.

Extended DB2 9 Native SQL Stored Procedures Support

Full support for DB2 9 native stored procedures (NSPs) with multiple versions is now available.

DB2 9 ROW CHANGE TIMESTAMP Support

The DB2 9 ROW CHANGE TIMESTAMP column type is now supported. For each row change (insert or update), DB2 generates a timestamp value for the column that corresponds to the time of the insert or update. When GENERATED BY DEFAULT is specified, the timestamp column values can be specified manually. Knowing how often updates occur during specific time ranges helps Database Administrators plan for data replication and auditing tasks. This enhancement also helps to perform the following tasks:

- Ensure data integrity and limit the time that locks are held when optimistic locking is implemented.
- Reduce application development costs when DB2 handles the timestamp logic automatically.

AUTOREPAIR Option in RC/Merger

An AUTOREPAIR option has been added to the Move and Copy options in the RC/Merger component. This option repairs the catalog data of the tablespace when the Relative Byte Address (RBA), VERSION, or ROW FORMAT differ in the source and target subsystems. Previously, a manual repair was required, but you can now set the AUTOREPAIR option to YES to detect any differences and perform the repair automatically.

Note: For more information about using this option, see the *CA RC/Migrator User Guide*.

Parallel Unload and Load Support

Parallel unload and load processing using CA Fast Unload and CA Fast Load utility statements is now supported. CA RC/Migrator can now generate unloads into multiple streams balancing the streams so that the jobs can complete in a shorter amount of time. This processing is useful for large tables with many partitions to unload all the data more quickly.

Use utility profile services (CA RC/Migrator Expert profile menu option U) to enable this processing.

To support this enhancement, a new @DEFAULT model is provided. Execute the CA RC/Migrator Load and Customize the Model Services DB2 catalog customization task to refresh this model. DB2 SYSADM authority is required to execute the DB2 catalog customization tasks.

Note: For more information about executing this task, see the *CA Database Management Solutions for DB2 for z/OS Implementation Guide*. If you have a previously customized version of the @DEFAULT model, apply changes from the new model to your customized model. Detailed instructions are provided in the *CA RC/Migrator User Guide*.

Important! To exploit parallel unload and load processing in existing models, update the models as described in the CA RC/Migrator customization member (RCMCUST). This feature is not supported in releases before Version 16.0.

PBG Tablespace Enhancements

When you copy partitions of a partition by growth (PBG) tablespace, RC/Merger now adds any extra partitions for you, if needed. You no longer have to use the Add Partition utility (RMMUTIL in *hlq.CDBASRC*) to add partitions to the target subsystem.

The following enhancements have been made:

- If the source grows beyond the first partition, RC/Merger automatically creates the appropriate number of partitions for the target subsystem.
- If the target subsystem has more partitions than the source, RC/Merger generates LOAD REPLACE statements in the analysis DDL. These statements delete the data in the extra partitions on the target side.
- If the number of partitions in the source tablespace exceeds the MAXPARTITION value of the target tablespace, RC/Merger alters the MAXPARTITION value of the target tablespace to accommodate the extra partitions. Any existing data in the target partitions is removed.

Set SQL Terminator in the Batch Processor

You can now set the SQL terminator character for an SQL statement in the Batch Processor component of CA RC/Migrator. The default SQL terminator is the semicolon (;). However, you can now override the default by specifying the following syntax before the SQL statement for which you are setting the new terminator:

```
--#SET TERMINATOR character
```

SPUFI interprets *character* as a statement terminator. The option to override the default SQL terminator is useful when you want to specify a different SQL terminator for one or more SQL statements.

CA RC/Query Enhancements

The following enhancements have been made to CA RC/Query in this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 temporal table support is now provided.
- DB2 9 roles and trusted context support is now provided.
- DB2 9 ROW CHANGE TIMESTAMP column support is now provided.
- A new ALL command is now provided.
- CA RC/Query can now generate a SPUFI-compatible HDDL. Object selection for SPUFI-compatible HDDL output is now possible.
- The index reports have been enhanced with columns added to the Index Detail (I-D) and Index List (I-L) reports. The I-D report now has a new Sparse column, and the I-L report now has a new Include column.
- The routine reports have been enhanced with a new package (PK) report option.
- The table reports have been enhanced with a new CLONE line command. History table information for temporal tables is now displayed in Table Drop Impact (T-DI) reports.
- Additional wildcard support has been added to the CA RC/Query header fields.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

Extended logging support is now provided in the following CA RC/Query reports:

- System Image Copy Detail (SY-IC)
- System List DB2 Utilities (SY-L)
- Table Column (T-C)

The new CATALOG and CATALOG TEST keywords are now supported in the repair utility.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA RC/Query now displays temporal history tables.

