



# CA Recovery Suite for DB2 for z/OS

Jan Marek, Jacek Rafalak – May 31, 2018 – 2.14

## Prague Technology Days

May 30 - June 1, 2018



# For Informational Purposes Only

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

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

# Content

1 CA LOG ANALYZER

2 CA Fast Recover

3 CA Merge/Modify

4 CA Quick Copy

5 CA Recovery Analyzer

6 Q&A

# CA Log Analyzer

## Overview

---

# CA Log Analyzer Overview

- CA Log Analyzer analyzes the DB2 log (active and archive data sets), image copies, BSDS datasets and SMF records to help DBAs to
  - Audit data changes
  - Recover data
  - Recover dropped objects
  - Back out errant updates
  - Replicate changes to other subsystems or non DB2 relational targets
  - Get detailed recovery information

# CA Log Analyzer Overview

- Reports can be generated by purpose using online panels navigation
- Or using a strategy
  - Strategy can save the last log RBA/LRSN so the next time strategy is executed it can RESUME the processing at the saved point
    - All inflight URs at this point
    - This all together makes it easy to continuously monitor activity on one or more tables

# CA Log Analyzer Overview – Authorizations

- DB2 authorization can be used to determine whether a user has the authority to review or use the log data for specific tables.
- When DB2 authorization is enabled
  - Table authorizations are used to determine a user's authority on data in the log.
  - Primary and secondary authid information is used.
  - SYSUSERAUTH, SYSDBAUTH, and SYSTABAUTH are read to determine whether the user has SELECT authority on a table that has information in the log.

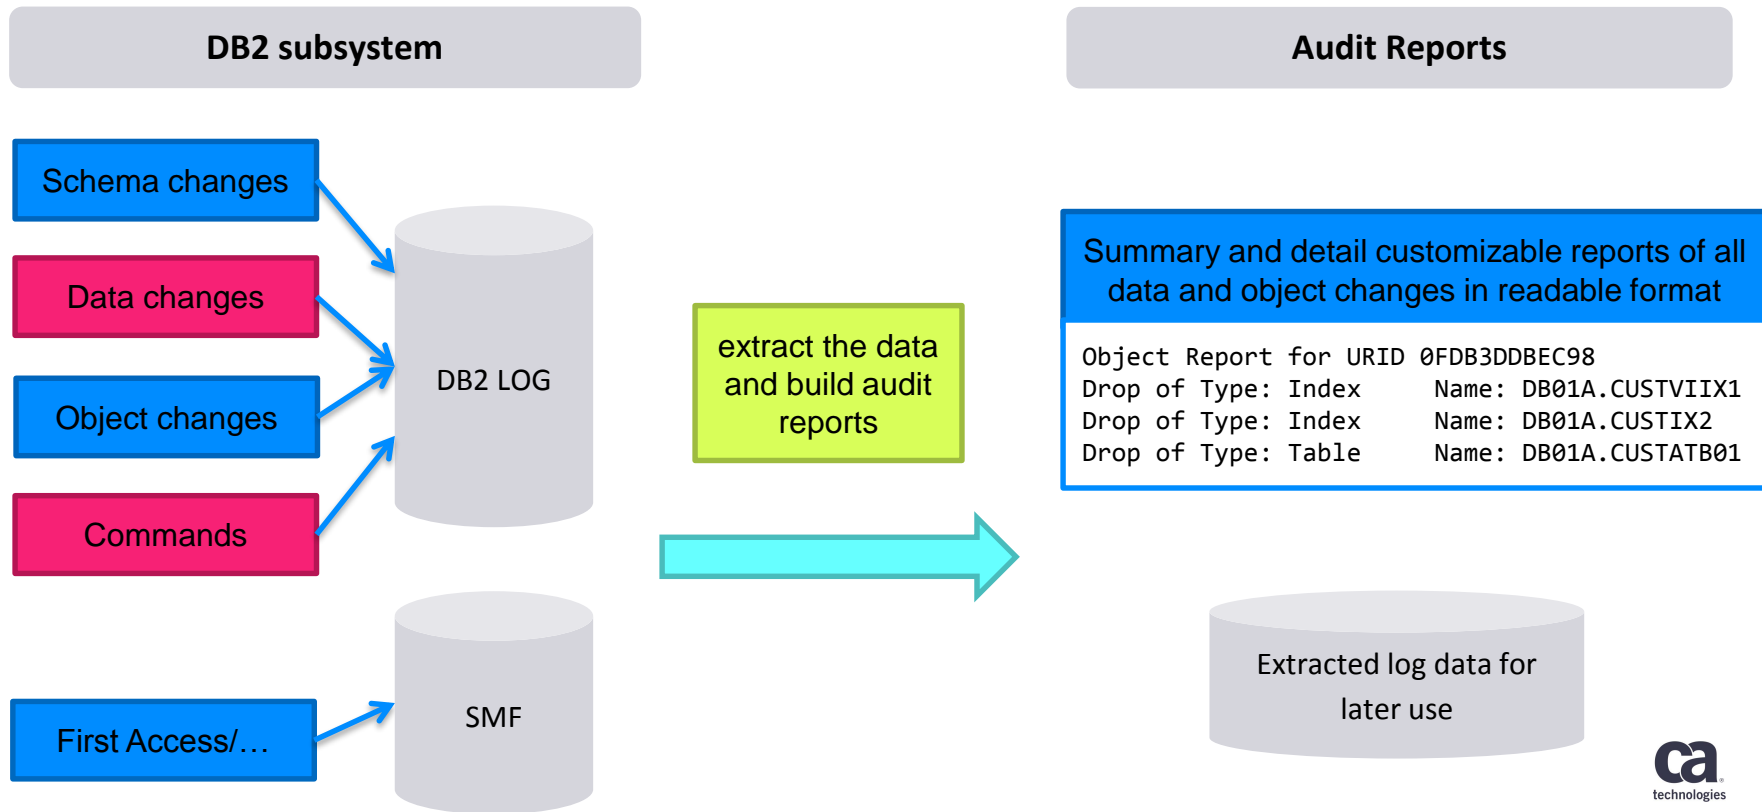
# Auditing

A thick, solid pink horizontal bar that underlines the word "Auditing".

# Log and SMF auditing

- DBA is required to keep track of all changes which impact objects as well as the data stored in the database. Some of those tracked changes must be recorded for later use
- CA Log Analyzer help with all those tasks by providing a variety of summary and detailed audit reports
  - You can build various customizable reports showing in summary or in detail what is being changed and who is doing the changes
  - You can use extracted log records to store them for later use (contains much less data than original DB2 log)
  - SMF Audit Facility provides reports on DB2 audit traces

# Log and SMF auditing



# Audit Object and Data Changes and DB2 commands

- Summary report
  - Number of updates, deletes, inserts, creates, alters or drops that occurred during the log processing range for the selected database objects
- Detail report
  - Shows a before and after image for each change that was made
  - If the level is not image copy and the object hasn't data capture changes only the partial image is shown
  - Key option shows updated columns in a separate section

# Audit Object and Data Changes and DB2 commands

- SQL REDO/UNDO
  - Shows actual SQL statement
  - The URID header from like in the Detail report could be included as comment in this report to identify who and when issued the statement
- DB2 Commands report
  - Shows issued DB2 commands with all details
- <https://youtu.be/Q2ush3lyDXM>

# SMF auditing

- Reports on DB2 audit traces recorded in the SMF data sets. The SMF Audit Trace reporting lets you generate reports on the following types of information:
  - Explicit grants and revokes
  - Creates, alters and drops
  - Access attempts denied
  - First access to object
  - First change to object
  - SQLID assignment or changes
  - Utilities performed
  - Binds performed
  - Start/stop trace activities

# Customized reports

- The Report Form Services facility let's you define and save report formatting options as a form for use with summary and detailed reports. You can customize
  - The report headings
  - Unit of recovery description
  - Summary and totals description
  - Table description
  - Update and updated row description
  - Object report description

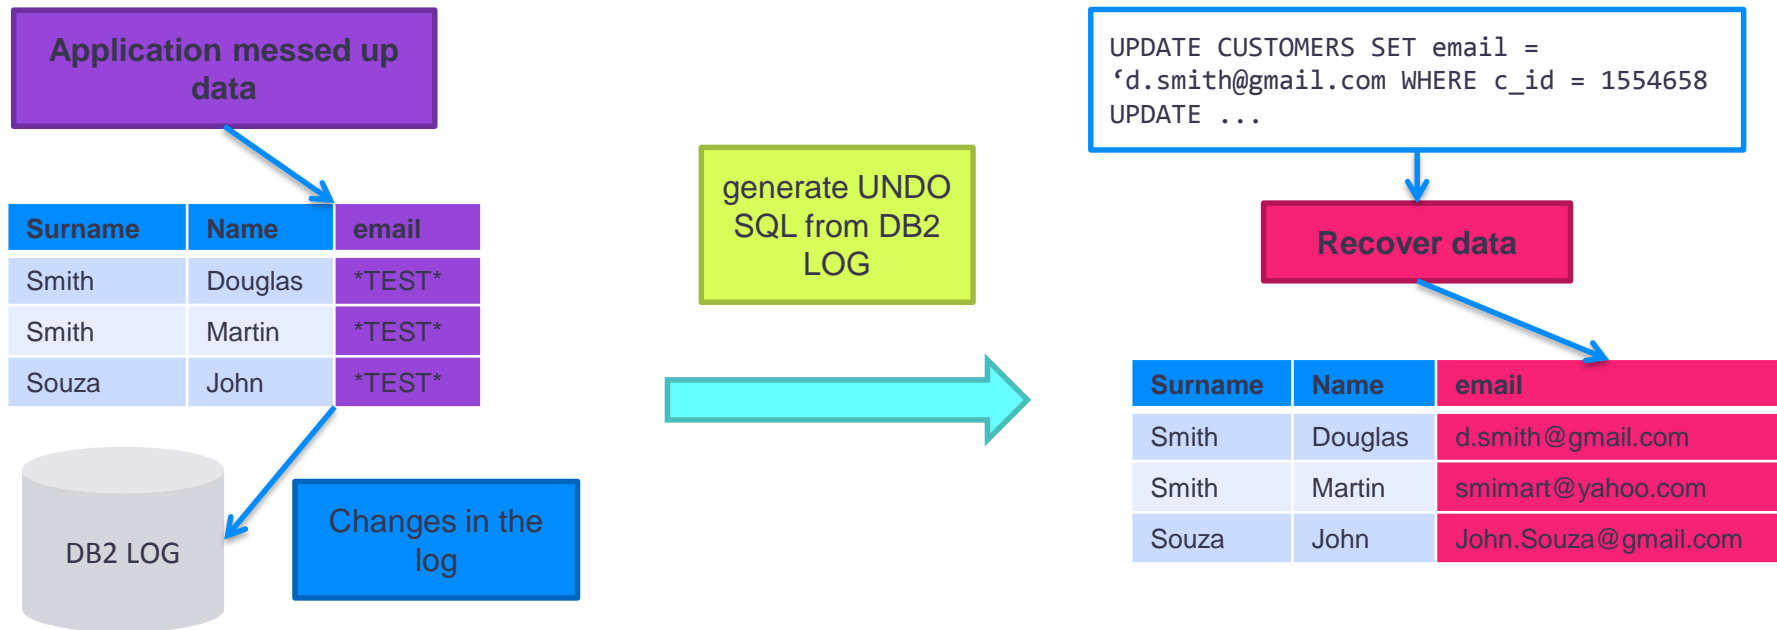
# **Data Recovery**

---

# Data recovery

- Application developers, while building the applications which access database objects, may mess up the data either on test subsystems or on production subsystems
- Or anybody else with the database access may do “something wrong”
- DBA is required to resolve this situation and recover the data into it's original state
- [https://youtu.be/i4ri\\_SAzsrM](https://youtu.be/i4ri_SAzsrM)

# Data recovery



[https://youtu.be/i4ri\\_SAzsM](https://youtu.be/i4ri_SAzsM)

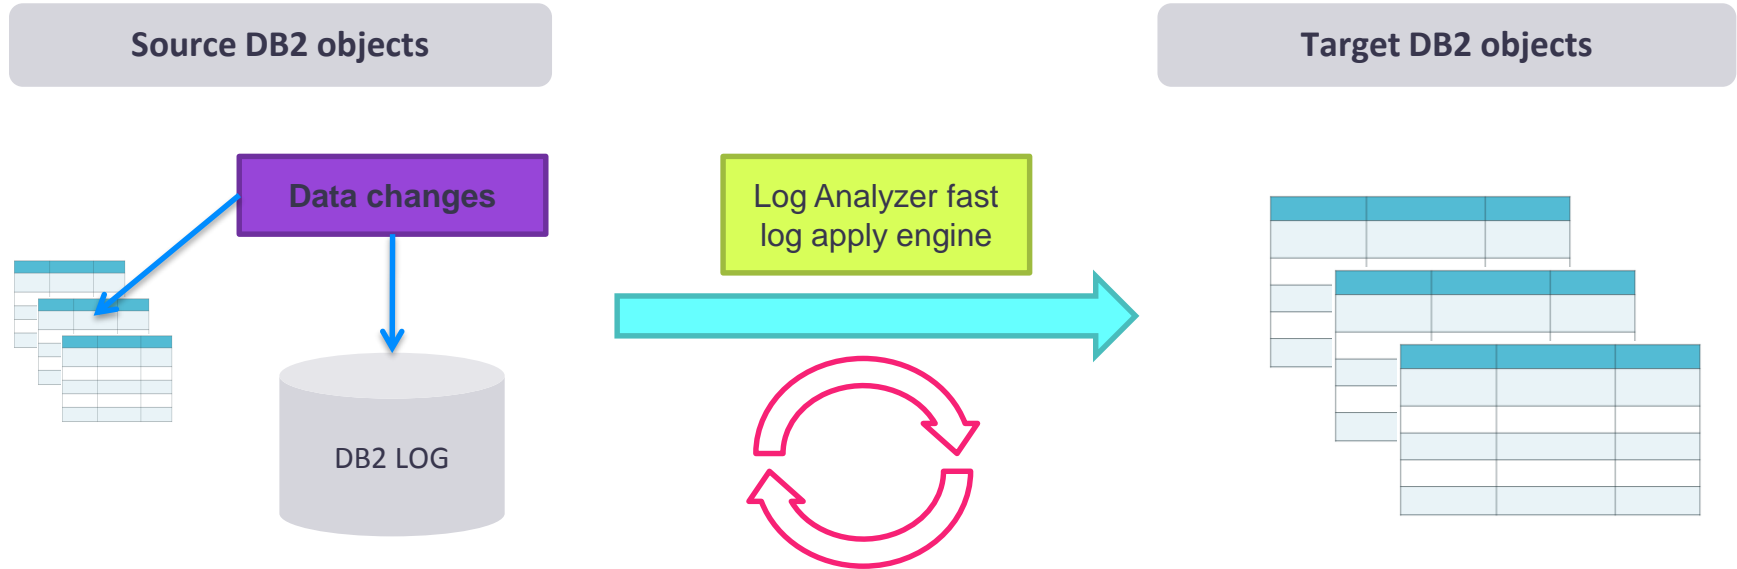
# Replicate data



# Replicate data

- DBA is required to replicate the production data into a test or shadow objects; to duplicate the data or to create a test environment similar to production
- CA Log Analyzer offers several options to propagate changes from the log of a source subsystem to objects in a target subsystem.
  - Using the strategy ensures the data is consistent
  - There is no outage on the source objects as the data comes from the DB2 log

