VSE/VSAM Fundamentals, Hints &Tips and Best Practices

Mikhail Zaslonko

z/VSE 50
50 years of innovation

http://www.ibm.com/zVSE
http://twitter.com/IBMzVSE

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are 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.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
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 subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

zVSE/VSAM Components

Virtual Storage Access Method (VSAM)

Catalog Management

- Maintains attributes of all files (clusters) defined to VSE/VSAM
- Allocates and maintains DASD Space

Open/Close Management

- Connects and Disconnects a cluster with an application program
- Ensures access integrity

Record Management

- Performs all I/O access to clusters and catalogs
- Manages Buffer Pools
- Ensures cluster's data integrity

VSAM Utilities

- IDCAMS
- IKQVCHK, IKQVEDA, IKQVDU, IKQPRED, etc.

VSAM Data Organization - Catalogs

Two Level Hierarchical Catalog structure

- One Master catalog per system
- Optional, User Catalogs . As many as required.
- Maximum one catalog per volume
- Each Catalog can own space on multiple volumes
- Multiple catalogs can own space on one volume
- VSAM Files are called CLUSTERS
- Catalog contain clusters in the VSAM space it owns
- A Catalog can be shared by multiple VSE systems

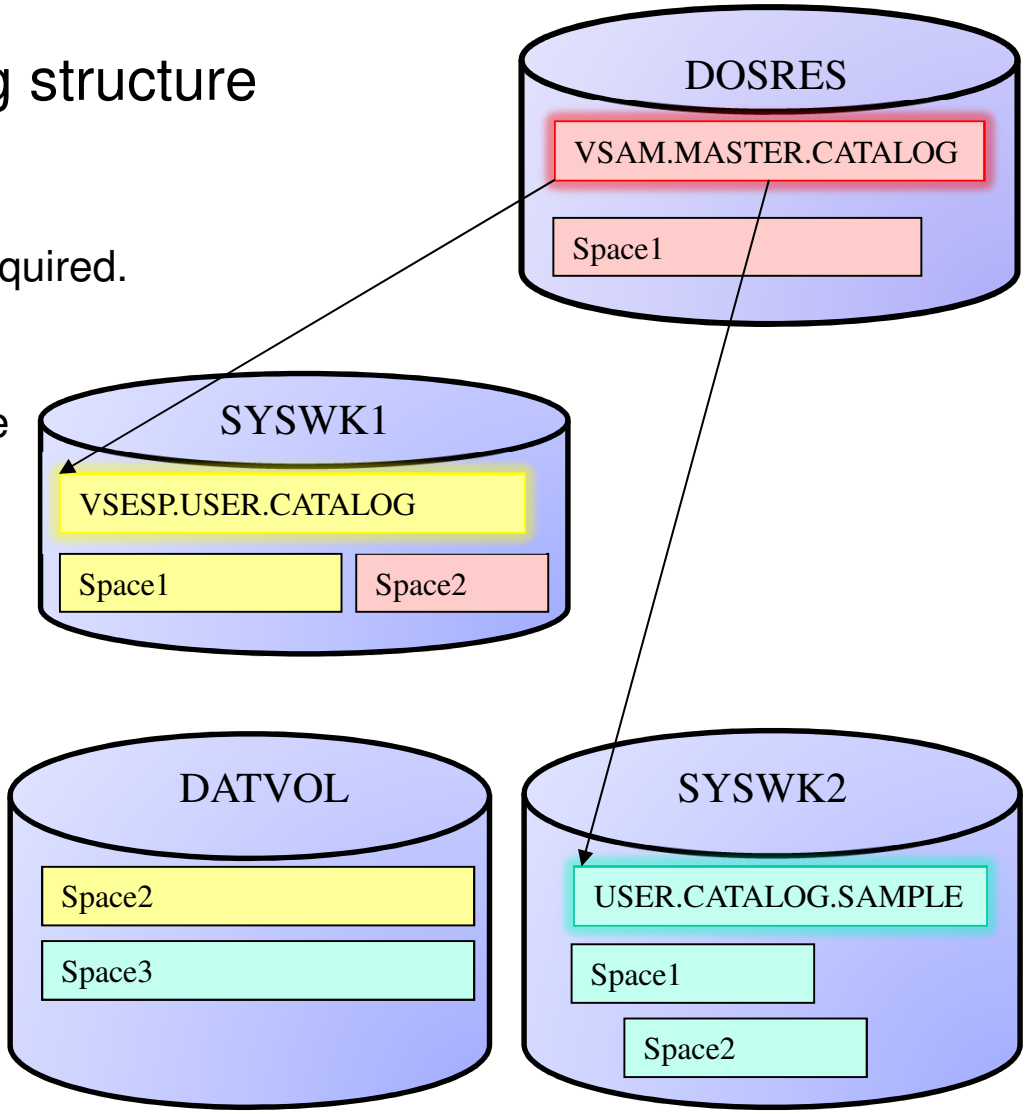

VSAM Data Organization - Catalogs

Master Catalog (IJSYSCT)

- One Master catalog per system. Defined during system installation, normally on DOSRES. Assigned (via DEF SYSCAT) during IPL.
- VSE.MESSAGES.ONLINE File (IESMSGS).
- Definitions of User Catalogs (unique entries for Master Catalog).
- Definition for VSE libraries in VSAM managed space (PRD1, PRD2).

User Catalog

- Optional. As many as required.
- Created with IDCAMS utility.
- Only one catalog per volume.
- Unique Catalog name.
- Contains VSAM clusters (user's data).