A new table type of H has been added to enable this support. The detail report for this table includes versioning information for the schema name and table name of the history table or system-maintained temporal table.

DB2 9 Roles and Trusted context Support

DB2 9 roles and trusted context support (GRANT and REVOKE privileges) is now provided for the following objects:

- Collection IDs
- Databases
- Functions
- Packages
- Plans
- Schemas
- Sequences
- Tables
- Views
- Distinct types
- Buffer pools
- Storage groups
- Tablespaces

Associating privileges with trusted contexts and roles helps improve overall security by controlling when privileges are made available, based on the trusted connections.

CA RC/Query also now provides drop impact and user authorization reports for roles.

DB2 9 ROW CHANGE TIMESTAMP Support

The DB2 9 ROW CHANGE TIMESTAMP column type is now supported. For each row change (insert or update), DB2 generates a timestamp value for the column that corresponds to the time of the insert or update. When GENERATED BY DEFAULT is specified, the timestamp column values can be specified manually. Knowing how often updates occur during specific time ranges helps Database Administrators plan for data replication and auditing tasks. This enhancement also helps to perform the following tasks:

- Ensure data integrity and limit the time that locks are held when optimistic locking is implemented.
- Reduce application development costs when DB2 handles the timestamp logic automatically.

ALL Command Support

A new ALL command is now supported. This command lets you copy any line command to all lines of a generated report, and then execute the command for all lines. You can include (=) and can exclude (=X) lines as needed.

The ALL command is useful when a change applies to many objects. You can now make changes without manually typing equal (=) for all items in the report list.

Note: For the utilities (like COPY, LOAD), there is one job step for each object. An option is provided to apply the changes to a single object or all objects. If more than 255 objects are selected, a new job statement is inserted in the JCL automatically after every 255 steps.

The following new model JCL members are used when a user specifies YES for USE THESE OPTIONS FOR ALL SELECTED OBJECTS:

- MJUTLGU (driver for unload)
- MJUTALLL (load for multiple objects)
- MJUTALLU (unload for multiple objects)

Configure these members as needed to use site-specific naming conventions for data sets that the unload and load utilities generate. The following standard naming conventions are used:

- %USERID..TABLE.UNLOAD.DATA.I%INCR (for UNLOAD)
- For load, %DATASET..I%INCR.

You can configure these model JCL members so that the unload and load utilities generate data sets with those names.

SPUFI-Compatible HDDL Output Support

You can now generate SPUFI-compatible Hierarchical DDL (HDDL) output. This enhancement lets you run scripts in online mode instead of batch mode, which can save processing time when the batch queues are busy.

A new field, SPUFI Compatible, is provided on the HDDL Request Options panel to enable this feature.

You can select the objects for which you would like to generate the SPUFI-compatible HDDL output. When you specify YES in the SPUFI Compatible field and select a set of objects, the HDDL command generates DDLs for the selected objects.

Index Report Enhancements

The following enhancements have been made to Index object reports:

Index Detail (I-D) report

The Sparse column has been added to the Index Detail (I-D) report. The column indicates whether an index is sparse. A sparse index has fewer index entries compared to other indexes; it may not have an entry for each data row in the table.

Index List (I-L) report

The Include column has been added to the Index List (I-L) report. The column indicates whether an index is an Include index. Include indexes have additional non-key columns called Include columns added to the set of key columns. Only unique indexes support Include columns. The use of Include columns reduces the need for additional indexes, lowering index maintenance and physical storage requirements.

Routine Report Enhancements

You can now view package details associated with a routine using the new PK line command. When you specify R (routines) as the primary object, the PK option is now available.

The Routine Package report lists the dependent packages and provides basic information about the related packages.

Table Report Enhancements

The following enhancements have been made to Table reports:

New CLONE line command

You can now execute a CLONE line command on a table (T) object to create its clone table. You can exchange data between the base table and the clone table by using the EXCHANGE command, whenever necessary.

Table Drop Impact (T-DI) Report

History table information for temporal tables is now displayed in Table Drop Impact (T-DI) reports. You can execute all line commands for table objects from the history table line in T-DI reports.

Wildcard Character Support

The asterisk (*) is now supported as a wildcard character in CA RC/Query header fields. The asterisk is identical to the percent sign (%) wildcard, which lets you substitute a string of zero or more characters in a DB2 query.

