



# z/OS® 1.12 New Feature - VSAM CA Reclaim

Barbara McDonald  
IBM DFSMS Product Management  
bawhite@us.ibm.com

August 2, 2010  
Session 8054



**SHARE** in Boston

## Disclaimer

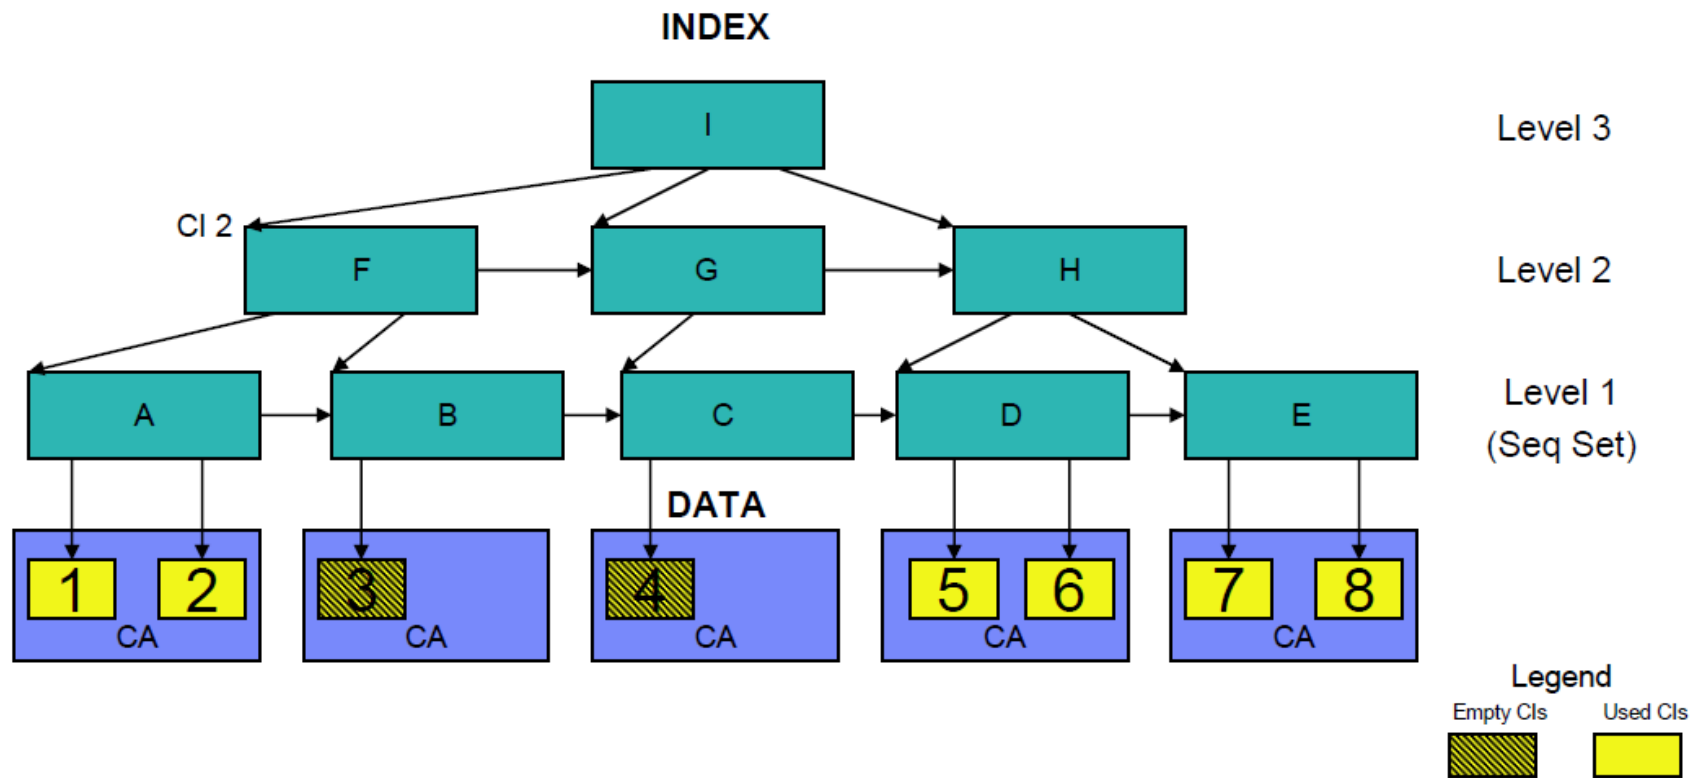
The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.