- May be shared by multiple VSE systems.

VSAM Data Organization - Catalogs

VSESPUC User Catalog (VSESP.USER.CATALOG)

- On-line System Files:
  - VSE.CONTROL.FILE (IESCNTL)
  - CICS Start-up Dataset (CSD)
  - Restart Dataset (RSD)
  - Global Catalog (GCD, CICS TS)
  - Local Catalog (LCD, CICS TS)
  - Transient data, Intra-partition dataset (TD.INTRA)
  - Temporary Storage (DFHTEMP)
  - Data Management Facility (DMF) file
  - Transaction Abend Dump Library (DFHDMPA / DFHDMPB)
  - On-line Problem Determination File (IESPRB)
  - VSE Primary Library (Alternate ICCF Library)
- System Work Files
- PTF.FILE (Used to apply PTFs from disk)
- Text Repository File (IESTRFL)
- VSE/VSAM Record Mapping Definitions (See e-business connectors)
- CICS REXX files (RFSDIR1, RFSPOL1, RFSDIR2, RFSPOL2)
- CICS Listener (EZACONF, EZACACH)

VSAM Data Organization - Catalogs

VSAM Catalog Contents

Catalog has internal format of VSAM KSDS file with Key-Ranges

- Key length – 44 bytes
- Record length – 512 bytes

High Key Range (True Name)

- Contains index of 44-character names to internal catalog CI#
- Used to address Cluster and Volume entries in the Low Key Ranage

Low Key Range

- Self-describing records (including cluster definitions for catalog itself)
- Catalog Control record
- Volume definitions (including dataspace, space-map and dataset-directory information)
- Cluster definitions (including Data, Index, AIX, Path and Upgrade-Set entries)

Index

- Used only to access True-name records

VSAM Data Organization - Catalogs

Catalog Recommendations

Naming conventions

- Cluster components naming (Data and Index) - explicitly or IDCAMS defaults. Dataset-Directory section in LISTCAT includes data and index, not cluster names.
- Name clusters (and catalogs, where applicable) to include application names.
- Exploit partition and system independent naming (% or %%) (SAM ESDS files only).