CA RC/Secure Enhancements

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

CA RC/Update Enhancements

The following enhancements have been made to CA RC/Update in this release:

- DB2 11 NFM and CM support is now provided.
- DB2 10 temporal table support is now provided.
- DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) support is now provided.
- DB2 10 unique index INCLUDE column support is now provided.
- DB2 9 clone support is now provided.
- DB2 9 roles and trusted context support is now provided.
- DB2 9 ROW CHANGE TIMESTAMP column support is now provided.
- Full support for native stored procedures (NSPs) with multiple versions is now provided.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided.

The following new table types are now supported:

- Temporal
- History (for system-period or bi-temporal tables)

When altering a table, you can now:

- Add or change a default attribute of a column defined AS ROW BEGIN, AS ROW END, or AS TRANSACTION START ID.
- Add a PERIOD definition.

You can also use the Alternate Catalog Mapping (ACM) analysis option in strategies.

When creating, altering, or templating a table, you can now add the BUSINESS_TIME WITHOUT OVERLAPS clause to a unique constraint. To enable this support, a new option, B, is provided on the Unique Constraints Management panel.

When creating, altering, or templating an index, you can now specifically include or exclude the BUSINESS_TIME WITHOUT OVERLAPS clause in the index key. To enable this support, a new field, BTWOO, is provided on the Index Create, Alter, and Template panels.

Note: For more information about using temporal tables, see the *CA RC/Update User Guide*.

DB2 10 TIMESTAMP Support

DB2 10 TIMESTAMP support is now provided. CA RC/Update can now process tables containing DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) data types.

DB2 10 Unique Index INCLUDE Column Support

DB2 10 New Function Mode (NFM) includes a new INCLUDE clause on the CREATE INDEX and ALTER INDEX statements. The INCLUDE clause lets you define nonkey columns in a unique index. These nonkey columns let queries use the unique index for index-only access, so you can eliminate indexes that were created solely to enable index-only access. Eliminating unnecessary indexes can help improve your system performance, simplify index maintenance, and decrease physical storage requirements.

DB2 9 Clone Support

DB2 9 cloned objects are now supported. Clones can provide maximum data availability. For example, you can replace tables quickly by applying changes to a clone table and then exchanging the clone table for the base table.

Note: For more information about cloned object support, see the *CA RC/Update User Guide*.

DB2 9 Roles and Trusted Context Support

DB2 9 support for roles and trusted context is now provided.

CA RC/Update alter, create, template, and drop options now support DB2 9 roles and trusted contexts. Associating privileges with trusted contexts and roles helps improve overall security by controlling when privileges are made available, based on the trusted connections.

DB2 9 ROW CHANGE TIMESTAMP Support

The DB2 9 ROW CHANGE TIMESTAMP column type is now supported. For each row change (insert or update), DB2 generates a timestamp value for the column that corresponds to the time of the insert or update. When GENERATED BY DEFAULT is specified, the timestamp column values can be specified manually. Knowing how often updates occur during specific time ranges helps Database Administrators plan for data replication and auditing tasks. This enhancement also helps to perform the following tasks:

- Ensure data integrity and limit the time that locks are held when optimistic locking is implemented.
- Reduce application development costs when DB2 handles the timestamp logic automatically.

Extended DB2 9 Native SQL Stored Procedures Support

Full support for DB2 9 native stored procedures (NSPs) with multiple versions is now available.

CA Recovery Analyzer Enhancements

The following enhancements have been made to CA Recovery Analyzer with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 11 extended logging support is now provided.
- DB2 10 temporal table support is now provided.
- A new Block Size field is now provided for disaster recovery strategies.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 11 Extended Logging Support

DB2 11 extended logging support is now provided. Existing 6-byte RBA and LRSN values do not provide sufficient log addressing capacity in larger systems, which are expected to exhaust these values soon. DB2 11 extends the RBA and LRSN values from 6 bytes to 10 bytes, significantly increasing the log addressing capacity.

CA Recovery Analyzer can now process the extended RBA and LRSN values.

Note: Before using a pre-17.0 strategy to process extended log values, use the current CA Recovery Analyzer online panels to reanalyze the strategy. The reanalysis regenerates the control statements for you. For more information about updating strategies, see the *CA Recovery Analyzer User Guide*.

DB2 10 Temporal Table Support

DB2 10 temporal table support is now provided. Temporal tables provide an efficient means to maintain versioned data and track data trends in your DB2 system. These tables have defined time periods that contain the start and end timestamp values for each row. The time periods indicate the time range when a row is valid.

DB2 supports three types of temporal tables:

- System-maintained temporal tables have a defined `SYSTEM_TIME` period. DB2 maintains the timestamp values in these tables. When a system-maintained table is updated, DB2 also automatically archives the older rows to an associated history table.
- User-maintained temporal tables have a defined `BUSINESS_TIME` period. The user is responsible for maintaining the timestamp values in these tables. User-maintained tables do not have an associated history table.
- Bi-temporal tables combine both types of time periods. These tables let you track user-specified period information and system-based historical information in the same table.