# Replicate data



<https://youtu.be/HhNvNGSP7qc>

# Replicate data using REDO SQL

- You can generate REDO SQL and run it against a target subsystem to propagate changes
  - You see the actual SQL statement submitted to target subsystem
  - Limited options to rename objects to a target subsystem
  - You do not have control over the SQLCODEs processing
  - More suited to one time requests
- Or use Load Format report and Apply utility
  - Load file containing data records for each update, delete and insert (only data, no SQL statements)
  - Control file that describes the format of the records in the load file
  - Reads records in data file and generates dynamic SQL statements to apply the changes on a target subsystem

# Replicate data using LOAD FORMAT and APPLY

- Main advantages of Apply Utility
  - Easy source-to-target mapping of objects
  - Apply engine is much faster, more efficient and consumes less CPU compared to generated REDO SQL file
  - You have control over SQLCODEs generated during the processing
    - Options to CONTINUE/TERMINATE/WARNING on a particular SQL code
  - Ultra fast INSERT only processing
    - Alternative to LOAD RESUME SHARELEVEL CHANGE

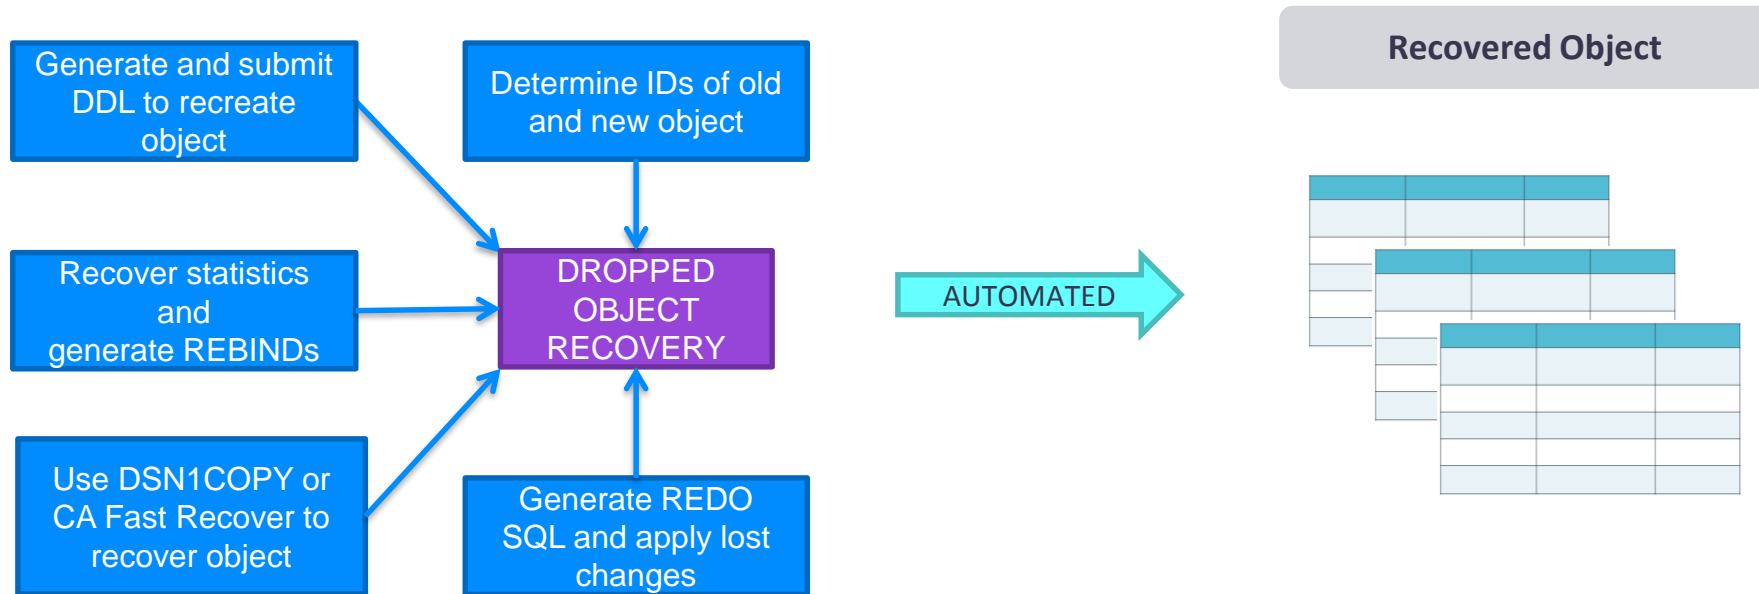
# Dropped object Recovery

---

# Dropped Object Recovery

- As a DBA you may be sometimes facing the situation where one or more objects in your critical database were accidentally dropped. Sometimes even restrict on drop does not help to prevent this very bad situation.
- It is critical to get the object recovered as soon as possible.
- CA Log Analyzer offers an automated and easy-to-use way to recover dropped object to its original state.

# Dropped Object Recovery



# Dropped Object Recovery

- Various miscellaneous options saved in your ISPF profile allows you to adjust the dropped object recovery process to your needs
  - Chose used utilities vendor
  - Whether to rebuild indexes
  - Include rebinds
  - Recover statistics/ generate runstats
  - Take image copy after recovery
  - Delete/keep temporary datasets used during the recovery process

<https://www.youtube.com/watch?v=nSxMbTEOWEs>

# Other reports

A thick, solid pink horizontal bar that underlines the text "Other reports". It starts at the left edge of the text and extends to the right edge of the text.

# Quiet Point Analysis report

- Locates ranges in the log with no data changing activity for a given list of objects.
- Such ranges are called “quiet” and the lowest RBA of the selected quiet range can be registered in SYSIBM.SYSCOPY as a true recovery point.
- Quiet points are alternative to QUIESCE where any outage caused by QUIESCE is not acceptable.
- <https://www.youtube.com/watch?v=M0q2OINa4cM>

# Change Data Capture

- Lets you determine how enabling or disabling DATA CAPTURE CHANGES will increase or decrease the amount of data written to your log data sets
  - You can also monitor when a table's data capture attribute has been changed and what was the impact on logging
- When DATA CAPTURE CHANGES is used, the log records contains enough information to build a full before and after image, so no image copy is needed and the reporting is much faster.

# Other reports

- Rollback Activity
  - Allows you to monitor the frequency and number of rollbacks that occur on the DB2 subsystem and who is doing updates that cause rollbacks
- Image Copy Frequency
  - Provides you with the monitoring of the frequency with which Image Copies are made in relation to the activity over the desired tablespace
- Commit Frequency
  - Allows you to monitor the frequency with which your applications are performing commits which is helpful to verify no long-running applications are not committing

# **CA Fast Recover**

---

# CA Fast Recover Overview

- CA Fast Recover
  - CA Fast Recover restores database objects from image copies and DB2 log
- Benefits
  - Rich database object specification capabilities (wildcarding)
  - Intelligent multi-tasking to maximize performance
  - Creation of image copies of the recovered objects with CA Quick Copy
  - Automatic index rebuilding
  - Three recovery engines
  - OBID translation for fast data migration and recovery of dropped objects
  - Cooperation with IBM RECOVER utility
  - Preview of the recovery process and verbose messages
  - Log Labels to identify log points
  - RBA format conversion
  - Ability to recover from BACKUP SYSTEM



Watch video



Watch video

# CA Fast Recover - Types of Recoveries Supported

- To current recoveries
  - Objects are recovered to the most recent state
  - Used when objects are lost completely
  - You can specify alternate VSAM file
- Index Rebuilds
  - Re-create indexes by reading and extracting data from tablespaces
- Point-in-Time recoveries
  - Objects are recovered to a given point-in-time in the past
  - Recovery from application errors
  - You can choose a copy which is not present In SYSCOPY (FULLDD)
- BACKOUT Recoveries
  - Undoing changes to given RBA/LRSN
- Forward Recoveries
  - Applying logs from given RBA/LRSN to recover from non-DB2 backups or for special purpose recoveries

# CA Fast Recover Wildcarding

- CA Fast Recover allows you to process multiple objects easily
- Wildcarding characters
  - `COPY TABLESPACE DALPHA?.%`
- LISTDEF (standard DB2 mean for wildcarding)
  - `LISTDEF MYLIST INCLUDE TABLESPACE DALPHA.%`  
`COPY LIST MYLIST`
- Special wildcarding keywords
  - `INDEX (ALL)` recovers all indexes
- DEFINE NO objects are automatically skipped when wildcarding

# CA Fast Recover

## Multi-tasking

- CA Fast Recover can process multiple objects or partitions in parallel to maximize performance
- Multi-tasking is applicable when
  - Single partitioned object is recovered
  - Multiple objects are recovered
  - Rebuilding multiple indexes
- Multi-tasking intelligence
  - CA Fast Recover automatically restricts multi-tasking to prevent multiple tasks competing for one resource, such as
    - Two or more tasks competing for one input image copy
    - Two or more tasks competing for one tape unit
- Keyword
  - MAXTASKS n

# CA Fast Recover

## Automatic Index Rebuilding

- CA Fast Recover can automatically rebuild indexes that have been defined with COPY NO option (COPY NO indexes)
  - You cannot create a copy of a COPY NO index
  - The only way to restore such index is to run REBUILD
  - Attempt to recover a COPY NO index results in an error message if the automatic index rebuilding is switched off
- Useful when wildcarding or when recovering mixture of COPY NO and COPY YES indexes
  - `RECOVER INDEX(ALL) TABLESPACE DBETA.MANYIXES`
- PARMLIB option
  - `AUTO-IX-REBUILD(YES)`

# CA Fast Recover - Recovery Engines

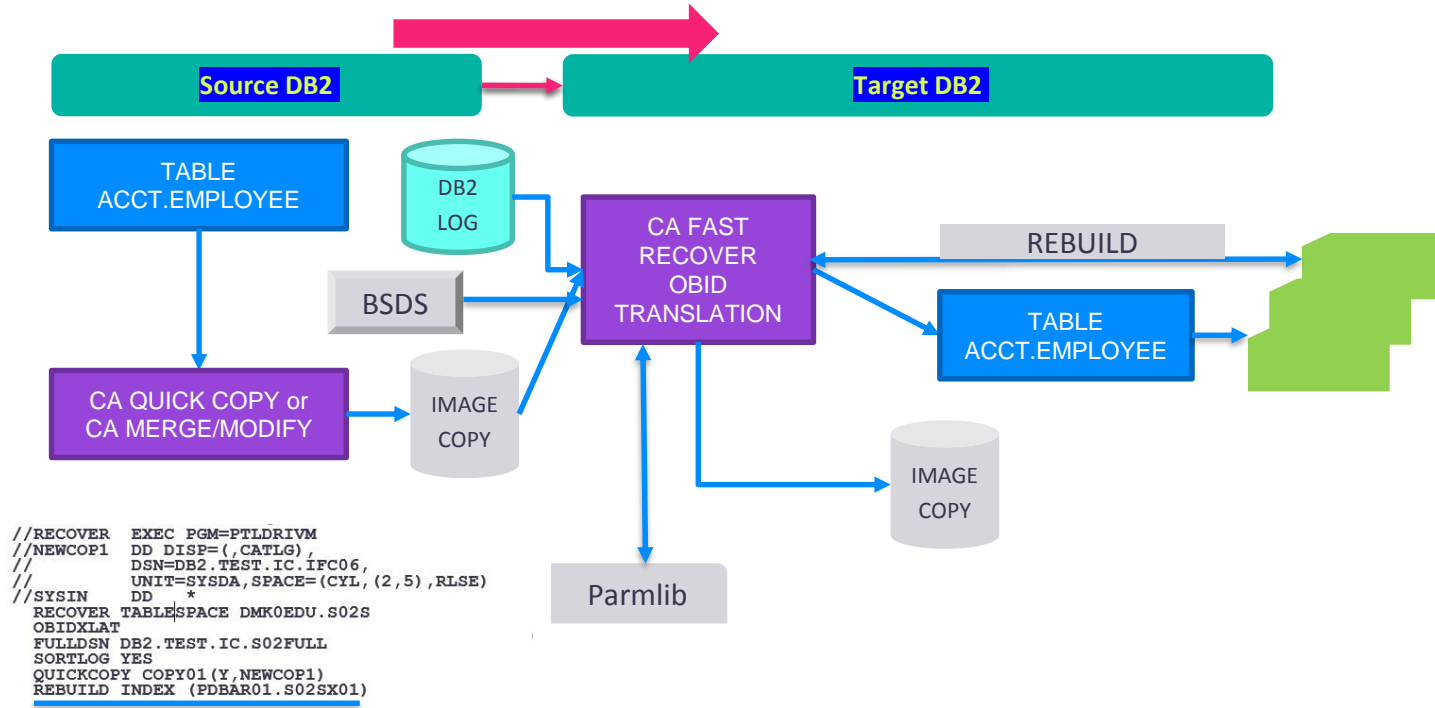
- CA Fast Recover offers three recovery engines
  - **SORTLOG NO** engine for standard recoveries
    - Standard 5-phase approach to the recovery process
    - Recovers tablespaces and LOBs from standard image copies
    - Rebuilds indexes
  - **SORTLOG YES** engine for efficient fully featured recoveries that include both rebuilding indexes and creating copies
    - This engine sorts log records in order to recover tablespaces, rebuild indexes, and copy the recovered objects in a single pass
    - Recovers tablespaces, indexspaces, and LOBs from standard image copies
  - **SORTLOG CONCURRENT** for recovery from SNAPSHOT copies
    - This engine sorts log records in order to recover tablespaces from SNAPSHOT copies, rebuild indexes, and copy the recovered objects in a single pass
    - Recovers tablespaces and LOBs from SNAPSHOT image copies