# VSAM CA Reclaim Overview

- **Problem**
  - VSAM KSDSs must be “reorganized” on a regular basis
    - Reclaim space previously used by deleted records
      - *Have the ability to Reclaim CIs, but not CAs*
    - ERASE’ing a range of keys can result in “orphaned” or “empty” CA’s if the new records inserted have higher key values
      - *Examples: Keys containing date and time stamps and keys generated using an increasing number such as a generation*
  - Over time “orphaned” CAs can cause performance issues
    - Fragmented space on volumes as data set grows and multiple volumes from extend (EOV)
    - Sequential processing degradation because sequence set must be read to determine the CA is empty
    - For direct processing, the number of index levels may grow larger than necessary, and result in more i/o's to the index component.
  - A “reorg” consists of
    - Closing the KSDS and taking it "offline" to the application
    - Back up the KSDS (if not DEFINEd as REUSE) to sequential data set using REPRO, EXPORT or some other method
    - Delete the KSDS and redefine
    - Reload the KSDS from the sequential copy using REPRO, IMPORT, or some other method

# VSAM CA Reclaim Overview

- Before CA Reclaim



# VSAM CA Reclaim Overview

- **Solution**

- Enhance VSAM and VSAM RLS to reclaim empty space (control areas) in KSDSs - applies to SMS and non-SMS data sets
  - Empty CAs are placed on a free list after the last record in the CA is erased
    - *List Anchor is in CI 2 of the index*
  - Reclaimed CAs may then be reused as new records are inserted anywhere in the data set
    - *CAs from free CA list are used first – rather than empty CAs beyond the high used RBA*
- CA Reclaim will not reclaim the following data CAs for a KSDS:
  - Partially empty CAs and CAs already empty before CA reclaim is enabled
  - CA with RBA 0 and the CA with the highest key of the KSDS
  - CAs in KSDSs with IMBED
  - Applications opening a data set with Global Shared Resources (GSR)
- Not a true online reorg, but it reduces the need to reorg the KSDSs