CA Recovery Analyzer now supports all three temporal table types.

New Block Size Field for Disaster Recovery Strategies

A new field is now provided for specifying the block size to use when copying archive logs as part of a disaster recovery strategy. When you create the strategy, you can select the archive log method. This method rebuilds a subsystem from the archive log copies and image copies that you ship to a remote recovery site. Specifying the block size within the strategy inserts this block size into the resulting recovery JCL. You can match the block size to the size of the archive log copy instead of using a default size that may not match. Matching the block sizes can improve the recovery performance.

The new Block Size field is provided on the PRA Disaster Recovery Options panel. You can enter the following values in this field:

NONE

Does not include a block size in the recovery JCL. This setting is the default.

int

Specifies a block size. The value that you enter is rounded up to the next multiple of 4,096 (4 KB) and inserted into the recovery JCL. You can also enter the value with a K, to indicate kilobytes (for example, 8K).

Limits: 8192 to 28672 (or 8K to 28K)

ZPARM

Inserts the current BLKSIZE ZPARM value for the subsystem into the recovery JCL.

Note: For more information about disaster recovery strategies and the archive log method, see the *CA Recovery Analyzer User Guide*.

CA Report Facility Enhancements

The following enhancements have been made to CA Report Facility with this release:

- DB2 11 NFM and CM support is now provided.
- The RUN PROC command is now supported within a procedure.
- The procedure error threshold can now be defined.
- The KEEP column name suffix can now be specified.
- The precedence of sources for variable resolution can now be specified.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

RUN PROC Command within a Procedure

CA Report Facility now lets you build a modular design and execute similar tasks from a separate procedure, but by any procedure where needed. This enhancement is enabled by using the RUN PROC command within another procedure without any interruption to processing. The number of embedded procedures and the level of embedding are not limited.

Note: To prevent recursion that can cause an infinite loop of processing of the procedures, do not allow a procedure to call itself either directly or indirectly. For more information about procedures, see the *CA Report Facility User Guide*.

Define the Procedure Error Threshold

CA Report Facility now lets you define a procedure error threshold in the system profile settings. Procedure error threshold indicates the error level at which a batch job stream processing stops across the system. The error level indicator matches the severity code on CA Report Facility messages: S (severe), E (error), or W (warning).

The Procedure Error Threshold option is available under Report/Query Controls on the System Profile panel for the TSO, CICS, and IMS systems.

CA Report Facility users can override this threshold for their sessions using the ENVSET PROCERRS subcommand.

Note: For more information about system profiles, see the *CA Report Facility Administration Guide*. For more information about the ENVSET subcommands, see the *CA Report Facility User Guide*.

Specify the KEEP Column Name Suffix

CA Report Facility now lets you specify what suffix to append to column names in the KEEP table. KEEP tables are DB2 tables that CA Report Facility users create to store detail data of their queries. The data can then be reused with other queries at other times.

CA Report Facility now supports two formats for KEEP table column names:

- CA Report Facility format—The column name includes a suffix naming the table from which the data item is sourced and a replicate index when that source is repeated in the source query. The table indicator is included only if there is more than one table in the source query. The source column name is truncated to accommodate the suffix.

Example: COLUMN_002_003 indicates that the column was sourced from the second table in the source query and the third repeat of the same column name.

- IBM QMF-compatible format—The column name is suffixed with its relative appearance number only if the name is repeated in the target table.

Example: COLUMN3 indicates that the column is the fourth iteration of the name COLUMN is the resulting KEEP table. The first iteration is not suffixed.

The KEEP Column Name Suffix option is available under Report/Query Controls on the System Profile panel for the TSO, CICS, and IMS systems.

CA Report Facility users can override this setting for their sessions using the ENVSET KEEPSFX subcommand.

Important! This setting changes how the current KEEP commands name columns in the KEEP tables. This change can break any queries that are based on current KEEP tables if the table is dropped and recreated. We recommend changing this setting only if an isolated need for compatible names is required.

Note: For more information about system profiles, see the *CA Report Facility Administration Guide*. For more information about the ENVSET subcommands, see the *CA Report Facility User Guide*.

Specify the Precedence of Sources for Variable Resolution

When resolving replaceable parameters in procedures, CA Report Facility always looks first for values that are provided in the command. If the command override value is not available, CA Report Facility searches object default values or global variable values. CA Report Facility now lets you specify whether object default values are searched before global variable values, or opposite.