# CA Fast Recover

## OBID Translation (OBIDXLAT)

| PROBLEM  | Solution  |
|--|---|
| <ul style="list-style-type: none"><li>• DBA is required to quickly migrate data from one DB2 system to another one, for example from development environment to test environment</li></ul> | <ul style="list-style-type: none"><li>▪ DBA Creates a backup of the data on source DB2 system for recovery purposes using either CA Quick Copy or CA Merge/Modify depending on the availability of the object required</li><li>▪ DBA reuses the backup also to migrate the data to the target DB2 system using OBID TRANSLATION function of CA Fast Recover</li></ul> |

# CA Fast Recover OBID Translation (OBIDXLAT)



# CA Fast Recover

## Preview of the Process and Verbose Messages

- Preview
  - Recovery means (image copies, log ranges etc.) that would have been used are displayed, nothing is really done.
  - PREVIEW YES
- Verbose messages
  - When ALLMSGS keyword is specified, CA Fast Recover provides detailed information on the recovery process. This includes (but is not limited to) the following:
    - Image copies processed
    - Log ranges processed
    - BSDS and DB2 log datasets used
    - Index rebuild statistics
    - Number of log records and pages processed
    - Image copies created
    - Catalog and object status updates
    - Multi-tasking

```
PFR0078I - IMAGE COPY 'PDBAR01.#K.D11A.MMTSIK06.AAA00.G0003V00'  
ON DISK, FILESEQNO = 0, VOLSER(S) = (CATLG),  
WITH DATE 18/05/23, STYPE ' ',  
TIMESTAMP 2018-05-23-10.14.55.400001R  
AND START RBA/LRSN OF X'000000000A6BB9B36188'  
WILL BE MERGED FOR TABLESPACE 'MMDBIK01.MMTSIK06'
```

```
PFR0156I - LOG SORT TASK 1 INFORMATION:  
ESTIMATED AVERAGE RECORD LENGTH.....1,837  
ESTIMATED NUMBER OF RECORDS.....2,222  
SORT MESSAGE DDNAME.....SL01MSG  
SORT WORK DDNAME(S).....SL01WKNN  
NUMBER OF DYNAMICALLY ALLOCATED WORK DDS.....3
```

```
PFR0113I - RECOVER STATEMENT PREVIEW COMPLETE. HIGHEST RETURN CODE = 0
```

# CA Fast Recover

## Use Log Labels to Identify Log Points

- Specify Log points for point-in-time recoveries with convenience
- The Log Label standalone utility and CA Recovery Analyzer can create log labels that identify log points with human-readable names
- Use TOLOGLABEL keyword to identify log point with previously created log label

```
RECOVER TABLESPACE  
DBANK.CPCRD  
TOLOGPOINT  
X'CADE014589781A80331A'
```

```
RECOVER TABLESPACE  
DBANK.CPCRD  
TOLOGLABEL  
CREDITCARDS.NOONLOGPOINT
```

### Example: Create a Log Label

This example creates the MYLOGLABEL log label at the time when the job is submitted:

```
.CALL LOG  
.DATA  
SSID = (DB01)  
END = (CURRENT TIME)  
LOGLABEL = (MYLOGLABEL)
```

# **CA Merge/Modify**

# Content

1

CA LOG ANALYZER OVERVIEW

2

LOG ACCUMULATION

3

CHANGE ACCUMULATION

4

MODIFY

5

MERGECOPY

# CA Merge/Modify Overview

# Using of CA Merge/Modify for Db2 for z/OS



## Create up-to-date image copies

Create up-to-date image copies without impacting live database objects by merging incremental copies with full image copies and applying changes from the archive and active DB2 logs.



## Create a copy of an image copy

This option increases the availability of tape drives during image copy processing, but does not compromise data security requirement for disaster recovery purposes.



## Manage and optimize your backup and recovery data

You can define granular criteria for deleting image copy records from SYSCOPY or can have the utility automatically delete unwanted recovery information from SYSCOPY and SYSLGRNX together with deleting the copy data sets.



## Create up-to-date full and incremental image copies

Create up-to-date full and incremental image copies by merging existing image copies.

# Log Accumulation

# Log Accumulation

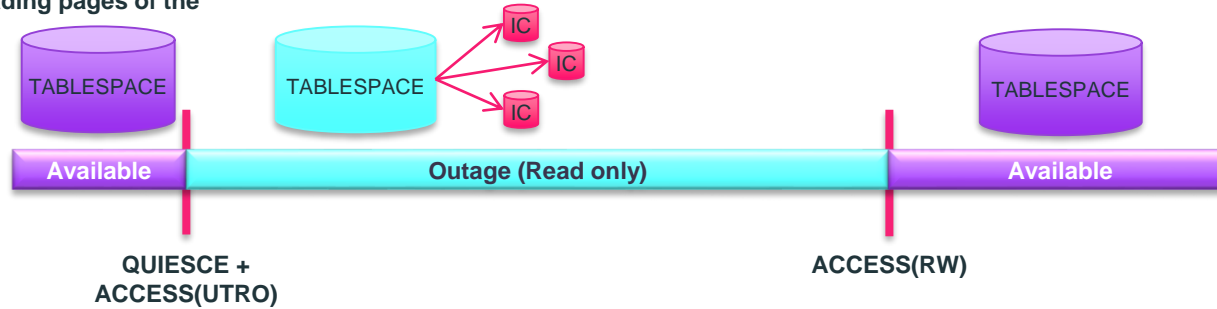
With **Log Accumulation component**, you can create consistent full image copies of table spaces, LOBs, and index spaces with minimum outage

- Useful when
  - Only minimum outage is tolerated for the purpose of creating backups
- Highlights
  - Minimum outage time (QUIESCE only) needed for consistent copy
  - Create up to 4 copies in one pass
  - Store copies directly on tapes
  - LISTDEF and dynamic allocation of the copies using TEMPLATE

# Log Accumulation Processing

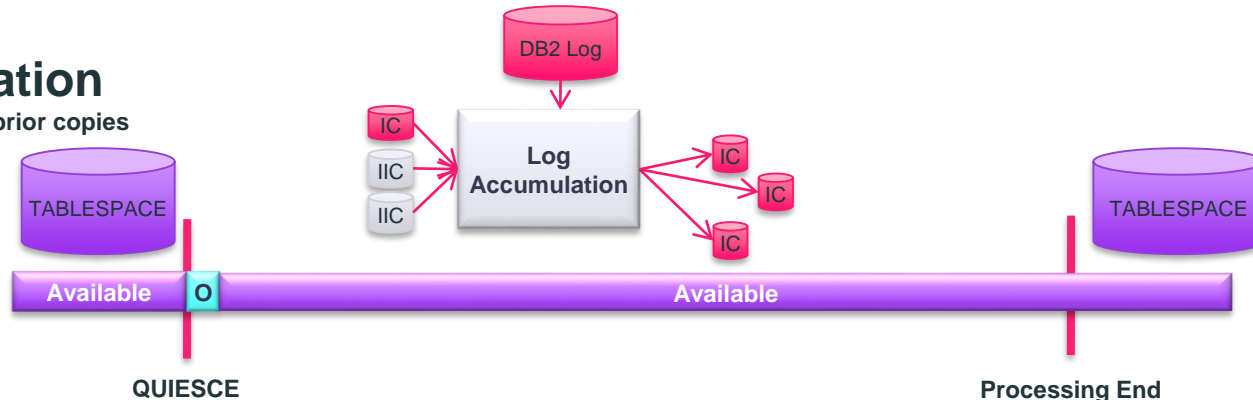
## COPY SHRLEVEL REFERENCE

Create copy by reading pages of the table space



## Log Accumulation

Create copy from prior copies and log



# Log Accumulation Sample JCL

- **Log Accumulation example**
  - Two table spaces
  - Image copies allocated dynamically with TEMPLATE

```
//ACCCOP    EXEC  PROC=CA@UTIL  
//SYSIN     DD  *
```

```
TEMPLATE ICTEMPLT DSN '&PMFID.&DB..&TS..IC.&UNIQ.'
```

```
ACCUMULATE
```

```
TABLESPACE DBETA.BOXES
```

```
COPIES COPY01(Y,ICTEMPLT,,LP)
```

```
TABLESPACE DBETA.BRIDGES
```

```
COPIES COPY01(Y,ICTEMPLT,,LP)
```

```
QUIESCE YES
```

# Log Accumulation Pros and Limitations

- **Pros**

- Minimum outage, apart from QUIESCE, the object is fully available during the processing. It is also possible to run the log accumulation without QUIESCE - SHRLEVEL CHANGE image copy.
- Object is not touched during the processing
- Copies can be stored on tapes
- When there are no prior copies, new copy can be created solely from the DB2 log
- It is possible to specify TORBA/TOLRSN

- **Limitations**

- More CPU intensive than COPY
- Not all types of objects are supported (XML, HASH)
- Cannot process catalog and directory objects
- Cannot use prior SNAPSHOT/FLASH copies

# Change Accumulation

# Change Accumulation

With **Change Accumulation** component, you can filter log records required to recover a group of database objects and store them in a single compact file called *change accumulation log data set*

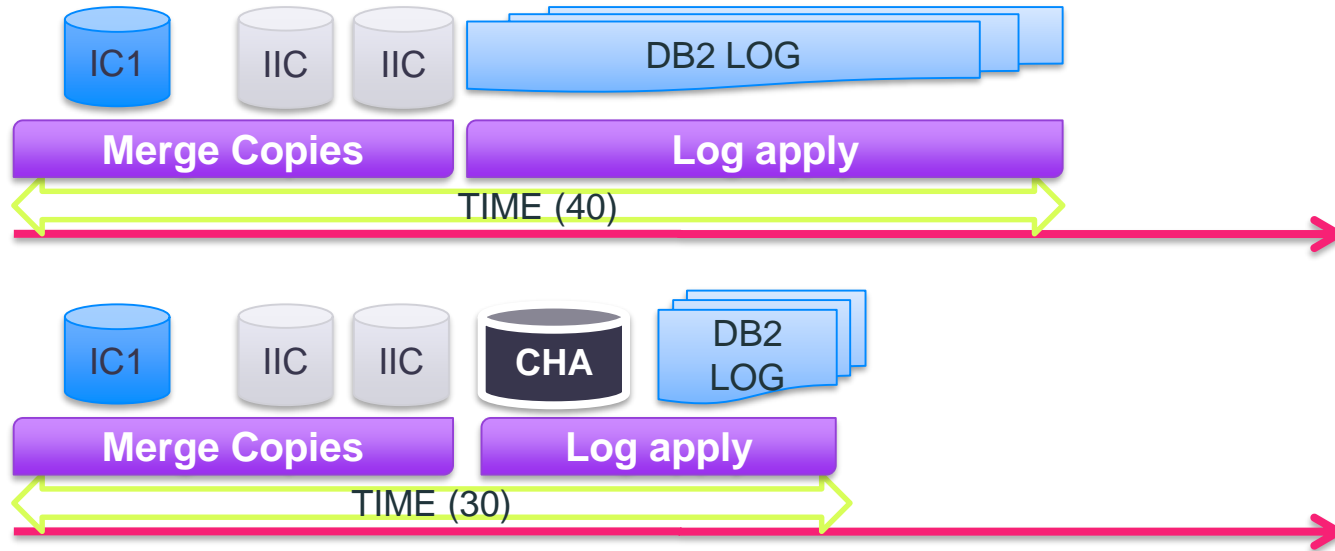
- Useful when
  - Number of log records processed by CA Fast Recover
- Benefits
  - Faster recovery (only CA Fast Recover)

# Change Accumulation – How it Works

- Find the last full image copy and its RBA/LRSN
- Determine current RBA/LRSN
- Read log records from the copy's RBA/LRSN up to the determined current RBA/LRSN
  - If the log record would be used to recover the object store it to the change accumulation log data set
- Register the change accumulation log data set to the PTI.ALOGFILE and PTI.ALOGRANGE tables

You can use ISPF interface to create a strategy to accumulate changes

# Change Accumulation – Faster Log Accumulation



# How to Use Change Accumulation with Log Accumulation and CA Fast Recover

When using the change accumulation log file with CA Fast Recover™, consider the following:

- You must have already accumulated the log records for the tablespaces being recovered, using **ACCUMULATE CHANGES**.
- Remember about two change accumulation recovery tables to track information about the change accumulation files. These tables must be located on the subsystem performing the recovery (ALOGRANGE, ALOGFILE)
- To perform the recovery with the records in the change accumulation file, specify **USE-CHANGE-ACCUM YES**. CA Fast Recover uses the correct change accumulation file for the objects being recovered. If the most recent changes have not been accumulated, CA Fast Recover accumulates them from the log data sets.