File organization

- Place static (once defined, multi-access) and dynamic (Reusable) files in separate catalogs (on separate volumes).
- Place batch vs on-line files in separate catalogs.
- Do not put all your eggs in one basket. (Do not keep large number of files in a single Catalog).
  - Slower performance.
  - Recovery aspect.
  - DY47482 – Listcat problems.

VSAM Data Organization - Catalogs

Catalog Recommendations (cont.)

Allocation issues

- Catalogs are limited to 16 extents, and can only expand on original volume.
- Can be monitored by checking Catalog self-descriptor cluster Listcat entry (See "*VSE/VSAM Commands*" under "*Appendix A. Interpreting LISTCAT Output*").
- Always specify secondary allocation in the Catalog definition. However, consider large enough primary allocation.
- Define catalog with "DEDICATE" at cluster level, and explicit catalog allocation at "DATA" level.
- When restoring an entire catalog, do not set the catalog volume as first volume in list.
- If catalog fills up, re-define with additional allocation at "DATA" component level.

Catalog consistency checking with IKQVCHK utility

- Use Catalog Checker (IKQVCHK) utility to verify User Catalog for internal consistency.
- Run IKQVCHK on regular basis.
- Refer to the support specialists in case of IKQVCHK reported issues.

Consider Catalog rebuild as a part of maintenance

- Provides Catalog consistency and better performance.

VSAM Data Organization – Clusters

VSAM Cluster Types

KSDS (Key-Sequenced Data Set)

- Direct or Sequential (Browse) access by, RBA, or Key
- Insert in key sequence.
- Contains data and index components.

ESDS (Entry-Sequenced Data Set)

- Sequential (Browse) Access
- Direct Access by Relative Byte Address (RBA)
- Insert only at end-of-file.
- Record update only allowed if record length does not change.
- SAM ESDS is a unique sub-set

RRDS (Relative-Record Data Set)

- No Index component. Records retrieved using Relative Record Number (RRN) as key.
- Direct or Sequential (Browse) access by RBA, or RRN (key)
- Fixed length records only
- Insert / Update in RRN sequence. (Cannot change length of record)

VSAM Data Organization – Clusters

VSAM Cluster Types

VRDS (Variable-Length Relative-Record Data Set)

- Application access identical to RRDS. Uses RRN as key to access / insert / update records.
- Allows variable length records.
- Contains data and index components.

Alternate Index

- KSDS file.
- Linked via Path to Base Cluster, which may be KSDS (keys) or ESDS (RBA)
- See details on page 23.

Cluster Definition

- Created via IDCAMS utility
- Cluster Name – unique for the Catalog
- Data Component Name
- Index Component Name (KSDS or VRDS)

VSAM Data Organization – Clusters

Control Interval format

Logical Records can be:

- Fixed (Average record size = maximum record size)
- Variable (Average record size < maximum record size)
- Spanned (larger than control interval)
- Compressed

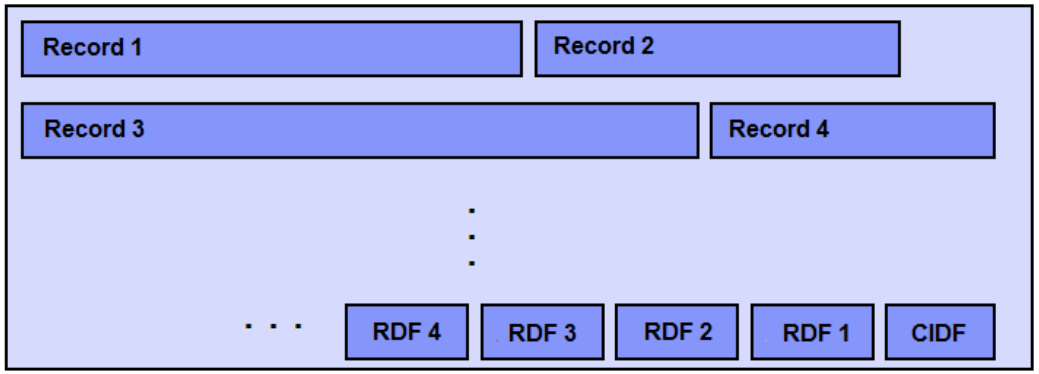

Control Interval Definition Field (CIDF):

- (four bytes) describes free-space in control interval

Record Definition Field (RDF):

- If records have different lengths - one RDF (three bytes) per record.
- If a series of records have the same length - two RDFs per sequence (*nn* records of *mm* length).
- Spanned record has two RDFs per CI. Identifies position of this segment in logical record (start, mid, end).