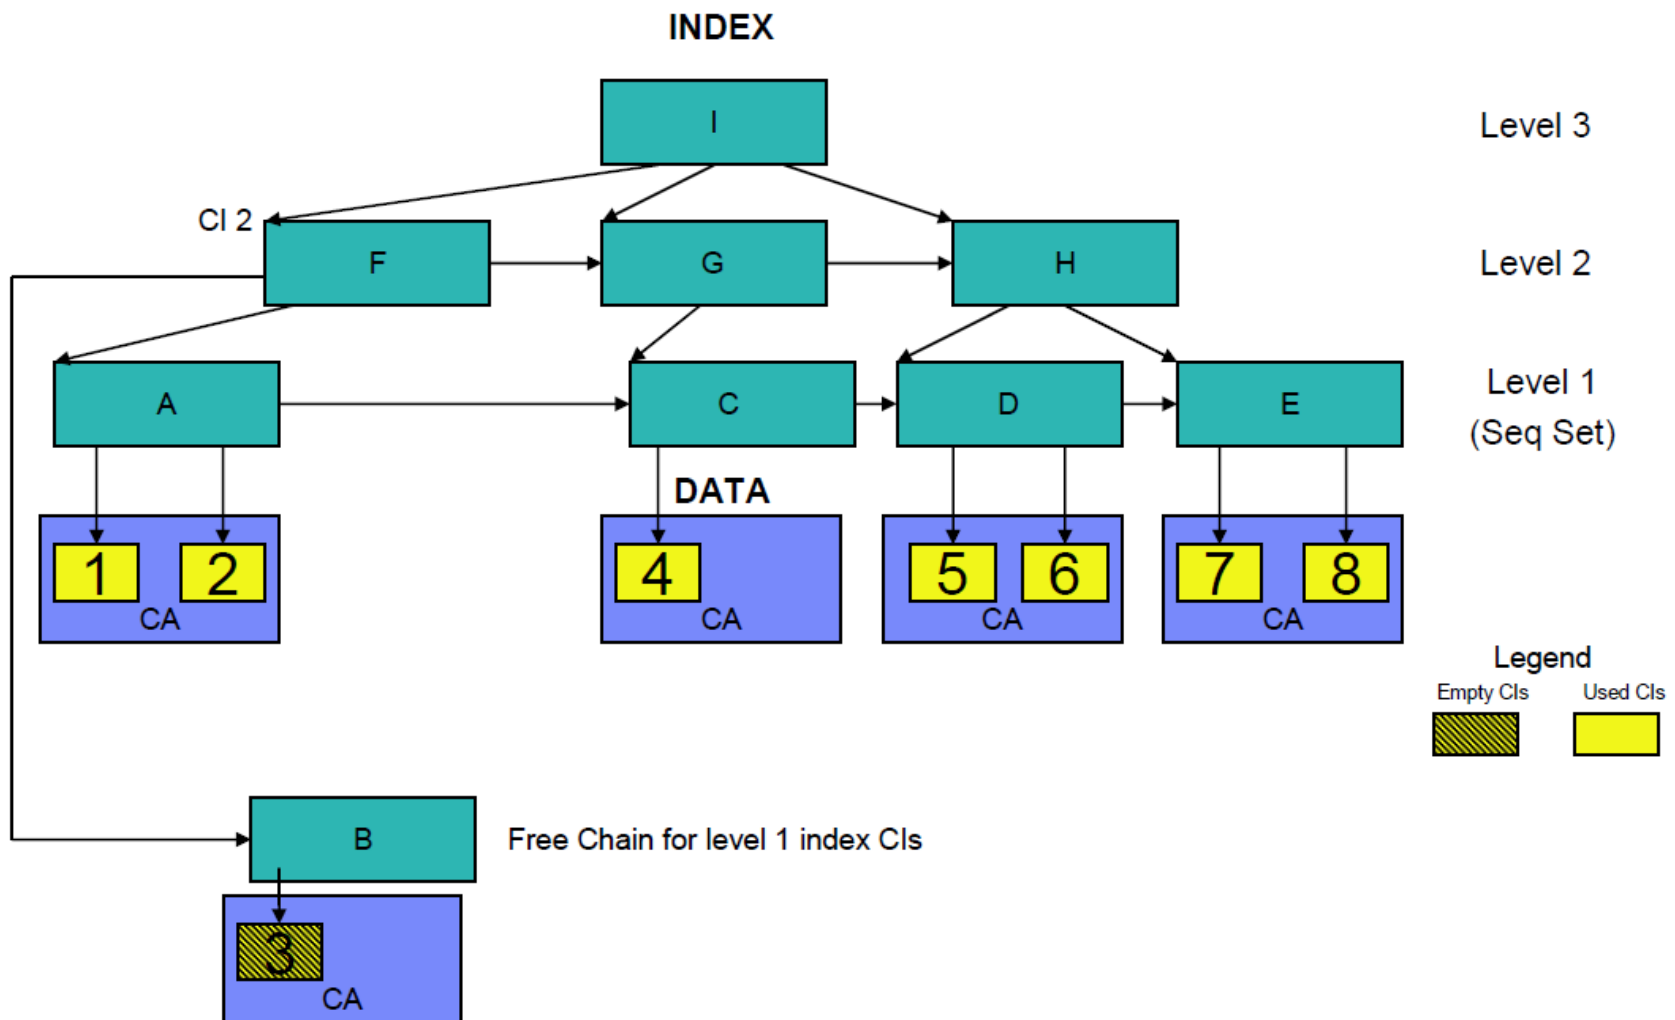
- **Benefits**

- Improved sequential and direct read performance
  - Sequential processing searches the entire index, even the empty CAs
- Fewer extents
  - Reclaimed CAs are reused rather than calling EOVS to get additional space
- Actual benefit will depend on the degree of VSAM data fragmentation and how the data is accessed
  - Anticipated KSDSs that are severely fragmented or rarely reorganized will see the most benefit
  - Applications that delete a large number of records from a narrow key range and then immediately re-insert them could potentially see some performance degradation