# MODIFY

# MODIFY - Highlights

- **Preview mode** that does not touch anything, but reports in detail what would have been changed (PREVIEW YES keyword)
- You can delete both SYSIBM.SYSCOPY **rows** and the image copy **data sets**
- Rich set of criteria that can be used to specify image copies to be deleted
- Multiple database objects can be processed with LISTDEF
- Syntax compatible with IBM
- You can not only delete SYSCOPY rows, but also **update** existing ones and **insert new** ones if needed

# MODIFY – Delete Obsolete SYSCOPY Rows

- In most cases MODIFY is executed regularly to delete obsolete rows in SYSIBM.SYSCOPY and from SYSIBM.SYSLGRNX. By deleting obsolete recovery-related data, recovery process becomes faster
- When deleting obsolete rows, you can specify either what rows will be deleted or what rows will be kept  
**[WHERE {sql-where-clause}]**
- Rows related to the associated indexes are deleted too
- If all copies are deleted, MODIFY will set the COPY PENDING status flag on affected database objects

# MODIFY – Syntax (continued)

- You can insert or modify rows in the SYSIBM.SYSCOPY table
  - RECOVER does not allow to recover from an image copy that is not registered in the SYSIBM.SYSCOPY table. In such case, you can use MODIFY to INSERT a row
  - You can use UPDATE to make changes of SYSIBM.SYSCOPY records if needed

## Insert Options:

```
[INSERT SET {column data} [USING ROW WHERE clause]]
```

## Update Options:

```
[UPDATE SET {column data WHERE clause}]
```

# MODIFY – SYSLGRNX

- When SYSIBM.SYSCOPY rows are deleted, certain SYSIBM.SYSLGRNX rows can become obsolete
- The SYSIBM.SYSLGRNX rows are not automatically deleted, but are staged for later deletion
- To delete staged SYSIBM.SYSLGRNX rows, specify the following syntax:

```
MODIFY SYSLGRNX [PREVIEW [NO|YES|SELECT|VERIFY] ]
```

# **SYSLGRNX deletion process – new approach based on calling IBM MODIFY**

- new approach based on calling IBM MODIFY
- MODIFY RECOVERY DELETE is executed -> all assigned records are deleted at the end of the process
- Output from the IBM MODIFY utility is displayed in the PTIIMSG dataset
- No need to run MODIFY SYSLGRNX
- PMM PARMLIB option -> IBM-LGRNX(YES)

# OR-DSNOTFOUND keyword

- Deletes rows from SYSIBM.SYSCOPY that reference cataloged image copy data sets that do not exist.
- OR-DSNOTFOUND option does not affect SYSLGRNX staging or deletion.
- OR-DSNOTFOUND can be used with DELETE AGE (see example)

This example deletes the image copy rows where the cataloged image copies are not found in the ICF catalog.

```
MODIFY RECOVERY TABLESPACE DALPHA. %  
DELETE AGE(22) OR-DSNOTFOUND
```

# MERGECOPY

# MERGECOPY

With **MERGECOPY** component, you can consolidate your incremental image copies and copy existing image copies

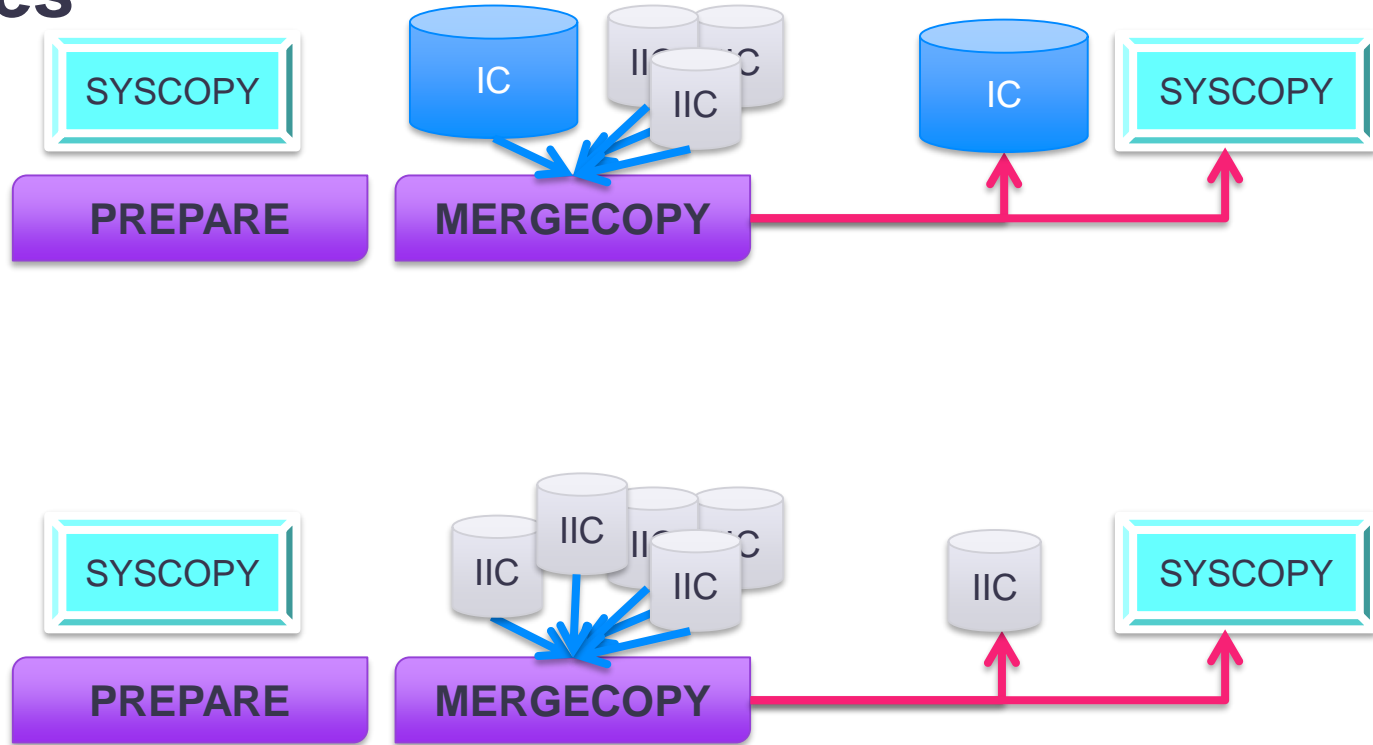
- Useful when
  - Outage required to recover or backup objects must be minimal
- Highlights
  - Database objects are not touched
  - CA Quick Copy fast I/O routines are used to process the image copies

- Highlights
  - By consolidating incremental image copies, you can speed up the recovery process as RECOVER utility will have fewer copies to process
  - By copying image copies, you can reduce outage time needed to create multiple backups of the database objects
  - You can also convert a snapshot copy to a standard image copy and store it on a tape

# MERGECOPY – Consolidate Incremental Image Copies

- In a recovery situation
  - RECOVER utility merges last full image copy with all subsequent incremental image copies. The database object is not available
  - If the copies are merged regularly by MERGECOPY, outage required to recover a database object can be reduced
- With MERGECOPY, you can
  - Consolidate multiple incremental image copies to one incremental image copy. In this case, RECOVER will have to process only two copies
  - Merge multiple incremental image copies to a new full image copy. RECOVER will only have to process one image copy

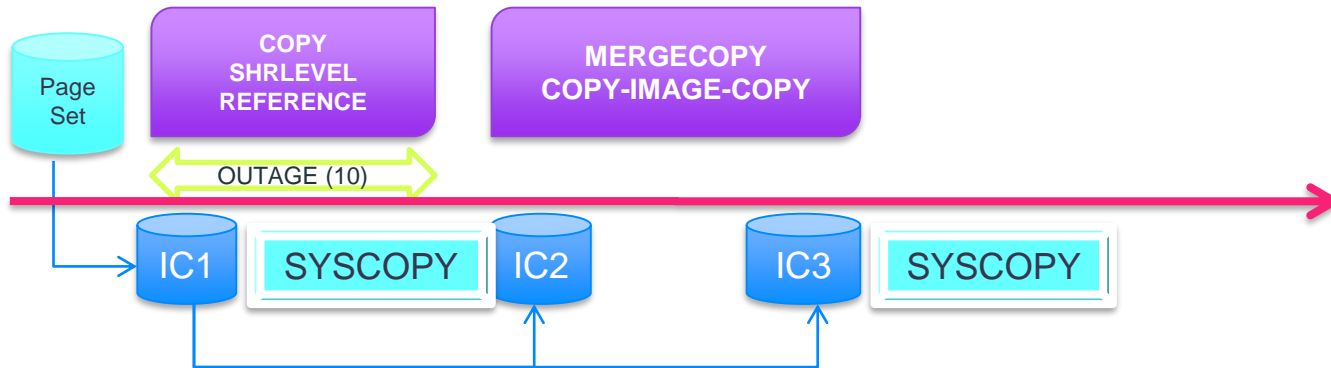
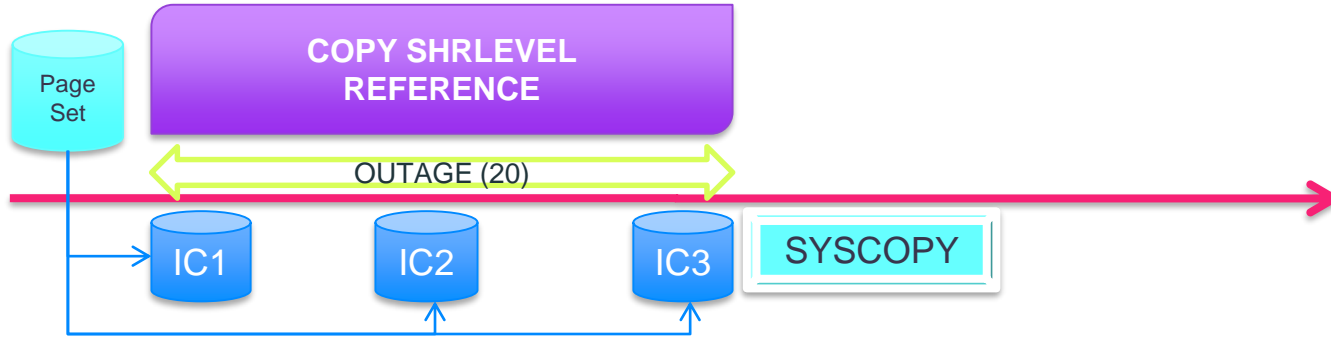
# MERGECOPY – Consolidate Incremental Image Copies



# MERGECOPY – Copy Image Copy

- If multiple image copies are needed for recovery
  - Run COPY utility and create multiple image copies. **The database object is in RO state for the time needed to create multiple copies.**
  - Run COPY utility to create one image copy. Then run MERGECOPY to create the rest of the copies. **The database object is in RO state for the time needed to create only one copy.**

**Note:** If you use CA Quick Copy for DB2 for z/OS to take SNAPSHOT point-in-time copies using the INSTANT option, you can use CA Merge/Modify™ for DB2 for z/OS to convert the INSTANT copy into a standard image copy that is registered in SYSIBM.SYSCOPY and that can be used in standard recoveries. This is done using the LAST-INSTANT parameter of the `COPY-IMAGE-COPY` keyword.



# **CA Quick Copy**

---

# CA Quick Copy Overview

- CA Quick Copy creates various types of image copies to backup database objects
- Benefits
  - Fast and zIIP-enabled copy engine that creates standard image copies
  - Creation of snapshot copies to minimize outage
  - Rich database object specification capabilities (wildcarding)
  - Two ways of dynamic allocation of image copy datasets
  - Control of object availability during the copy process
  - Intelligent multi-tasking to maximize performance
  - Auditing. Checking the object being copied for page integrity errors
  - Syntax compatible with IBM
  - Changelimit based on RTS tables
  - Preview of the copy process
  - Automatic switch to full image copy
  - Seamless integration with other CA Utilities
  - Other functions

# CA Quick Copy - Types of Copies

- Standard Full Image Copies
- Standard Incremental Image Copies
  - Spacemap-based true incremental copies –MODIFYBITS YES, not valid for SHRLEVEL CHANGE
  - Spacemap-based merged incremental copies - MODIFYBITS NO, MERGE-COPY YES
  - LRSN/RBA-based incremental copies – LRSNCOPY YES, START\_RBA (SYSCOPY)
    - If you specify [MODIFYBITS YES](#), CA Quick uses the SYSCOPY row from the most recent full or incremental copy
- SNAPSHOT Copies
  - Exploit point-in-time copy features of modern DASD hardware
  - CPiTS
- Concurrent Copies
  - Use DFSMSdss to create full image copies

# CA Quick Copy

## Fast and zIIP-enabled Copy Engine

- To create standard image copies, CA Quick Copy uses high performance EXCP NO copy engine
- Features and Benefits
  - Highly efficient producer-consumer-controller concept based copy engine that minimizes idle waiting for DASD or tape
  - Workload is offloaded to a zIIP processor if one is available
  - Engine uses high-performance synchronization of its tasks
  - High performance Media Manager interface is used to read the database objects
  - Up to 8 copies of a database object can be created simultaneously

# CA Quick Copy SNAPSHOT Copies

- CA Quick Copy can exploit point-in-time copy features of modern DASD hardware
- Snapshot copies for recovery purposes
  - Creation of **temporary SNAPSHOT** image copies that are automatically converted to standard image copies. This reduces outage time and allows to store the resulting copies to tapes.
  - Creation of **instant SNAPSHOT** image copies that are registered to the SYSCOPY table and can be used for recovery.
- Snapshot copies for non-recovery purposes
  - Creation of **unregistered consistent SNAPSHOT** image copies with full object availability. Such copies cannot be used for recovery but can be used for data migration. **This is limited to universal tablespaces.**

# CA Quick Copy Wildcarding

- CA Quick Copy allows you to process multiple objects easily
- Wildcarding characters
  - `COPY TABLESPACE DALPHA?.%`
- LISTDEF (standard DB2 mean for wildcarding)
  - `LISTDEF MYLIST INCLUDE TABLESPACE DALPHA.%  
COPY LIST MYLIST`
- Special Wildcarding keywords
  - `INDEX-ALL` (includes all associated indexes)
  - `EXCLUDE` (excludes database objects)
  - `CREATOR` (objects with given creator)
  - `SCOPE ALL|PENDING` (includes only objects in copy pending status)
- `DEFINE NO` and `COPY NO` objects are automatically skipped
- OLM (Predefined lists of objects stored permanently)

# Object List Manager - syntax

```
OLIST {CREATE [creator.]listname |  
      INCLUDE IN | EXCLUDE FROM [creator.]listname base-selection  
      [expansion-options] [type-restriction] [verbose|noverbose] |  
      LIST [creator.]listname [SELECTIONS | OBJECTS | ALL] [type-restriction] |  
      EXPORT [creator.]listname TO LISTDEF listdefname DD ddname |  
      DROP [creator.]listname}
```