Index:

- One low-level index record per data CA, with a pointer for each data control interval, containing the highest index value of that CI.

VSAM Data Organization – Clusters

Data Control Interval with 25 fixed length (80-byte) records:

```
000000  F0F0F1F0 F0F1F4F5 F2F7F0F1 F9F4F7F0  F2F2F2F2 F4F3F6F2 F8F5F2F1 F7F24040  *00100145270194702222436285217 2  *
000020  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                                *
000040  40404040 40404040 40404040 40404040  F0F0F1F0 F6F9F1F4 F4F6F2F2 F7F5F2F0  *                0010691446227520*
000060  F1F1F4F3 F1F9F6F4 F9F7F9F9 F7F54040  40404040 40404040 40404040 40404040  *11431964979975                  *
000080  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                                *
0000A0  F0F0F1F0 F7F1F3F3 F3F8F9F6 F9F4F4F4  F6F2F5F7 F0F8F0F2 F0F7F6F3 F8F64040  *00107133389694446257080207638 6  *
0000C0  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                                *
0000E0  40404040 40404040 40404040 40404040  F0F0F1F0 F8F9F6F1 F8F7F6F5 F3F1F5F2  *                0010896187653152*
000100  F6F1F4F9 F0F9F1F2 F2F9F0F6 F3F24040  40404040 40404040 40404040 40404040  *61490912290632                  *
000120  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                                *
000140  F0F0F1F0 F9F1F6F7 F7F7F8F1 F7F3F9F3  F2F3F9F1 F6F2F8F7 F1F0F0F6 F8F54040  *00109167778173932391628710068 5  *
000160  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                                *
000180  40404040 40404040 40404040 40404040  F0F0F1F1 F0F5F5F3 F3F3F5F9 F3F8F4F6  *                0011055333593846*
0001A0  F0F7F7F5 F2F4F3F9 F4F7F9F8 F3F84040  40404040 40404040 40404040 40404040  *07752439479838                  *
0001C0  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                                *
                .
                .
                .
0007E0  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *                                *
000800 to  000FDF  same as previous line
000FE0  00000000 00000000 00000000 00000000  00000000 00080019 400050 07D00826  *                          &     *
```