# VSAM CA Reclaim Overview



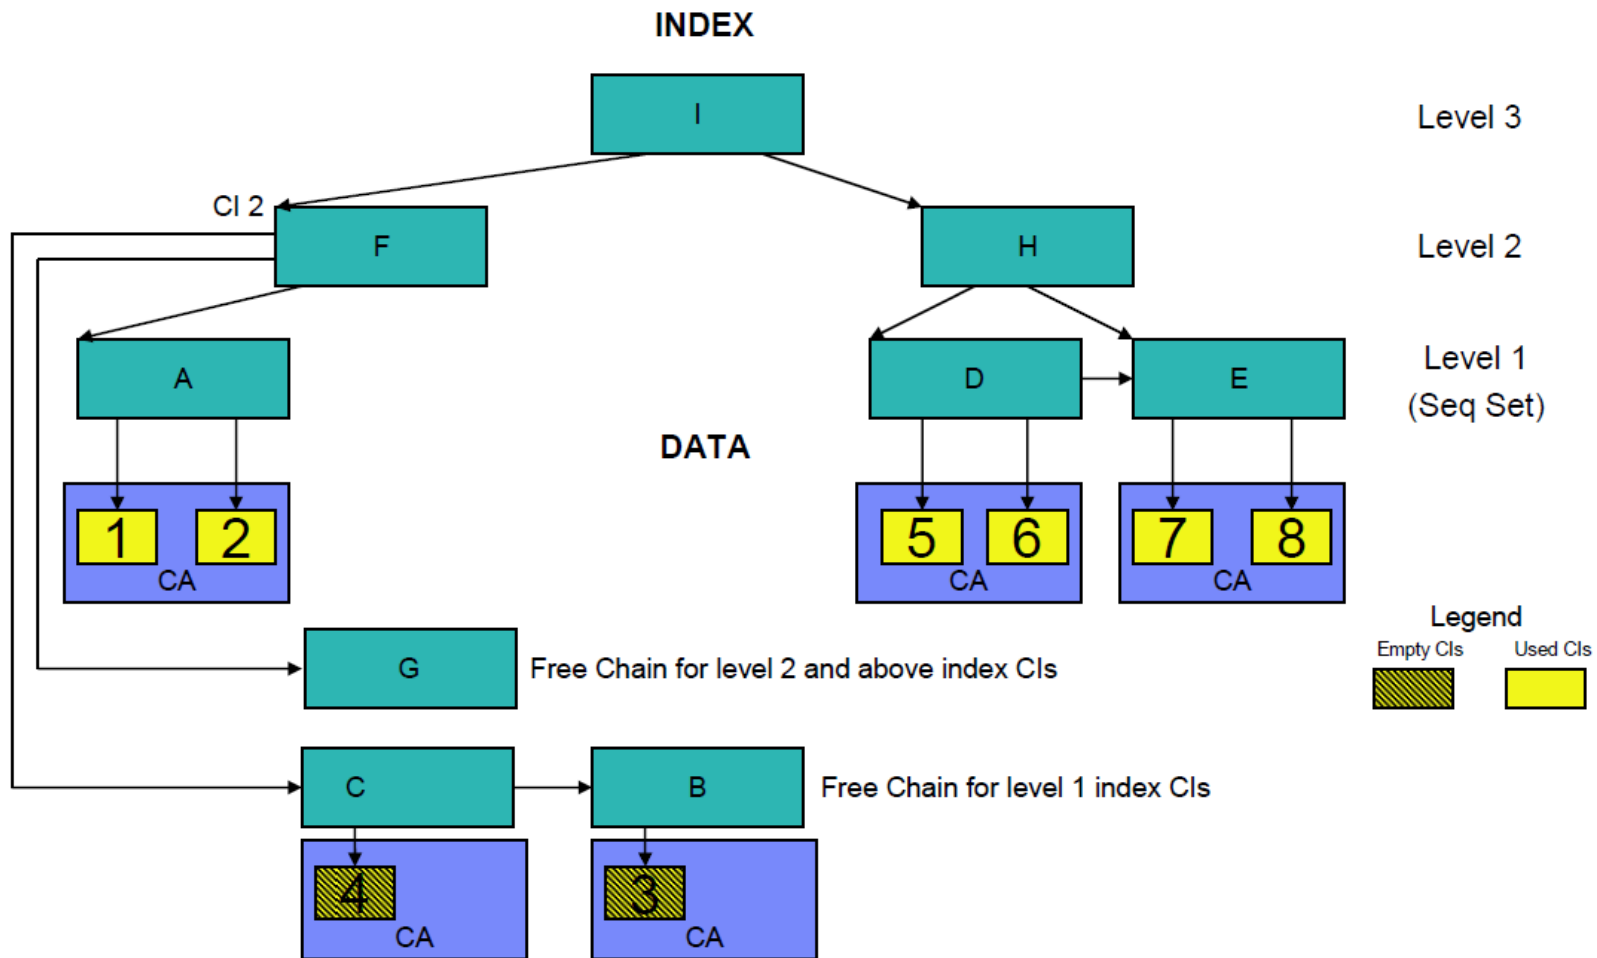
- After CA Reclaim – 1 free CA



# VSAM CA Reclaim Overview



- After CA Reclaim – 2 free CAs



# CA Reclaim Usage and Invocation – System Level



- CA Reclaim is disabled by default on a system level, but is enabled by default for all KSDSs without having to redefine the data set
- New system level parameter in SYS1.PARMLIB(IGDSMSxx):  
**CA\_RECLAIM(NONE | {DATACLAS|DATACLASS})**
  - When system level NONE is specified (default), CA Reclaim is disabled for all KSDSs on this system
  - When system level “DATACLAS|DATACLASS” is specified, the setting of new data class parameter: CA\_Reclaim(Y/N) is used when the data set is defined to determine whether the data set is eligible for CA reclaim
- CA Reclaim can be enabled/disabled on a system level
  - **SETSMS CA\_RECLAIM(NONE | {DATACLAS|DATACLASS})**
    - Takes effect immediately without an IPL
    - SYSPLEX – must route using RO \*ALL, SETSMS
  - PARMLIB option CA\_RECLAIM() or SET SMS=xx or RLS VSAM address space recycle
    - Takes effect after the IPL or recycle



## Changes to RLS Sysplex Cache Manager (SCM)

- There are three new SCM messages issued for SETSMS CA\_RECLAIM(NONE | {DATACLAS|DATACLASS}) :
  - **IGW467I DFSMS CA\_RECLAIM PARMLIB VALUE SET DURING SMSVSAM ADDRESS SPACE INITIALIZATION ON SYSTEM: system-name CURRENT VALUE: parm-value**