## base-selection:

```
[DATABASE dbname [SCHEMA schema]]  
[INDEX [creator.]indexname [DATABASE dbname]]  
[INDEXSPACE [dbname.]indexspacename [SCHEMA schema]]  
[TABLE [creator.]tname [DATABASE dbname]]  
[TABLESPACE [database.]tablespacename [CREATOR creator]]  
[LISTDEF listdefname]
```

## expansion-options:

```
[PART WHOLE | ALL | 0 | n[:m]]  
[RI]  
WITH [INDEXSPACE [PARTED] [NONPARTED] [COPYYES | COPYNO]]  
[XML]  
[LOB]  
[ARCHIVE]  
[HISTORY]  
[ALL]
```

## type-restriction:

```
[TYPE [EXCLUDE] [BASE] [TABLESPACE | INDEXSPACE] [XML | LOB] [ARCHIVE | HISTORY | CLONE | NOCLONE]]
```

New enhancements



# CA Quick Copy

## Dynamic Allocation of Image Copy Datasets

- When copying hundreds of objects, creation of JCL DD statements can become difficult
- CA Quick Copy provides two ways of dynamic allocation
  - TEMPLATE
  - PQCPARM
- TEMPLATE
  - Standard DB2 mean for dynamic allocation of image copies
  - Recommended to be used when creating new JCLs
  - Required for LISTDEF
- PQCPARM
  - CA Quick Copy specific
  - Allows simple conversion of JCLs that use static allocation to dynamic allocation by just removing DD statements from JCL
  - Configure using PARMLIB member

# CA Quick Copy

## Control Availability of Objects

- CA Quick Copy allows you to control availability of the object during the copy process
- **Full availability** (SHRLEVEL CHANGE, slowest)
  - Object is available (RW). Good for 24x7 and catalog copies
  - Resulting copy is considered inconsistent (it may contain data that was not committed yet). Recovery to consistent state requires applying logs
- **Read only availability** (SHRLEVEL REFERENCE, fast)
  - Object is available for reading (RO) and also QUIESCED before the copy. Resulting copy is always consistent and sufficient to recover to consistent state
- **No availability** (SHRLEVEL NONE, fastest)
  - Object is not available and resulting copy is always consistent.

# CA Quick Copy

## Intelligent Multi-Tasking

- CA Quick Copy can process multiple objects or partitions in parallel
- Multi-tasking improves performance when data is stored on different DASDs or tapes
- Multi-tasking is applicable when
  - Single partitioned object is copied
  - Multiple partitioned or non-partitioned objects are copied
- Single partitioned object
  - Each copy subtask processes **subset of partitions**
- Multiple objects (or wildcarding)
  - Each subtask processes **all partitions** of given object (or non-partitioned object)

# CA Quick Copy

## Multi-Tasking Considerations

- The more subtasks, the more main storage is required
- Multi-tasking is automatically adjusted to prevent contentions (such as two subtasks competing for one tape unit)
- Number of tasks and number of partitions determines number of image copy datasets created and registered to the SYSCOPY table
- To precisely control how many partitions will be stored to one image copy dataset use the PARTS-PER-COPY keyword

# CA Quick Copy - Auditing

- CA Quick Copy can check the database object for malformed pages or other similar defects
- Every page processed is passed to the audit routine which performs page-type specific checks
- Auditing is not valid for SHRLEVEL CHANGE copies
- Syntax
  - `AUDIT {0, 1, 2}`
    - Header and trailer bytes
    - Broken page bit
    - I/O error bit
- Audit level determines thoroughness of the checks (and performance impact)

# CA Quick Copy

## Syntax Compatible with IBM

- CA Quick Copy understands most of the IBM COPY syntax, the existing JCLs and syntax can be migrated with minimum effort
- Example of syntax that both IBM COPY and CA Quick Copy will accept

```
TEMPLATE MYTEMP
    DSN TSDUMP.&DB..&TS..D&DATE..T&TIME..&ICTYPE.
    UNIT VTape RETPD 30 STACK YES
    DISP (NEW,CATLG,CATLG)

LISTDEF MYLIST INCLUDE TABLESPACE DALPHA.

COPY LIST MYLIST
    COPYDDN(MYTEMP)
    SHRLEVEL REFERENCE
    PARALLEL
    FULL YES
```

# CA Quick Copy

## Preview of the Process

- CA Quick Copy allows to preview the copy process
  - Syntax is checked for correctness
  - The copied objects are verified against DB2 catalog
  - No copy is performed
- Preview allows to quickly spot possible syntax errors and object statuses that may prevent copy creation

# CA Quick Copy

## Automatic Switch to Full Image Copy

- The `FULL-AUTO` keyword specifies whether to automatically convert an incremental image copy to a full image copy (if the activity against the tablespace warrants the conversion).
- The following warrants converting an incremental image copy to a full image copy:
  - A load was performed
  - A reorganization was performed
  - A utility was terminated
  - A TORBA or TOCOPY recovery was performed.
  - A full image copy does not exist for each of the copy types (LP, LB, RP, RB) for which an incremental copy is requested (since the last LOAD or REORG was performed).

# CA Quick Copy

## Seamless Integration with Other CA DB2 Utilities

CA Quick Copy can be used from other CA DB2 Utilities

- **CA Fast Load** and **CA Rapid Reorg** use CA Quick Copy to create copies after loading or reorganizing data
- **CA Fast Recover** uses CA Quick Copy to create image copy after recovery or during SORTLOG YES recovery
- **CA Merge/Modify** uses CA Quick Copy to copy or merge image copies or to create up-to-date new copy
- **CA Recover Analyzer** generates JCLs that execute CA Quick Copy
- **CA Fast Unload** can unload data from standard or SNAPSHOT copy previously created by CA Quick Copy

# CA Quick Copy

## Other Functions 1

- Change Limits
  - CA Quick Copy can automatically determine whether to create full, incremental, or no copy using percentage of pages changed since last registered image copy
  - Keyword `CHANGELIMIT (pct1,pct2)`
  - CA Quick Copy uses Real Time Statistic to make a decision
- Exit Routine
  - CA Quick Copy allows you to pass every copied block to a user-exit routine
  - The exit routine can access or even modify the block passed.
  - The block is written to the copy **after** the exit routine is called
  - Syntax
    - `COPY TABLESPACE DGAMMA.GARMENTS`  
`COPIES COPY01(Y,CP1,MYEXIT,LP)`

# CA Quick Copy

## Other Functions 2

- Object Ordering
  - CA Quick Copy can order the objects by their size
  - Larger objects are processed first
  - This prevents idle subtasks when multi-tasking and waiting for the last subtask to complete
  - Automatically disabled when there is no multi-tasking
  - Syntax: `OBJ-ORDER SIZE`
- Status Messages
  - For copy steps that process many large objects, CA Quick Copy allows to periodically display copy status messages.
  - This allows you to see that the utility is still alive
  - Syntax: `COPY-STATUS n` (display a message every n minutes)
  - The message `PQC0133I STATUS: 00000270540 4K BLOCKS PROCESSED FOR PQC3DSDB PQC3DSTS`

# CA Quick Copy

## Other Functions 3

- Alternative to COPYDDN and RECOVERYDDN

```
COPY LIST PRDACCTG01
COPIES COPY01(Y,TMPACT01,,LP)
      COPY02(Y,TMPACT02,,RP)
```

- (Register Y/N, DDNAME, exit routine, ICBACKUP)
- Copying DB2 Catalog
  - Special type of execution (DB2CATALOG is not a wildcard)
  - Template and COPYDDN required

```
TEMPLATE TEMPLP11
DSN(HLQ.&DB..&SN..L.D&JU..T&TI.)
DISP (NEW,CATLG,CATLG) STORCLAS SCWRKD VOLCNT 25
UNIT SYSDA

COPY DB2CATALOG
COPYDDN(TEMPLP11) SHRLEVEL CHANGE INDEX-ALL PREVIEW NO
```

# CA Quick Copy

## Other Functions 4

- Differential Backup
  - First job in the chain creates a FULL COPY
    - COPY TABLESPACE DALPHA.ASSETS  
FULL YES, MODIFYBITS YES
  - Subsequent jobs create differential copies
    - COPY TABLESPACE DALPHA.ASSETS  
FULL NO, MODIFYBITS NO, MERGE-COPY YES
  - Only 1 incremental copy is always kept in SYSIBM.SYSCOPY
- LRSN-Based Incremental Image Copy
  - For TRACKMOD NO objects
  - Specify LRSNCOPY YES in the syntax
  - LRSNCOPY processing is set automatically for TRACKMO NO objects

# CA Quick Copy

## Other Functions 5

- QUIESCE fallback for SHRLEVEL CHANGE copies

For SHRLEVEL CHANGE copies, this feature specifies whether QUIESCE is invoked for tablespaces or tablespace partitions when the START\_RBA cannot be determined. The objects stay available and the image copy is not created. To use this feature, specify PQC-QUIESCE-FALLBACK (NO) in the UTIL parmlib. If you do not set this parameter, the default is YES. (PTF SO01660)

Do not confuse this parameter with the PARMLIB option

PQC - Quiesce Before Copy.....> NO (No/Yes)

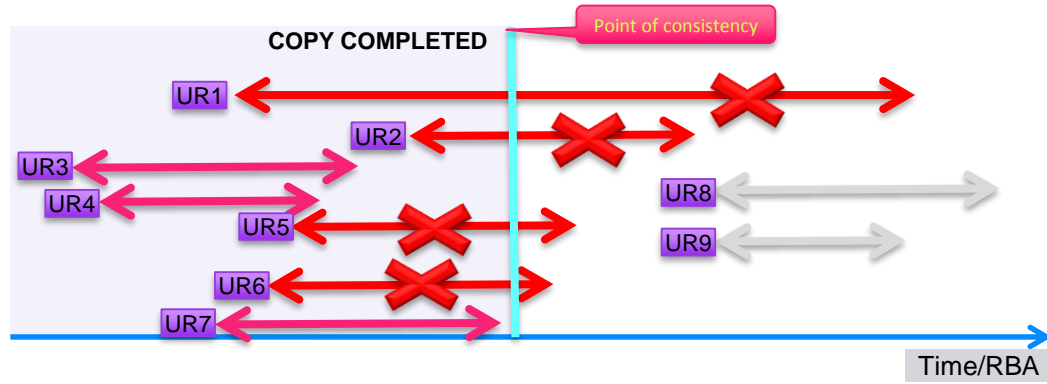
# Consistent Point in Time Service

# Overview

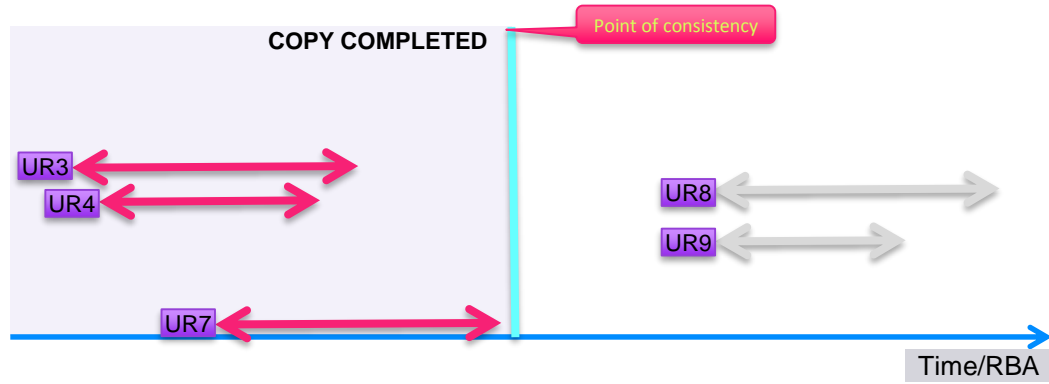


- **What CPITS does in a nutshell**
  - CPITS is a service that allows to create **consistent image copy for UTS** with full database object availability
  - General approach
    - Take an existing full SHRLEVEL CHANGE image copy and make it consistent using information from Db2 Log
  - Making the copy consistent
    - Revert all data changes that were performed by units of recovery (URs) that were not committed at given point of consistency
  - Main phases of the process
    - Open previously created SHRLEVEL CHANGE snapshot copy
    - Forward **log apply** to ensure that all updates are in the copy and not only in buffers
    - **Log Screening** to identify URs that were not committed
    - **Log undo** to revert data changes that were not committed

# URs



# URs



# CA Quick Copy JCL

- JCL Step

```
//COPY02 EXEC PT@UTIL  
//SYSIN DD *  
TEMPLATE T1  
DSN(SHK8.TEMP.MIG001.&DB..&TS..P&PART.)
```

```
COPY TABLESPACE DBANK.CPCRD  
SSCOPYDDN(T1)  
SNAPSHOTACCESS INSTANT  
SHRLEVEL CHANGE  
CONSISTENT  
SNAPSHOTIWAIT NO
```

```
COPY TABLESPACE DBANK.CPCRD2  
SSCOPYDDN(T1)  
SNAPSHOTACCESS INSTANT  
SHRLEVEL CHANGE  
CONSISTENT  
SNAPSHOTIWAIT NO
```