The Variable resolution order option is available under Report/Query Controls on the System Profile panel for the TSO, CICS, and IMS systems.

CA Report Facility users can override this setting for their sessions using the ENVSET VARORDER subcommand.

Note: This setting can affect existing procedures and batch jobs. We recommend that you analyze the impact of promoting global variables before you change this setting.

For more information about system profiles, see the *CA Report Facility Administration Guide*. For more information about the ENVSET subcommands, see the *CA Report Facility User Guide*.

CA SQL-Ease Enhancements

The following enhancements have been made to CA SQL-Ease with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 10 index probing is now supported.
- DB2 10 row and column access control is now supported.
- A Create Explain Table(s) facility is now provided.
- Conditionally indented SQL substatements are now supported.
- A new DB2 Profiles Services facility is now provided to facilitate the management of DB2 profiles.
- Explain table indexes can now be automatically created during a future explain.
- Physical Rule 1083 is now supported.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

DB2 10 Index Probing Support

DB2 10 index probing is now supported. CA SQL-Ease Enhanced Explain now reports on statistics that the DB2 optimizer uses during index probing.

DB2 uses index probing to obtain more accurate filter factor estimates for matching predicates when conditions prevent the DB2 optimizer from obtaining an accurate estimate. Improved filter factor estimates can help stabilize the access path.

Note: For more information about how DB2 uses index probing, see the *IBM DB2 for z/OS Performance Topics* guide.

Index probing is supported through the following new features:

- CA SQL-Ease enhanced explain Index Probe Statistics Report
- DSN_COLDIST_TABLE and DSN_KEYTGDIST_TABLE explain tables

The Index Probe Statistics report shows how index probing influenced the access path that the DB2 optimizer selected during SQL execution. For example, you can determine which runtime statistics the DB2 optimizer applied to the query.

The Index Probe Statistics Report consists of two subreports. The first subreport focuses on column distribution statistics. This report is generated only when statistics are present in the DSN_COLDIST_TABLE explain table. The second subreport focuses on key target distribution statistics. This report is generated only when statistics are present in the DSN_KEYTGTDIST_TABLE explain table. Each report is generated on a per statement basis.

Two new explain tables, DSN_COLDIST_TABLE and DSN_KEYTGTDIST_TABLE facilitate the generation of this report. Create these tables using the new Create Explain Table(s) facility.

DB2 10 Row and Column Access Control Support

DB2 10 row and column access control is now supported.

This feature lets you manage access to a table at the row and column level by using row permissions and column masks. Controlling access in this way allows you to enforce stricter security policies that prevent users from accessing rows and columns that contain sensitive information.

The SQL Dependency Analysis report now includes object support for column masks and row permissions. Generate this report using the CA SQL-Ease Enhanced Explain facility.

Note: For more information about generating enhanced explain reports, see the *CA SQL-Ease User Guide*.

New Create Explain Table(s) Facility

A Create Explain Table(s) facility is now provided. This facility expands your ability to manage the explain output and generate enhanced explain reports from within CA SQL-Ease.

This functionality is enabled through a new main menu option, Create Explain Table(s). You can also specify SQLEASE EXPLTB from an edit session on the SQL-EASE EDIT ENTRY PANEL.

CA SQL-Ease now supports the creation of the following explain tables:

PLAN_TABLE

Contains access path information for an explained SQL statement.

DSN_STATEMNT_TABLE

Contains the estimated cost of execution for an explained SQL statement.

DSN_FUNCTION_TABLE

Contains user-defined functions, including name and schema, for an explained SQL statement.

DSN_COLDIST_TABLE

Contains the column distribution statistics for an explained SQL statement.

DSN_KEYTGDIST_TABLE

Contains the key target distribution statistics for an explained SQL statement.

Note: For more information about how to create explain tables, see the *CA SQL-Ease User Guide*.

Conditionally Indent SQL Substatements

You can now control how your SQL statements display. SQL substatements can now be conditionally indented, making complex output easier to read.

This new functionality is enabled through the new PPINDENT parameter in the PPA parmlib member.

Note: For more information about editing product-specific parmlib members, see the *Implementation Guide*.

New DB2 Profile Services Facility

A new DB2 Profile Services facility is now provided. This facility simplifies DB2 profile management by facilitating the process of managing profiles and profile attributes in the DSN_PROFILE_TABLE and DSN_PROFILE_ATTRIBUTES tables. For example, this facility helps verify that only valid profiles and attributes are added to the profile tables.