    - System-name System on which the SMSVSAM address space was initialized
    - Parm-value Current value of the CA\_Reclaim parameter
  - **IGW467I DFSMS CA\_RECLAIM PARMLIB VALUE CHANGED ON SYSTEM:**
    - System-name
    - OLD VALUE: old-value
    - NEW VALUE: new-value
  - **IGW467I DFSMS CA\_RECLAIM PARMLIB VALUE ON SYSTEM: System-name CURRENT VALUE: parm-value**
    - The value did not change

# CA Reclaim Usage and Invocation - Dataset Level



- When system level “DATACLAS|DATACLASS” is specified, the setting of a new data class parameter is used to determine whether the data set is eligible for CA reclaim
- CA Reclaim is enabled by default for individual KSDSs
  - ISMF data class keyword CA\_Reclaim(Y/N)
  - DATACLAS attribute processed only when a catalog entry is defined (DEFINE, redefine, MIGRATE, RECALL)
    - Value is saved in the catalog regardless of the SYS1.PARMLIB setting
  - If DATALCAS is changed afterward, the change will not take effect unless the KSDS is re-DEFINED (or ALTER)
- CA Reclaim can be enabled/disabled on a data set level
  - **ALTER RECLAIMCA|NORECLAIMCA**
  - The ALTER specification for a KSDS changes the resultant cataloged CA reclaim attribute obtained from the data class
    - It will take effect at the first OPEN following the CLOSE of all open ACBs against the data-set control block structure
  - In case of error:  
**IDC3180I CA RECLAIM ALTER IS NOT VALID FOR THE DATA SET TYPE – non KSDS or IMBED**