This creates a snapshot copy

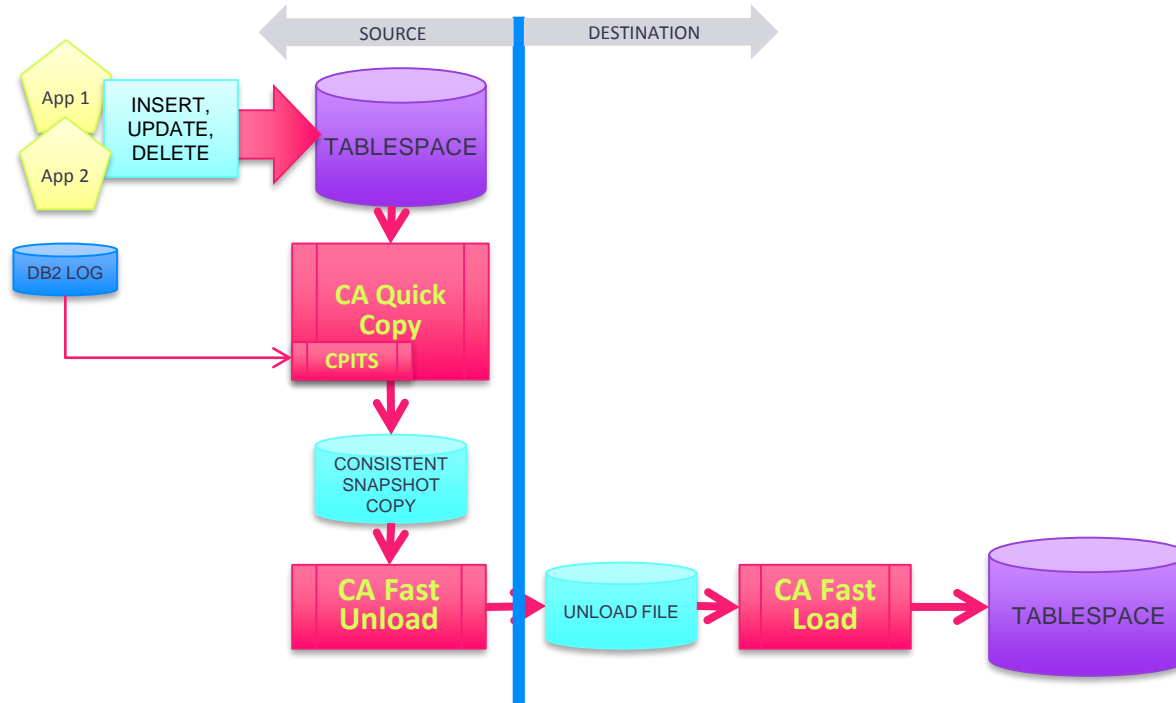
Object will be fully available

Consistent copy, call CPITS  
(new keyword)

Don't wait until the snapshot  
copy is complete

# Scenario

## Unloading and Migrating data



# CA Recovery Analyzer

---

# CA Recovery Analyzer Overview



- DB2 Recovery Automation
  - Automates recovery of single object or of a complete DB2 system
  - Just specify DB2 objects to be recovered. CA Recovery Analyzer checks recoverability of the objects and creates recovery JCLs automatically according to given parameters
- ISPF Interface
  - You prepare your recovery strategies using comfortable ISPF panels with online help always available
- Utilities
  - CA Recovery Analyzer allows you to use CA or IBM utilities to perform backup and recovery
- Customization
  - Specify various options to customize your recovery strategy



# CA Recovery Analyzer Overview



- Recovery Performance
  - Recovery analyzer can order recovered objects by RBA or by order the image copies occur on tapes to prevent check pending statuses or slow tape reads (Optimize Recovery)
  - Recovery analyzer can avoid recovery of objects that haven't been changed since the last image copy
  - Recovery analyzer can generate multiple recovery JCLs that can be submitted in parallel
- Profiles
  - Setup Recovery Analyzer global options for every user



# CA Recovery Analyzer

## Overview of Functions

### Strategy Services

Prepare in advance for recovery of a group of DB2 objects. Once a recovery strategy is created, it can be reused any time later. This is the core function of the Recovery Analyzer

### Quick Recovery

Quickly select table spaces or index spaces. Then automatically generate and submit recovery JCL. Suitable for ad hoc recoveries.

### System Recovery

Recover DB2 catalog, volumes, storage groups, boot strap datasets, and log datasets

### Disaster Recovery

Prepare for recovery and recover a DB2 subsystem in case of disaster

### Special Recovery Services

Display objects in restricted states, perform DSN1COPY translation (XLAT), recover broken pages and ranges, display and resolve in-doubt threads, create full pack backup of volumes

### Reporting Services

Analyze all your existing recovery strategies and generate reports. Report objects that are in multiple strategies, objects not present in any strategy, and objects eliminated from the reports

### Point-in-Time Recovery

Recover whole DB2 subsystem to a given point-in-time

# CA Recovery Analyzer - Strategies

Strategies are the base units of work for CA Recovery Analyzer. Most of the functions are applied on strategies

- Types of strategies
  - **Strategy services strategy** ([recovery strategy](#)) used to recover DB2 objects
  - **Disaster recovery strategy** for recoveries of DB2 subsystems from DR sites
  - **Reporting strategy** for reporting
  - **Point in time recovery strategy** for point in time recoveries of DB2 subsystems
- [Recovery strategies](#) are the most commonly used

# CA Recovery Analyzer – ISPF panels

The image displays four screenshots of the CA Recovery Analyzer ISPF panels, arranged in a 2x2 grid. Each panel is a window titled 'tpx.ca.com' with a menu bar (QWS3270, Edit, View, Options, Tools, Help) and a status bar at the bottom.

**Top Left Panel:** Shows the 'PRA Recovery Strategy Services' panel. It displays a strategy named 'DEM001' with creator 'KALMI05'. The strategy is described as 'DEM001 DEMO STRATEGY 01' and 'KALMI05.XTMP (DEM001)'. The status bar indicates 'Connected to tpx.ca.com port 23'.

**Top Right Panel:** Shows the 'PRA Options Selection Menu' panel. It lists various options for recovery, including 'Global Options', 'Copy/Recover/Check Options', 'Quick Copy Options', 'Fast Recover Options', 'Fast Load Options', 'FDR Options', 'Image Copy 'Before' DSN Qualifiers', and 'Image Copy 'After' DSN Qualifiers'. The status bar indicates 'Connected to tpx.ca.com port 23'.

**Bottom Left Panel:** Shows the 'Global Strategy Recovery Options' panel. It displays a list of options for recovery, including 'Utility Vendor', 'Include Indexes', 'Recovery Point', 'Recover via Logonly', 'Avoid Recovery', 'Optimize Recovery', 'Allow non-DB2 backups', 'Allow Full Pack Backups', 'Image Copy Before', 'Runstats', 'Quiesce', 'Check Utility', and 'Image Copy After'. The status bar indicates 'Connected to tpx.ca.com port 23'.

**Bottom Right Panel:** Shows the 'PRA Special Recovery Services' panel. It displays a list of options for recovery, including 'Display Restricted and Advisory Table / Index Spaces', 'DSN1COPY XLAT and Recover', 'Broken Page Recovery', 'Error Range Recovery', 'Scan ICF Catalog And Report on Datasets on a Volume', 'Display And Resolve Indoubt Threads', and 'Full Pack Backup Registration Processing'. The status bar indicates 'Connected to tpx.ca.com port 23'.

# Strategy Services

# CA Recovery Analyzer - Recovery Strategy

- List of Objects
  - A recovery strategy contains list of objects that are backed up and recovered as a group
  - Table spaces, indexes, LOBs, and indexes associated to table spaces
  - Object can be specified with names, wildcards, plans, and packages
- Options
  - **Global options** that determine what utilities will be used, whether associated indexes are included in the strategy, if image copies will be created before and/or after recovery, if a quiet point will be created, etc.
  - **Object-specific options** that can override the global options where appropriate
  - **Utility-specific options** that control behavior of the utilities

# CA Recovery Analyzer

## Recovery Strategy Operations

- Objects specified in the recovery strategy can be subject of the following operations
  - **Image copy before recovery** as a safety precaution in case recovery fails
  - **Quiet Point** (with QUIESCE) to establish point of consistency
  - **Recovery** to recovery objects from backups
  - **RUNSTATS** to collect statistics
  - **Check** to verify correctness of the recovery
  - **Image copy after recovery** to establish a new recovery base
- All operations are optional. **This allows to use Recovery Analyzer also only to create backup image copies if needed.**

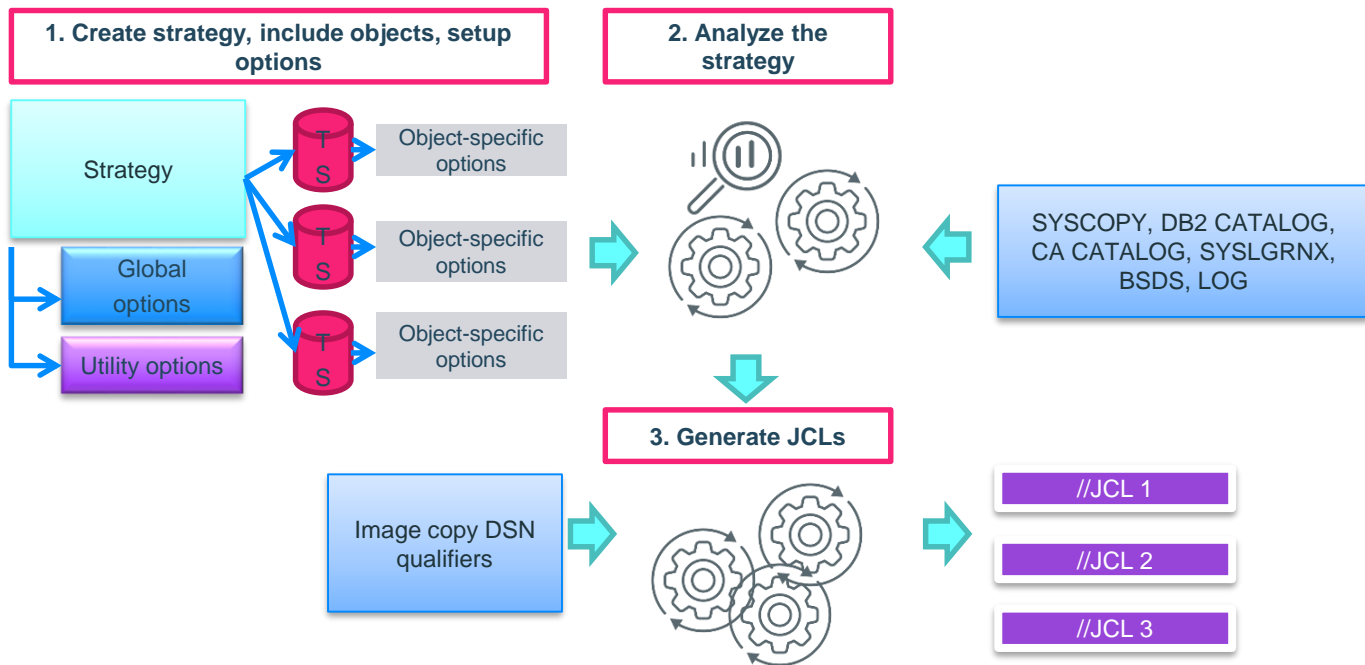
# CA Recovery Analyzer

## Recovery Strategy Workflow

1. Create a new recovery strategy
2. Include objects in the recovery strategy
3. Setup global recovery strategy options
4. Setup utility options
5. Setup object-specific options when appropriate
6. Analyze the recovery strategy to determine if there are any recovery impediments
7. Generate and submit JCLs to recover objects included in the strategy

# CA Recovery Analyzer

## Recovery Strategy Diagram



# CA Recovery Analyzer - Recovery Strategy

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRRSM 18.0 ----- PRA Recovery Strategy Services ----- 15/07/28 09:00
COMMAND ===== SCROLL ===== PAGE

Strategy ===== * Creator ===== KALMI05 ----- KALMI05 >
DB2 SSID ===== DB06 LOCATION ===== LOCAL VERSION ===== V10NF

0 Strategy Description Creator P O User Date Time
DEMO01 DEMO STRATEGY 01 KALMI05 N U <== STRATEGY CREATION
KALMI05 U KALMI05 15/07/28 07:16
----- BOTTOM OF DATA -----

Valid CMDs: For Strategies: (A,B,C,D,L,Q,U,X) For Versions: (B,D,E,S,V)
Press ENTER to process Press END (PF3/15) to quit

Connected to tpx.ca.com port 23 12/2 NUM IBM-3278-2-E - AS5T5277
```

## 1. Strategy creation

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRRSM 18.0 ----- PRA Strategy Recovery - Review Objects ----- 15/07/28 09:02
COMMAND ===== SCROLL ===== PAGE

STRATEGY ===== DEMO01 DESCRIPTION ===== DEMO STRATEGY 01
CREATOR ===== KALMI05 SHARE OPTION ===== U (U,Y,N)
SRC SSID ===== DB06 MULTIPLE JOB OPTION ===== N (P,Y,N)
REVIEW OPTIONS ===== N (Y,N) ----- KALMI05

Job Line Obj SpaceNm/ Database Prtn Ts/Ix Index Name / Optns
# Type Type Package Name Numb Creator Plan/Collection Set
001 S TS ALBUMS DALPHA PDBAR01 N
001 S TS ANTIQUES DALPHA PDBAR01 N
001 S TS ASSETS DALPHA PDBAR01 N
001 S TS BARS DBETA N
001 S TS BOXES DBETA N
001 S TS BRIDGES DBETA N
001 S IX XBRIDGES DBETA CRBETA XBRIDGES N
----- BOTTOM OF DATA -----

Valid Line Commands: A,B,D,E,IX,0,RI,RIA,TS,Z Valid Commands: Add,Header
(PF7/PF8) Scroll Up/Down Press END to save strategy

Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T5277
```

## 2. List of objects

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRPOPTS 18.0 ----- PRA Options Selection Menu ----- 15/07/28 09:07
OPTION =====

----- Strategy / Recovery Scenario Options ----- KALMI05

1) Global Options
2) Copy/Recover/Check Options
3) Quick Copy Options
4) Fast Recover Options
5) Fast Load Options
6) FDR Options
7) Fast Check Options
8) Profile Options
9) Image Copy 'Before' DSN Qualifiers
10) Image Copy 'After' DSN Qualifiers

Sets Recovery Scenario Options
Common to IBM and CA Utilities
Sets Common Utility Options
Options Specific to CA Quick Copy
Options Specific to CA Fast Recover
Options Specific to CA Fast Load
Options Specific to Innovation FDR
Options Specific to CA Fast Chk
User Dependent Options - Not Saved in Strategy