This facility supports the following management tasks for profiles that are based on optimization parameter or modeling functions:

- Create, enable, disable, update, browse, delete, or obtain the status of a profile
- Start or stop profiles and determine whether profiles are running

Note: You cannot create or update profiles that are based on the MONITOR or ACCEL functions. However, you can delete, enable, start, stop, display, or view status reports for these profiles.

This functionality is accessed through the new DB2 Profile Services option that is provided on the CA SQL-Ease main menu.

The following new panel-specific and line commands also support this functionality:

STARTPROF

Starts all enabled profiles and displays all DB2 command messages. The status reports for the enabled profiles follow these messages.

STOPPROF

Stops all active profiles.

DISPROF

Shows whether profiling is active or inactive.

S

Switches the profile status from enabled to disabled, or from disabled to enabled.

ST

Displays a status report for the selected profile.

CMA

Creates a MODELING profile using parameters that were captured from the modeled DB2 subsystem.

Note: This command requires access to Thread Termination\Dynamic DSNZPARM, which is a Value Pack component. The Value Pack components are provided at no additional charge with the CA Database Management Solutions for DB2 for z/OS.

SQLEASE PS

Lets you access the DB2 Profile Services option from an edit session in CA SQL-Ease.

If a profile is started, the Access Path report identifies the profile ID that was applied to the explain statement. This information is provided in all of the Access Path report formats.

Note: For more information about how to use this facility, see the *CA SQL-Ease User Guide* and the online help.

Create Explain Table Indexes Automatically During a Future Explain

Explain table indexes can now be automatically created during a future explain.

Indexes can help retrieve table data more quickly. Creating the indexes automatically reduces the possibility that the index is set up incorrectly.

This new functionality is enabled through the new AUTOCRIX parameter in the PPA parmlib member.

Note: For more information about editing product-specific parmlib members, see the *Implementation Guide*.

New Physical Rule 1083

You can now use new Physical Rule 1083 to identify opportunities for index optimization, which can help reduce index maintenance and overhead. This rule is part of the default rule set and displays in the PPA Explain Physical Rules Report. This rule applies to systems running only on DB2 10 or higher.

Note: For more information about the rule-based Expert System application, see the *CA Plan Analyzer User Guide* or *CA SQL-Ease User Guide*.

CA Subsystem Analyzer Enhancements

The following enhancements have been made to CA SQL-Ease with this release:

- DB2 11 NFM and CM support is now provided.
- You can now view the amount of space that each collection interval uses.

DB2 11 Support

This product has been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM).

View the Collection Interval Size

You can now view the amount of space that each collection interval uses. This information helps prevent storage issues by showing how much space an interval uses, especially when collection options are changed.

This functionality is provided in the new SIZE (MB) column on the Datastore Interval Display.

Value Pack Enhancements

The following enhancements have been made to the Value Pack components with this release:

- DB2 11 NFM and CM support is now provided.
- DB2 10 DSNZPARM support is now provided in Thread Termination\Dynamic DSNZPARM.
- DB2 10 TIMESTAMP support is now provided in the Batch Processor.
- Increased active log data set support is now provided in Thread Termination\Dynamic DSNZPARM.
- Data sharing support is now provided for the -DISPLAY CA command.
- Named relationship sets are now supported in RI Manager.

DB2 11 Support

The following Value Pack components have been enhanced to run in DB2 11 New Function Mode (NFM) with a converted catalog and in DB2 11 Conversion Mode (CM):

- Batch Processor
- RI Manager
- Thread Termination\Dynamic DSNZPARM

DB2 10 DSNZPARM Support

DB2 10 DSNZPARM support is now provided in Thread Termination\Dynamic DSNZPARM for parameters in the following Thread Termination\Dynamic DSNZPARM categories:

- Utilities Parameters
 - CHECK_FASTREPLICATION
 - CHECK_SETCHKP
 - DB2SORT
 - FLASHCOPY_COPY
 - FLASHCOPY_LOAD
 - FLASHCOPY_REBUILD_INDEX
 - FLASHCOPY_REORG_INDEX
 - FLASHCOPY_REORG_TS
 - UTIL_TEMP_STORCLAS

- Performance and Optimization Parameters
 - PARA_EFF
 - PLANMGMTSCOPE
- Storage Parameters
 - DPSEGSZ
 - MAXTEMPS
 - LOB_INLINE_LENGTH
 - MAXTEMPS_RID
- Additional Parameters
 - CATDDACL
 - CATDMGCL
 - CATDSTCL
 - CATXDACL
 - CATXMGCL
 - CATXSTCL
 - DISALLOW_DEFAULT_COLLID
 - IMPDSSIZE
 - INDEX_IO_PARALLELISM
 - REALSTORAGE_MANAGEMENT
 - XML_RANDOMIZE_DOCID
 - DDLTOX
- Tracing and Checkpoint Parameters
 - SMFCOMP
- Security Parameters
 - SECADM1
 - SEPARATE_SECURITY
 - SECADM2
 - SECADM1_TYPE
 - SECADM2_TYPE
- Routine (Stored Procedure) Parameters
 - BIF_COMPATIBILITY

Support for these DSNZPARMS lets you dynamically update DB2 10 DSNZPARM values without having to reassemble the DSNZPARM module.