# CA Reclaim Usage and Invocation

- ISMF Data Class Define/Alter Panel Changes
  - CA Reclaim
    - Y - reclaim free CAs (Default)
    - N - do not reclaim free CA s

```

Panel Utilities Scroll Help
-----
DGTDCDC6          DATA CLASS DEFINE/ALTER          Page 5 of 5
Command ===>

SCDS Name . . . : USER6.MYSCDS
Data Class Name : DC1

To DEFINE/ALTER Data Class, Specify:
Shareoptions Xregion . . . (1 to 4 or blank)
                Xsystem . . . (3, 4 or blank)
Reuse . . . . . N (Y or N)
Initial Load . . . . . R (S, R or blank)
BWO . . . . . (TC, TI, NO or blank)
Log . . . . . (N, U, A or blank)
Logstream Id . . . . .
FRlog . . . . . (A, N, R, U or blank)
RLS CF Cache Value . . . . A (A, N or U)
RLS Above the 2-GB Bar . . N (Y or N)
Extent Constraint Removal N (Y or N)
CA Reclaim . . . . . Y (Y or N)

Use ENTER to perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.

```

# CA Reclaim Usage and Invocation

- ISMF Data Class Messages
  - CA Reclaim

```
DGTMDC91 -----ISMF MESSAGE-----
HELP
COMMAND ===>
  MESSAGE NUMBER:  DGTDC091

  SHORT MESSAGE:   Inv CA Reclaim Value

  LONG MESSAGE:    CA Reclaim must be Y (Yes) or N (No)

  EXPLANATION:
    You have entered an invalid CA Reclaim value.

  SUGGESTED ACTION:
    Correct and reenter a valid value for CA Reclaim.
    CA Reclaim must be Y (YES), or N (NO).
  Use ENTER to continue with ISMF HELP; Use END to return to ISMF.
```

# CA Reclaim Usage and Invocation

- **Recommendations**

- KSDSs can be classified into three groups with respect to CA Reclaim:
  - CAs are never completely erased. Using CA reclaim will have no impact.
  - CAs are emptied, however, the similar record key ranges are re-inserted in a timely manner, reusing the CA. Using CA reclaim may cause a slight performance impact.
  - CAs are emptied and erased key ranges are not reinserted in a timely manner (if ever). Using CA reclaim will provide performance and space improvements.
- Most KSDSs will likely benefit from CA reclaim, but to give more control to applications, all KSDSs can be enabled/disabled on an individual basis.

# CA Reclaim Usage and Invocation

- **Recovery**

- If an ABEND or Cancel occurs when CA reclaim is in progress, VSAM and RLS will try to complete the CA reclaim
  - The philosophy is to complete the interrupted process as soon as feasible to minimize any possible complications
  - The recovery will complete the interrupted CA reclaim even after the user has issued the SETSMS command to disable CA reclaim
  - The CA reclaim algorithm is such that when it is interrupted, the worst case will be wasted DASD space, no worse than without CA reclaim
  - There are no data integrity problems caused by interrupted CA reclaim

- **Coexistence / Toleration**

- DFSMS Components require toleration APARs for earlier releases to coexist with CA Reclaim
  - VSAM OA26466
  - VSAM RLS OA25108, OA32711, OA33397
  - SMS OA27557 (V1R10 only)
  - Catalog OA32713
  - HSM OA29142
  - DSS OA29022
  - AMS OA26256
- **NOTE: May cause data integrity problems if not on**

## CA Reclaim Reporting - LISTCAT

- LISTCAT output will indicate if CA Reclaim is active for this data set
  - Based on value at DEFINE or after ALTER – not impacted by PARMLIB or SETSMS
  - NOTE: ALTER does not take affect until all ACBs are closed but does change the catalog value immediately

### SMSDATA

```
STORAGECLASS ---SXPXS04  MANAGEMENTCLASS--BKDAILY
DATACLASS -----ABCD  LBACKUP ---0000.000.0000
CA RECLAIM -----(YES)
BWO STATUS-----00000000  BWO TIMESTAMP---00000 00:00:00.0
BWO-----NO
```