Enter Selection to Process Press END (PF3/15) to Cancel

Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T5277
```

## 3. Options

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRRSM 18.0 ----- PRA Fast Recover Options ----- Top of data
COMMAND ===== SCROLL ===== PAGE

Strategy ===== DEMO01 Creator ===== KALMI05 ----- KALMI05

Fast Recover Options:

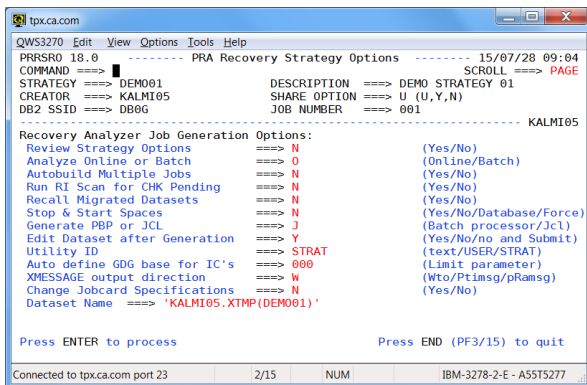
ABEND-TASK ===== N (Yes/No) More: +
ALLMSGs ===== N (Yes/No)
ALTERNATE ===== N (Yes/No)
ARCHIVE-LOG-EXCP ===== Y (Yes/No)
COPY-TYPE ===== (Local/Recovery)
DMS/OS ===== N (Yes/No)
EDITPROC ===== P (Problem/Supervisor)
ESTIMATED-KEYS ===== 0 (0 - 2,147,000,000)
ESTIMATED-RECS ===== 0 (0 - 2,147,000,000)
IDCAMS-CATALOG ===== Y (Yes/No)
LOG-BACKUP ===== N (Yes/No)
MAXTAPES ===== 05 (1 - 32)
MAXTASKS ===== 3 (1 - 6)

NOTE: Refer to the Fast Recover User Guide for an explanation of all options.
Press END (PF3/15) to Process SAVE to save defaults

Connected to tpx.ca.com port 23 2/17 NUM IBM-3278-2-E - AS5T5277
```

## 4. Fast Recover options

# CA Recovery Analyzer - Recovery Strategy



QWS3270 Edit View Options Tools Help  
PRRSR 18.0 PRA Recovery Strategy Options 15/07/28 09:04  
COMMAND ==> DEMO01 SCROLL ==> PAGE  
STRATEGY ==> DEMO01 DESCRIPTION ==> DEMO STRATEGY 01  
CREATOR ==> KALMI05 SHARE OPTION ==> U (U,Y,N)  
DB2 SSID ==> DB06 JOB NUMBER ==> 001  
----- KALMI05

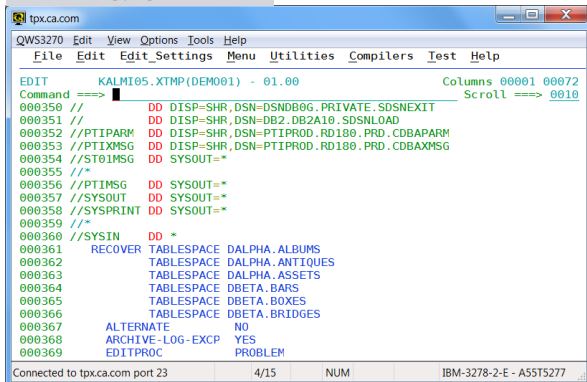
Recovery Analyzer Job Generation Options:

|                               |                            |                         |
|-------------------------------|----------------------------|-------------------------|
| Review Strategy Options       | ==> N                      | (Yes/No)                |
| Analyze Online or Batch       | ==> 0                      | (Online/Batch)          |
| Autobuild Multiple Jobs       | ==> N                      | (Yes/No)                |
| Run RI Scan for CHK Pending   | ==> N                      | (Yes/No)                |
| Recall Migrated Datasets      | ==> N                      | (Yes/No)                |
| Stop & Start Spaces           | ==> N                      | (Yes/No/Database/Force) |
| Generate PBP or JCL           | ==> J                      | (Batch processor/Jcl)   |
| Edit Dataset after Generation | ==> Y                      | (Yes/No/no and Submit)  |
| Utility ID                    | ==> STRAT                  | (text/USER/STRAT)       |
| Auto define GDG base for IC's | ==> 000                    | (Limit parameter)       |
| XMESSAGE output direction     | ==> W                      | (Wto/Ptmsg/pRmsg)       |
| Change Jobcard Specifications | ==> N                      | (Yes/No)                |
| Dataset Name                  | ==> 'KALMI05.XTMP(DEMO01)' |                         |

Press ENTER to process Press END (PF3/15) to quit

Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T5277

## 5. Strategy generation



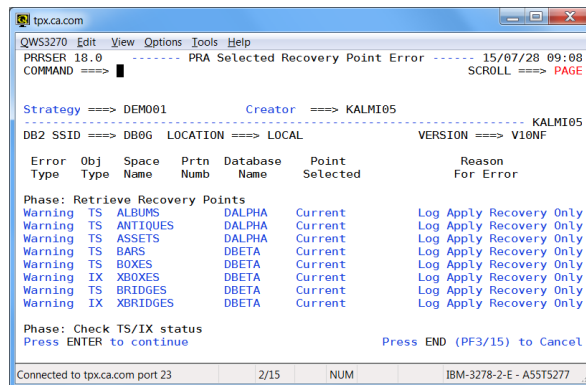
QWS3270 Edit View Options Tools Help  
File Edit Edit Settings Menu Utilities Compilers Test Help

EDIT KALMI05.XTMP(DEMO01) - 01.00 Columns 00001 00072  
Command ==> Scroll ==> 0010

```
000350 // DD DISP=SHR,DSN=DSNDB06.PRIVATE.SDSNEXIT
000351 // DD DISP=SHR,DSN=DB2,DB2A10.SDSNLOAD
000352 //PTIPARM DD DISP=SHR,DSN=PTIPROD.RD100.PRD.CDBAPARM
000353 //PTIXMSG DD DISP=SHR,DSN=PTIPROD.RD100.PRD.CDBAXMSG
000354 //ST01MSG DD SYSOUT=*
000355 /*
000356 //PTIMSG DD SYSOUT=*
000357 //SYSOUT DD SYSOUT=*
000358 //SYSPRINT DD SYSOUT=*
000359 /*
000360 //SYSIN DD *
000361 RECOVER TABLESPACE DALPHA,ALBUMS
000362 TABLESPACE DALPHA,ANTIQUES
000363 TABLESPACE DALPHA,ASSETS
000364 TABLESPACE DBETA,BARS
000365 TABLESPACE DBETA,BOXES
000366 TABLESPACE DBETA,BRIDGES
000367 ALTERNATE NO
000368 ARCHIVE LOG-EXCP YES
000369 EDITPROC PROBLEM
```

Connected to tpx.ca.com port 23 4/15 NUM IBM-3278-2-E - AS5T5277

## 7. Generated JCL



QWS3270 Edit View Options Tools Help  
PRRSR 18.0 PRA Selected Recovery Point Error 15/07/28 09:08  
COMMAND ==> SCROLL ==> PAGE

Strategy ==> DEMO01 Creator ==> KALMI05  
----- KALMI05  
DB2 SSID ==> DB06 LOCATION ==> LOCAL VERSION ==> V10NF

| Error Type                      | Obj Type | Space Name | Prtn Num | Database Name | Point Selected | Reason For Error        |
|---------------------------------|----------|------------|----------|---------------|----------------|-------------------------|
| Phase: Retrieve Recovery Points |          |            |          |               |                |                         |
| Warning                         | TS       | ALBUMS     |          | DALPHA        | Current        | Log Apply Recovery Only |
| Warning                         | TS       | ANTIQUES   |          | DALPHA        | Current        | Log Apply Recovery Only |
| Warning                         | TS       | ASSETS     |          | DALPHA        | Current        | Log Apply Recovery Only |
| Warning                         | TS       | BARS       |          | DBETA         | Current        | Log Apply Recovery Only |
| Warning                         | TS       | BOXES      |          | DBETA         | Current        | Log Apply Recovery Only |
| Warning                         | IX       | XBOXES     |          | DBETA         | Current        | Log Apply Recovery Only |
| Warning                         | TS       | BRIDGES    |          | DBETA         | Current        | Log Apply Recovery Only |
| Warning                         | IX       | XBRIDGES   |          | DBETA         | Current        | Log Apply Recovery Only |
| Phase: Check TS/IX status       |          |            |          |               |                |                         |

Press ENTER to continue Press END (PF3/15) to Cancel

Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T5277

## 6. Result of analysis

# CA Recovery Analyzer

## Recovery Strategy Notes

- Multiple strategies can be selected for processing
- JCL generated from the recovery strategies are kept and can be submitted later if appropriate
- History of JCLs generated is tracked within the recovery strategy
- If many objects are included in the strategy, recovery analysis can be performed in batch
- Indexes associated to recovered table spaces can be easily included in a recovery strategy by setting a single option

# Quick Recovery

# CA Recovery Analyzer

## Quick Recovery

Quick Recovery provides comfort when performing ad hoc recoveries

The workflow is the following:

1. Select objects
2. Set options
3. Recovery analysis
4. Generate and submit JCLs

A quick recovery can be saved as a recovery strategy if needed

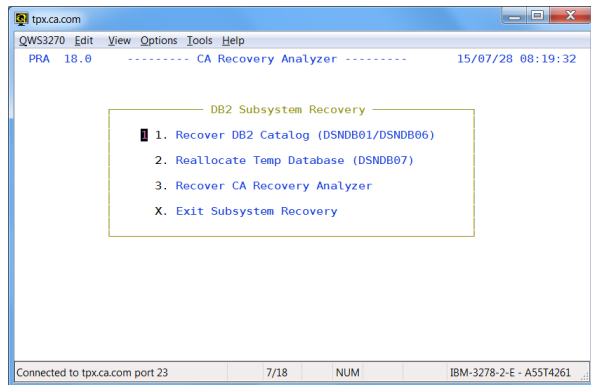
# System Recovery

# CA Recovery Analyzer

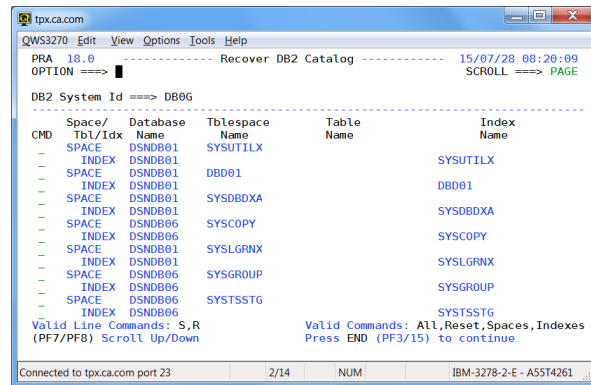
## Catalog Recovery

- With catalog recovery, you can easily recover single, multiple, or all DB2 catalog objects
- Workflow
  - Select DB2 Catalog Recovery
  - Select catalog objects to recover
  - Setup options (same as for the recovery strategy)
  - Review list of objects
  - Recovery scenario is analyzed to determine recovery impediments
  - Generate and submit JCLs

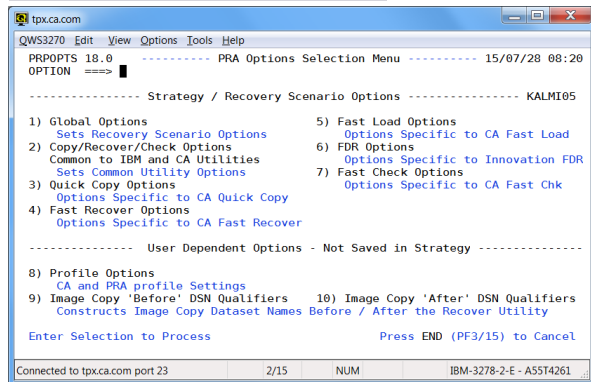
# CA Recovery Analyzer - Catalog Recovery



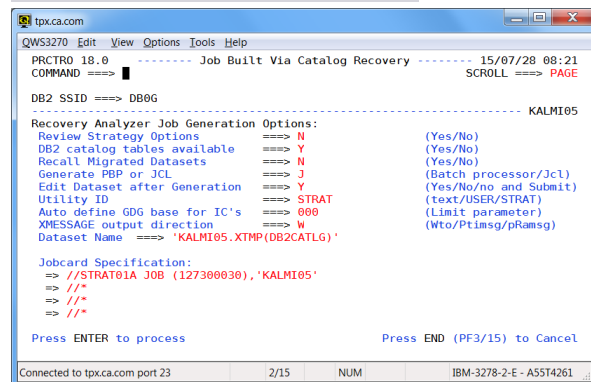
## 1. Select catalog recovery



## 2. List of catalog objects



## 3. Options



## 4. JCL generation

# CA Recovery Analyzer

## Recovery of Volumes and Storage Groups

With system recovery, you can do the following:

- **Recover DB2 objects residing on a particular volume or list of volumes.** Recovery Analyzer will scan the volumes for DB2 objects. Then you can specify the same options as for a recovery strategy and generate recovery JCL.
- **Recover DB2 objects that belong to a particular single storage group.** Recovery analyzer will list objects in the storage group. Then you can specify the same options as for a recovery strategy and generate JCL.

# CA Recovery Analyzer BSDS and Log Datasets

- **BSDS Recovery**
  - You can recover single (by copying from good one) or both BSDSs (by using archive log)
  - You can add/remove/modify records in BSDSs
  - You can compare two BSDSs
- **Log Recovery**
  - You can recover active or archive log datasets from good copies

# Disaster Recovery

# CA Recovery Analyzer

## Disaster Recovery Function Overview

- Disaster recovery functions cover
  - Preparation of backups for disaster recovery
  - Performing the disaster recovery
- **Prepare for disaster recovery** by creating a transportable backup package of that contains the following:
  - Copies of the DB2 subsystem libraries
  - Catalog, directory, BSDS, CA tables
  - Archive logs
  - Other libraries (QMF, DBRMLIB, ...)
- **Perform disaster recovery** by using the previously created transportable backup package

# CA Recovery Analyzer

## Disaster Recovery Method

- Recovery analyzer uses **Archive log method** for disaster recovery
- This method rebuilds the subsystem by using archive logs and image copies that are copied and shipped to the remote recovery site. This method minimizes data loss because you lose only those changes that were made after the RBA of the last archive log
- Disaster recovery backups and recovers DB2 system objects. User data backup is covered by the recovery strategies.