Note: For more information about the supported DSNZPARMs and set commands, see the *Thread Termination\Dynamic DSNZPARM User Guide*.

DB2 10 TIMESTAMP Support

DB2 10 TIMESTAMP support is now provided. Batch Processor can now process tables containing DB2 10 TIMESTAMP WITH TIME ZONE and TIMESTAMP(*precision*) data types.

Note: This support was added in Version 16.0 after GA.

Increased Active Log Data Set Support

Additional active log data set support is now provided in Thread Termination\Dynamic DSNZPARM. This support is provided for users running DB2 10 or higher. Log data sets that contain extended RBA values are supported for users running DB2 11.

Thread Termination\Dynamic DSNZPARM now lets you display and dynamically add up to 93 defined active log data sets from the DB2 BSDS list. Doing so allows you to manage your active log data sets without scheduling DB2 downtime.

This functionality is enabled through the Add Active Logs option on the Thread Termination\Dynamic DSNZPARM main menu.

Note: Active log data sets can no longer be deleted. For more information about *adding* active log data sets, see the *Thread Termination\Dynamic DSNZPARM User Guide*.

Data Sharing Support for -DISPLAY CA ACTIVE Command

After the version 16.0 GA release of the CA Database Management Solutions for DB2 for z/OS, data sharing support was added for the -DISPLAY CA ACTIVE command. This command is executed through the DB2 Command Processor. The ACTIVE option shows the utility status and the unload phase and count. For example, you can see which utilities are running and how much progress they have made. Being able to view this information across all members of a DB2 data sharing system helps you monitor your utility jobs more efficiently.

Note: For more information about the -DISPLAY CA command and its syntax options, see the *Value Pack Reference Guide*.

Named Relationship Sets

You can use RI Manager to name and save a set of referential integrity (RI) relationships for a table, including user-defined relationships. For example, in addition to the DB2 primary key, there can be several application managed relationships. You can define several user RI relationships for the same table. This functionality lets you create test data for a table using different foreign table relationships.

Chapter 2: Documentation

This chapter describes the documentation set for the CA Database Management Solutions for DB2 for z/OS. Updated guides are available from the Download Center on CA Support Online.

This section contains the following topics:

[CA HTML Bookshelf](#) (see page 103)

[Common Guides](#) (see page 105)

[Product-Specific Documentation](#) (see page 105)

[Edition Numbers](#) (see page 107)

CA HTML Bookshelf

This release contains the CA HTML bookshelf, which is an HTML help system that provides access to all deliverables in the product documentation set in both HTML and PDF. HTML provides robust online viewing and search capabilities, while PDF provides a print-friendly option.

The HTML bookshelf features include:

- A single help screen that displays all documentation for this release.
- An all-in-one search tool that searches the entire documentation set and returns matches found in both the HTML and PDF formatted documentation, without the need for a specialized .PDX index file.
- Additional links for using the bookshelf, downloading Acrobat Reader, and contacting CA Technologies.

Note: You must have Adobe Reader 8 or above to view the PDF files in the bookshelf.

Access the Bookshelf

You can access the bookshelf from any page on CA Support Online (CSO) using the Support By Product or Documentation links in the left navigation column.

To access the bookshelf

1. Log in to CA Support Online.
2. Click Documentation in the left navigation column. This option lets you view a list of all available bookshelves.
3. Type **CA Database Management Solutions for DB2 for z/OS** in the Select a Bookshelf drop-down list.

4. Select the bookshelf that you want to open from the list of available bookshelves. Click Go.

The bookshelf opens in a separate window.

5. Select the guide that you want to open from the bookshelf. You can view the HTML version or you can download the PDF.

Access the Guides

If your documentation set has a large number of guides, the initial display may not list the guides.

To access the guides through the bookshelf

1. Click Show All in the upper right corner above the All Documentation section.

All the guides are listed on the bookshelf.

Note: If you click Hide All, the list of documents disappears. You can select from the letter bar to list only those documents whose titles begin with the letter you select.

2. Click the HTML or PDF link next to the document you want to open.

The document opens.

Search the Bookshelf

The bookshelf includes a search facility that helps you locate information throughout the set.