**'YES' (default) means CA reclaim is enabled.**

# CA Reclaim Reporting - LISTCAT



- LISTCAT output will indicate the number of CA Reclaims (listed under the INDEX component)
  - **REC-DELETED** – number of CA reclaims since the dataset was created
  - **REC-INSERTED** – number of CA reclaimed and reused CAs
- Two fields in the INDEX have always been 0 prior to CA Reclaim
- Totals since the data set was created
  - In contrast, the numbers in the SMF64 record are totals since the last time the SMF64 record was cut.
  - Values are cumulative (ie always increasing)
- NOTE: values don't necessarily reflect the total number of CAs in the dataset

```

LISTCAT ALL ENTRIES(VSAMTST.CARVSC12.KSDS01)          05510005
0 INDEX ---- VSAMTST.CARVSC12.KSDS01.INDEX
IN-CAT --- SYS1.ICFCAT.PERFVSAM
HISTORY
  DATASET-OWNER----(NULL)  CREATION-----2010.203
  RELEASE-----2  EXPIRATION-----0000.000
  PROTECTION-PSWD----(NULL)  RACF------(NO)
ASSOCIATIONS
  CLUSTER--VSAMTST.CARVSC12.KSDS01
ATTRIBUTES
  KEYLEN-----19  AVGLRECL-----0  BUFSPACE-----0  CISIZE-----512
  RKP-----0  MAXLRECL-----505  EXCPEXIT------(NULL)  CI/CA-----49
  SHROPTNS(3,3)  RECOVERY  UNIQUE  NOERASE  NOWRITECHK  NOIMBED  NOREPLICAT  UNORDERED
  NOREUSE
STATISTICS
  REC-TOTAL-----75  SPLITS-CI-----6993  EXCPS-----5510170  INDEX:
  REC-DELETED-----49020  SPLITS-CA-----1  EXTENTS-----1  LEVELS-----3
  REC-INSERTED-----48951  FREESPACE-%CI-----0  SYSTEM-TIMESTAMP:  ENTRIES/SECT-----5
  REC-UPDATED-----0  FREESPACE-%CA-----0  X'C6519A1460F5925C'  SEQ-SET-RBA-----0
  REC-RETRIEVED-----0  FREESPC-----438272  HILEVEL-RBA-----32768
ALLOCATION
  SPACE-TYPE-----TRACK  HI-A-RBA-----476672
  SPACE-PRI-----19  HI-U-RBA-----38400
  SPACE-SEC-----1
VOLUME
  VOLSER-----MLSV00  PHYREC-SIZE-----512  HI-A-RBA-----476672  EXTENT-NUMBER-----1
  DEVTYPE----X'3010200F'  PHYRECS/TRK-----49  HI-U-RBA-----38400  EXTENT-TYPE----X'00'
  VOLFLAG-----PRIME  TRACKS/CA-----1
EXTENTS:
  LOW-CCHH---X'0051000B'  LOW-RBA-----0  TRACKS-----19
  HIGH-CCHH---X'0052000E'  HIGH-RBA-----476671
    
```



## CA Reclaim Reporting - EXAMINE

- Existing EXAMINE DATATEST or DATATEST INDEXTEST message IDC11724I
  - **IDC11724I DATA COMPONENT CA NOT KNOWN TO SEQUENCE SET** continues to display orphan CAs, but RC changed from 8 to 4
  - NOINDEXTEST continues to return RC = 8
- EXAMINE will not distinguish whether the orphan CAs were created by interrupted CA splits or CA reclaim.
- New EXAMINE messages will be generated to show information related to CA reclaim.

## CA Reclaim Reporting - EXAMINE

- How to decide whether a KSDS should be converted to use CA reclaim

- IDCAMS EXAMINE DATATEST will show the number of empty unreclaimed CAs

**IDC01728I FOUND nnnnnn EMPTY CONTROL AREAS THAT HAVE NOT BEEN RECLAIMED**

- Can be used to determine if data set is good candidate for CA Reclaim
- Number will not change if CA Reclaim enabled
  - Old Empty CAs remain
  - Only new empty CAs are reclaimed

## CA Reclaim Reporting - EXAMINE

### **IDC11768I [CI SPLIT | CA RECLAIM] IN PROGRESS**

- Caused by ABEND, CANCEL, or a system failure
- Informational, not a data integrity issue
- Does not leave orphan CA
- Corrected in subsequent access
- Different from IDC11778I