# CA Recovery Analyzer

## Preparing for Disaster Recovery

- Workflow
  - Create a new disaster recovery strategy
  - Specify backups of archive log and BSDS
  - Specify libraries to be backed up
  - Specify global options for the disaster recovery strategy
  - Specify image copy options
  - Generate backup and recovery JCLs
  - Submit the generated backup JCLs. The JCLs will perform the backup and also generate recovery JCLs including #INFO member that describes steps needed to recover

# CA Recovery Analyzer

## Disaster Recovery Backup JCLs

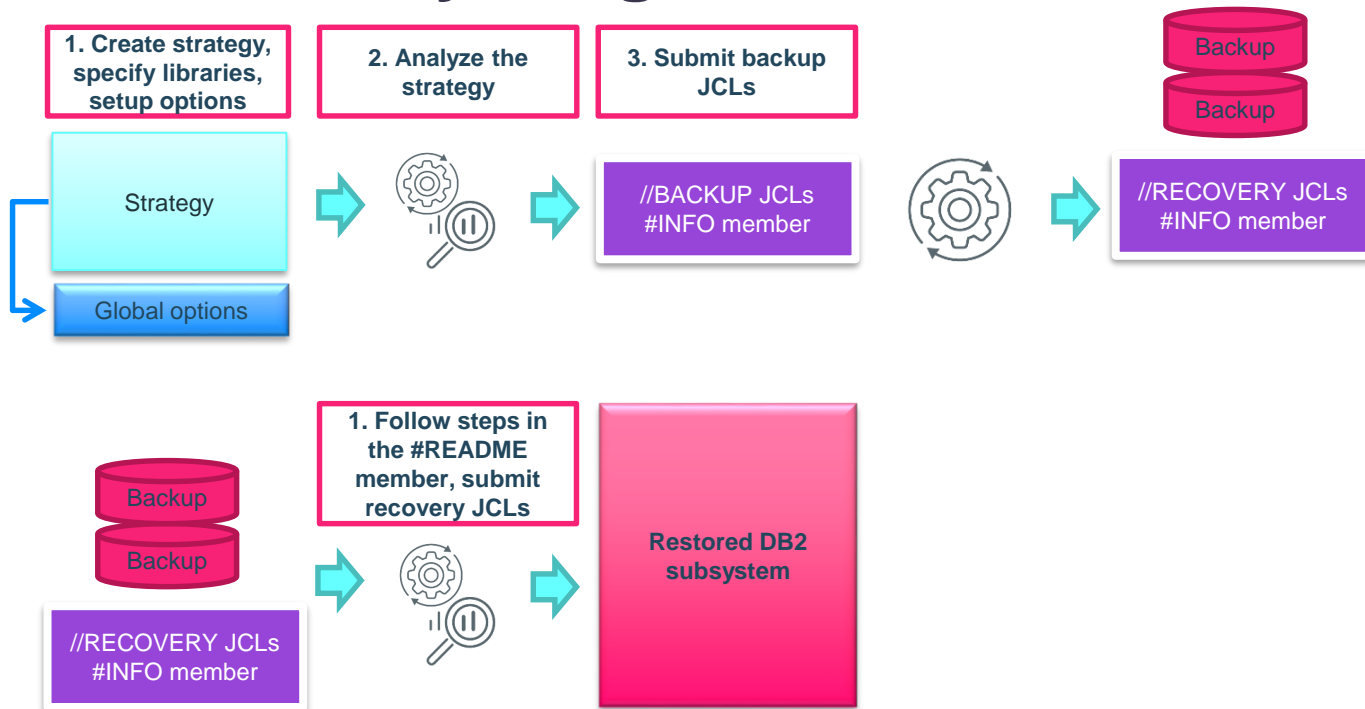
- **SSID@DNS** (Optional) Generates an SSID#DNS member that contains JCL to clean up the MVS catalog and remove old data sets at the recovery site.
- **SSID@DSN** Backs up the system files, copying any DB2 or system libraries that were specified in the strategy. This member also generates an SSID#DSN member to restore these files.
- **SSID@CDC** Backs up the DB2 catalog and archives the current active log. This member also generates the following recovery members: SSID#DBC, SSID#RDC, SSID#PRA, and SSID#ASG. These members build the DB2 catalog structure, restore the catalog contents from the backup, restore CA Recovery Analyzer, and realign the storage group definitions.
- **SSID@ARC** Backs up the bootstrap data sets and the archive logs. A data sharing environment has one SSID@ARC member for every member in the group. This member also generates the SSID#ARC and SSID#CAN members to restore these items.
- **SSID@JCL** Backs up the restore members and saves them to the same data set that contains the @ sign members.
- **SSID@TPL** (Optional) Creates a list of tapes that are needed for the recovery. The list cross-references the volumes and data set names for you. The list is saved to the data set that you specified on the Disaster Recovery Options panel.

# CA Recovery Analyzer

## Performing Disaster Recovery

- The steps needed are described in the #INFO member generated during preparation
- The steps in general involve the following
  - Submission of the previously generated recovery JCLs
  - Additional steps (ZPARM changes, DB2 commands)

# CA Recovery Analyzer Disaster Recovery Diagram



# CA Recovery Analyzer - Disaster Recovery

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRDSM 18.0 ----- PRA Disaster Recovery Services ----- 15/07/31 08:13
COMMAND ==> SCROLL ==> PAGE

Strategy ==> * Creator ==> KALMI05
DB2 SSID ==> DB06 LOCATION ==> LOCAL VERSION ==> V10NF

0 Strategy Description Creator User Last Update Date Time
- DISAS01 DR DEMO KALMI05 15/07/31 08:12
  'KALMI05.XTMP (#README)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB06@DNS)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB06@SDNS)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB06@CDC)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB06@ARC)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB2@ARC)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB3@ARC)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB4@ARC)' KALMI05 15/07/31 06:41
  'KALMI05.XTMP (DB06@JCL)' KALMI05 15/07/31 06:41
***** BOTTOM OF DATA *****
Valid CMDS: For Strategies: (A,D,L,U) For Members: (B,D,E,S,V)
Press ENTER to process Press END (PF3/15) to quit

Connected to tpx.ca.com port 23 19/2 NUM IBM-3278-2-E - AS5T4140
```

## 1. Strategy creation

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRDRGO 18.0 ----- PRA Disaster Recovery Options ----- 15/07/31 08:18
COMMAND ==> Strategy ==> DISAS01 Creator ==> KALMI05

Global Disaster Recovery Options:
Image copies used at DR site ==> RP (LP,LP,RP,RB)
Volume XREF/Skip DNS Member ==> NONE (NONE/mbr/Enter "S" for Selection)
Backup Utility ==> D (Dfmsdss/Fdr)
Backup Device Name ==> 3490 (3480,3490,Card,Tape,...)
Alias Name for Backup ==> BCKP (Eight Character Alias for Backup)
Nbr of Days back ==> 014 (1-999)
Recover Entire CA DB ==> N (Yes/No)
Tape pull list work file ==>
Archive Log Copy Options:
Use which archive logs at DR ==> C (Copied,pti Log,arch1,arch2,Erase)
Nbr of Logs to copy to disk ==> 004 with 00 copied in Active Logs
2 Archive Logs from 1 ==> N (Yes/No)
Retain Archives on 1 Tape ==> Y (Yes/No)
Utility Vendor ==> I (C(ca)/IBM)
Block Size ==> NONE (NONE,ZPARM,8K-28K,8192-28672)

Press END (PF3/15) to process SAVE to save defaults

Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T4140
```

## 3. Options

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
RPDRRV 18.0 ----- PRA Disaster Recovery - Review Objects ----- 15/07/31 08:17
COMMAND ==> Strategy ==> DISAS01 DESCRIPTION ==> DR DEMO
CREATOR ==> KALMI05 SHARE OPTION ==> U (U,Y,N)
SRC SSID ==> DB06 REVIEW GLOBAL OPTIONS ==> N (Y,N)

Cmd Dataset Name / Dataset Level Backup DSN (optional) Rmt Vol Rmt Alias
DB2G BSDS 1 DB2BACK.BKP.DB2SYS
DB2G BSDS 2
DB3G BSDS 1
DB3G BSDS 2
DB4G BSDS 1
DB4G BSDS 2
DB2 System Libraries DB2BACK.BKP.DB2SYS
DB2.DB2A10.GA.UI27713.SDSNLOAD VOL444
DB2.DB2A10.GA.SDSNLOAD VOL444
QMF Libraries DB2BACK.BKP.QMF
DB2.DB2A10.QMF VOL444
DBRM Libraries DB2BACK.BKP.DBRM
DB2.DB2A10.DBRM VOL444

Valid Commands: Insert,I#,Delete Press END to save strategy
(PF7/PF8) Scroll Up/Down

Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T4140
```

## 2. List of datasets

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRDRGO 18.0 ----- PRA Disaster Strategy Options ----- 15/07/31 08:20
COMMAND ==> Strategy ==> DISAS01 DESCRIPTION ==> DR DEMO
CREATOR ==> KALMI05 SHARE OPTION ==> U (U,Y,N)
DB2 SSID ==> DB06

Disaster Recovery Job Generation Options: PDS Member
Copy DB2 Catalog Tables ==> Y (Yes/No) DB06@CDC
Copy Archive Logs ==> Y (Yes/No) DB06@ARC
MVS Catalog Cleanup ==> Y (Yes/Scratch) DB06@DNS
Copy System Datasets ==> Y (Yes/No) DB06@SDNS
Utility ID ==> STRAT (text/USER/STRAT)
XMESSAGE output direction ==> W (Wto/Ptmsg/prmsg)
Dataset Name ==> 'KALMI05.XTMP'
Enter PDS with no member name or GDG Base with (+1)

Jobcard Specification:
=> //STRAT01A JOB (127300030),'KALMI05'
=> /*
=> /*
=> /*

Press ENTER to process Press END (PF3/15) to Cancel

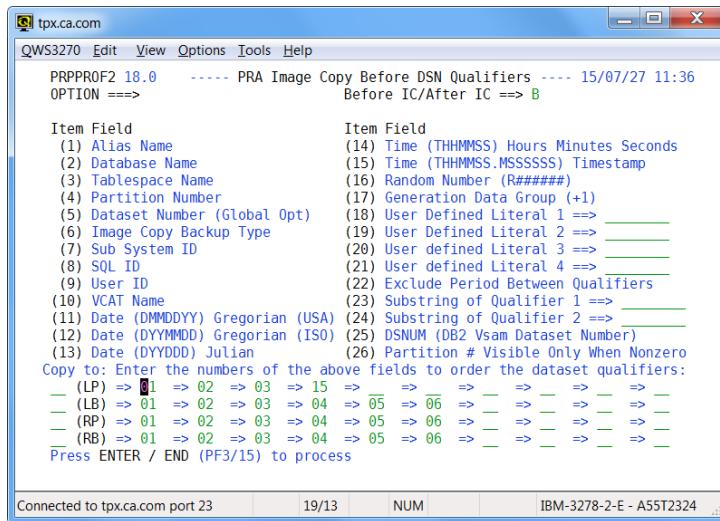
Connected to tpx.ca.com port 23 2/15 NUM IBM-3278-2-E - AS5T4140
```

## 4. JCL Creation

# General Topics

# CA Recovery Analyzer Image copies

- Image copies can be created with CA Quick Copy or with IBM COPY utilities
  - Dynamic allocation of the image copies is possible only with CA Quick Copy using PQCPARM dynamic allocation
- 
- Allocation using TEMPLATE is not possible, however Recovery Analyzer has elaborate system for creating DSN parameters of DD statements for static allocation



The screenshot shows the CA Recovery Analyzer interface with the following content:

```
tpx.ca.com
QWS3270 Edit View Options Tools Help
PRPPROF2 18.0 ----- PRA Image Copy Before DSN Qualifiers ---- 15/07/27 11:36
OPTION ===== Before IC/After IC ==> B

Item Field                                Item Field
(1) Alias Name                            (14) Time (THHMMSS) Hours Minutes Seconds
(2) Database Name                         (15) Time (THHMMSS.MSSSSSS) Timestamp
(3) Tablespace Name                       (16) Random Number (R#####)
(4) Partition Number                      (17) Generation Data Group (+1)
(5) Dataset Number (Global Opt)           (18) User Defined Literal 1 ==> _____
(6) Image Copy Backup Type                 (19) User Defined Literal 2 ==> _____
(7) Sub System ID                         (20) User defined Literal 3 ==> _____
(8) SQL ID                                (21) User defined Literal 4 ==> _____
(9) User ID                               (22) Exclude Period Between Qualifiers
(10) VCAT Name                            (23) Substring of Qualifier 1 ==> _____
(11) Date (DMMDDYY) Gregorian (USA)       (24) Substring of Qualifier 2 ==> _____
(12) Date (DYYMDD) Gregorian (ISO)        (25) DSNUM (DB2 Vsam Dataset Number)
(13) Date (DYYDDD) Julian                 (26) Partition # Visible Only When Nonzero

Copy to: Enter the numbers of the above fields to order the dataset qualifiers:
-- (LP) => 01 => 02 => 03 => 15 =>          =>          =>          =>          =>
-- (LB) => 01 => 02 => 03 => 04 => 05 => 06 =>          =>          =>          =>
-- (RP) => 01 => 02 => 03 => 04 => 05 => 06 =>          =>          =>          =>
-- (RB) => 01 => 02 => 03 => 04 => 05 => 06 =>          =>          =>          =>
Press ENTER / END (PF3/15) to process

Connected to tpx.ca.com port 23      19/13      NUM      IBM-3278-2-E - A55T2324
```

# CA Recovery Analyzer

## PRA#LOAD

CA Recovery Analyzer must have access to current system information to effectively perform recoveries for DB2 objects and volumes where the DB2 datasets reside

- PRA#LOAD is a batch job that retrieves and saves this information in product tables created for that purpose
- PRA#LOAD gathers object information from the DB2 catalog and dataset information from the MVS catalog or VTOC
- PRA#LOAD must be run for each DB2 subsystem where the CA Recovery analyzer will be used
- The PRA#LOAD job is created during the post-install customization step and placed in the CDBASAMP dataset
- The PRA#LOAD job should be scheduled for at least daily execution. In a highly dynamic environment where DB2 objects are created or dropped often, the schedule frequency should be higher



## **Jacek Rafalak**

Principal Product Owner

Jacek.Rafalak@ca.com



@cainc



in



# Thank You.