To search the bookshelf

1. Enter your search criteria in the Search field in the upper right corner of the bookshelf and press Enter.

The search returns HTML results listed by topic and PDF results listed by guide. The results are sorted by date so that the most recently updated topics or PDFs appear at the top of the list. To find a topic in a PDF, open the PDF and view the list of topics within the PDF that match the search criteria.

2. (Optional) Click Sort by Relevance.

The list is reordered so that the HTML topics or PDFs that contain the most matches appear at the top of the list.

Common Guides

The following common guides are provided for the CA Database Management Solutions for DB2 for z/OS:

- *Batch Processor Reference Guide*
- *General Facilities Reference Guide*
- *Implementation Guide*
- *Installation Guide*
- *Message Reference Guide (AGT to PRR)*
- *Message Reference Guide (PS to ZLB)*
- *Release Notes*
- *RI Manager User Guide*
- *Thread Termination\Dynamic DSNZPARM User Guide*
- *Utilities Quick Reference Guide*
- *Value Pack Reference Guide*

Product-Specific Documentation

The product-specific documentation is organized into the following functional areas, which parallel the major tasks performed at DB2 sites:

- Backup and Recovery
- Database Administration
- Database Performance Management
- Reporting and Information Management

Backup and Recovery

The following documentation is provided for the backup and recovery solutions:

- *CA Fast Check for DB2 for z/OS User Guide*
- *CA Fast Recover for DB2 for z/OS User Guide*
- *CA Log Analyzer for DB2 for z/OS User Guide*
- *CA LogCompress for DB2 for z/OS User Guide*
- *CA Merge/Modify for DB2 for z/OS User Guide*

- *CA Quick Copy for DB2 for z/OS User Guide*
- *CA Recovery Analyzer for DB2 for z/OS User Guide*

Database Administration

The following documentation is provided for the database administration solutions:

- *CA Endeavor SCM Interface for DB2 for z/OS User Guide*
- *CA Fast Index for DB2 for z/OS User Guide*
- *CA Fast Load for DB2 for z/OS User Guide*
- *CA Fast Unload for DB2 for z/OS User Guide*
- *CA Partition Expert for DB2 for z/OS User Guide*
- *CA RC/Extract for DB2 for z/OS User Guide*
- *CA RC/Migrator for DB2 for z/OS User Guide*
- *CA RC/Query for DB2 for z/OS User Guide*
- *CA RC/Secure for DB2 for z/OS User Guide*
- *CA RC/Update for DB2 for z/OS User Guide*

Database Performance Management

The following documentation is provided for the database performance management solutions:

- *CA Bind Analyzer for DB2 for z/OS User Guide*
- *CA Database Analyzer for DB2 for z/OS Reference Guide*
- *CA Database Analyzer for DB2 for z/OS User Guide*
- *CA Detector for DB2 for z/OS Reference Guide*
- *CA Detector for DB2 for z/OS User Guide*
- *CA Index Expert for DB2 for z/OS User Guide*
- *CA Insight DPM for DB2 for z/OS Batch Report Reference Guide*
- *CA Insight DPM for DB2 for z/OS System Reference Guide*
- *CA Insight DPM for DB2 for z/OS User Guide*
- *CA Insight DPM for DB2 for z/OS Writing Requests Reference Guide*
- *CA NSM Database Option for DB2 for z/OS User Guide*
- *CA Plan Analyzer for DB2 for z/OS Reference Guide*

- *CA Plan Analyzer for DB2 for z/OS User Guide*
- *CA Rapid Reorg for DB2 for z/OS Quick Reference Guide*
- *CA Rapid Reorg for DB2 for z/OS User Guide*
- *CA SQL-Ease for DB2 for z/OS User Guide*
- *CA Subsystem Analyzer for DB2 for z/OS Quick Reference Guide*
- *CA Subsystem Analyzer for DB2 for z/OS Reference Guide*
- *CA Subsystem Analyzer for DB2 for z/OS User Guide*

Reporting and Information Management

The following documentation is provided for the reporting and information management solutions:

- *CA Compile/PRF User Guide*
- *CA InfoRefiner Concepts Guide*
- *CA InfoRefiner Reference Guide*
- *CA InfoRefiner and CA InfoTransport Workstation and Remote Components Installation Guide*
- *CA InfoTransport User Guide*
- *CA Report Facility Administration Guide*
- *CA Report Facility User Guide*

Edition Numbers

Occasionally, we must update documentation outside of a new or updated release. To indicate a minor change to the documentation that does not invalidate it for any releases that it supports, we update the edition number on the cover page. First editions do not have an edition number.