### **IDC11778I CA RECLAIM IN PROGRESS**

- Caused by ABEND, CANCEL, or a system failure
- Informational, not a data integrity issue
- Temporary condition
- Index CI 2 is dumped. Processing continues

## CA Reclaim Reporting - VERIFY RECOVER

- **IDCAMS VERIFY RECOVER option**
  - RECOVER is new parameter
    - Completes an interrupted CA reclaim, in addition to what IDCAMS VERIFY currently does
    - May cause index update, so no other applications should have the data set OPEN
    - Default is no recover
  - The use of this command is optional in that
    - The following cases will do CA reclaim recovery automatically:
      - *Non-RLS or RLS VSAM PUT*
      - *An ERASE that starts a new CA reclaim*
      - *An index search that runs into an index CI that was involved in an interrupted CA reclaim*
      - *RLS recovery that automatically gets control as a result of the interruption will complete the interrupted CA reclaim*
    - Even if CA reclaim recovery is not done, with the partially completed CA reclaim the data set is still 100% accessible.

## CA Reclaim Reporting – SMF Records

- SMF 64 record contains new fields for CA Reclaim
  - written at CLOSE and EOVS

3 SMF64DAU FIXED(32); /\* Change in the number of CA-reclaimed control areas reused in the KSDS since the last EOVS or CLOSE\*/

3 SMF64RLM FIXED(32); /\* Number of CAs Reclaimed \*/

## CA Reclaim Updated Publications

- z/OS V1R12.0 DFSMS: Using Data Sets (SC26-7410-10)
- z/OS V1R12.0 DFSMS Access Method Services for Catalogs (SC26-7394-11)
- z/OS V1R12.0 DFSMS Using the Interactive Storage Management Facility (SC26-7411-07)
- z/OS V1R12 DFSMS: Implementing System Managed Storage (SC26-7407-07)
- z/OS V1R12.0 DFSMS Storage Administration Reference (for DFSMSdfp, DFSMSdss, DFSMShsm) (SC26-7402-13)
- z/OS V1R12.0 DFSMS Using the New Functions (SC26-7473-07)
- z/OS V1R12.0 MVS System Commands (SA22-7627-23)
- z/OS V1R12.0 MVS System Messages, Vol 7 (IEB-IEE) (SA22-7637-20)
- z/OS V1R12.0 MVS Initialization and Tuning Reference (SA22-7592-21)
- z/OS V1R12.0 DFSMSdfp Diagnosis Reference (GY27-7618-11)

# Summary


- **CA Reclaim**

- Detects Empty CAs
- The empty CAs are removed from the active index structure and moved to free chains
- Uses the empty CAs when needed before using new CAs
- Reduces or eliminates reorgs in support of 24x7 operations
- Applications that use VSAM key-sequenced data sets (KSDS) can benefit from improved performance, minimized space utilization, and improved application availability
- Performance improvements are anticipated for many applications using CICS®, VSAM, VSAM RLS, IMS™ VSAM, and Catalog processing

**Thank You!**



# Trademarks and Disclaimers



**SHARE**  
Technology • Connections • Results

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

CICS*	FICON*	Lotus*	System Storage
DB2*	FlashCopy*	MQSeries*	Tivoli*
DFSMS	GDDM*	Multiprise*	TotalStorage*
DFSMSdftp	GDPS*	OMEGAMON*	Virtualization Engine
DFSMSdss	geoManager*	OS/390*	VisualAge*
DFSMSHsm	HiperSockets	Parallel Sysplex*	VM/ESA*
DFSMSmm	HyperSwap	PR/SM	VSE/ESA
DFSORT	IBM*	QMF	VTAM*
DFSMS	IBM logo*	RACF*	WebSphere*
DS4000	ImagePlus*	Rational*	z/Architecture*
DS6000	IMS	RMF	z/OS*
DS8000	Intelligent Miner	System i	z/VM*
Enterprise Storage Server*	Language Environment*	System z	z/VSE
ESCON*		System z9	zEnterprise
			zSeries*
			zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

## NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided "AS IS" without warranty of any kind.

## Trademarks and Disclaimers (continued)



### NOTES:

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.