'08' - Left-hand RDF
'0019' - 25 fixed-length records

'40' = Right hand RDF
'0050' = 80-byte Records

'07D0' = Displ to Free Space
'0826' = Length of Free Space

VSAM Data Organization – Clusters

Data Control Interval variable length records:

```
000000  F0F0F1F0 F0F1F4F5 F2F7F0F1 F9F4F7F0  F2F2F2F2 F4F3F6F2 F8F5F2F1 F7F24040  *00100145270194702222436285217 2  *
000020  40404040 40404040 40404040 40404040  40404040 4040F0F0 F1F0F6F9 F1F4F4F6  *                      0010691446*
000040  F2F2F7F5 F2F0F1F1 F4F3F1F9 F6F4F9F7  F9F9F7F5 40404040 40404040 40404040  *2275201143196497997 5          *
000060  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                              *
000080  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                              *
0000A0  40404040 40404040 40404040 40404040   0404040 40404040 4040C1E2 C4C6C1E2  *                         ASDFAS*
0000C0  C4C6C1E2 C4C6C1E2 C4C60000 00000000  00000002 27520F00 00000000 00000000  *DFASDFASDF                    *
0000E0  0003E840 40404040 40404040 40404040  40404040 4040F0F0 F1F0F7F1 F3F3F3F8  *  Y                   0010713338*
000100  F9F6F9F4 F4F4F6F2 F5F7F0F8 F0F2F0F7  F6F3F8F6 40404040 40404040 40404040  *9694446257080207638 6          *
000120  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                              *
000140  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *                              *
000160  40404040 40404040  0404040 40404040  40404040 40404040 4040C1E2 C4C6C1E2  *                         ASDFAS*
000180  C4C6C1E2 C4C6C1E2 C4C60000 00000000  00000009 69444F00 00000000 00000000  *DFASDFASDF           |        *
0001A0  0003E840 40404040 40404040 40404040  40404040 40404040 4000E6C5 C9C7C5D3  *  Y                      WEIGEL*
                .
                .
                .
000F60  00000000 00000000 21177F00 00000000  00000000 0003E840 40404040 40404040  *         "          Y         *
000F80  40404000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *                              *
000FA0  00000000 00000000 0000B900 00280000  5C00004D 00006E00 00950000 4900004F  *               *  (  >  n     |*
000FC0  00005A00 00AF0000 C900004B 00007C00  00DE0000 5000002D 0000F000 00DE0000  *  !     I  .  @     &     0     *
000FE0  9D00009B 00008E00 00F50000 B60000F1  00006000 00E40000 00 000036 0F830025  *      5     1  -  U       c  *
```

'00' = Single RDF
'00C0' = 192 byte record

'00' = Single RDF
'0036' = 54-byte Record

'0F83' = Displ to Free Space
'0025' = Len of Free Space

VSAM Data Organization – Clusters

Control Area format

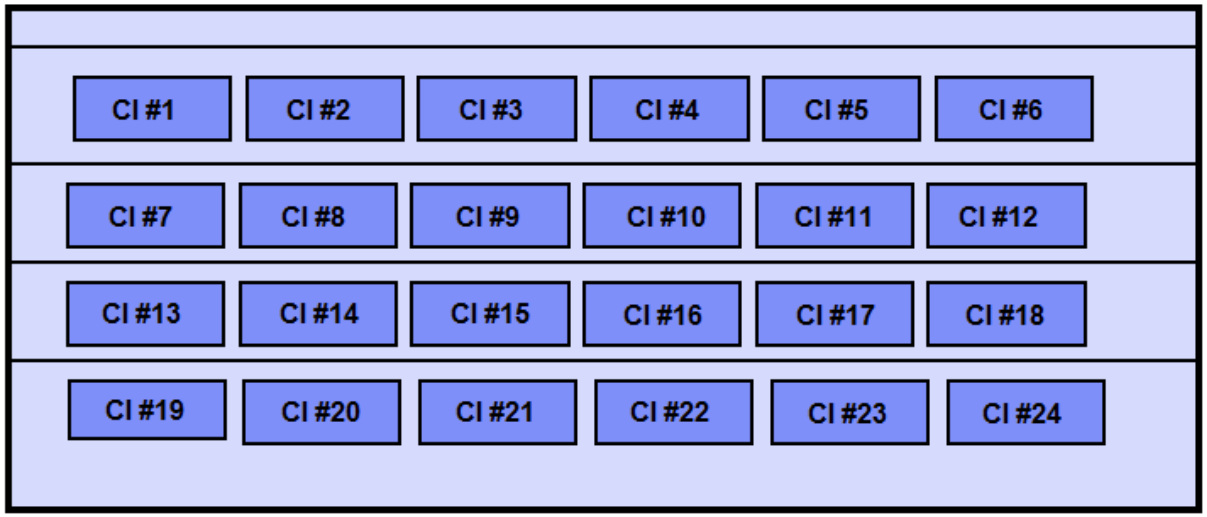

- SAM ESDS (non-CA format. CI can be split at end of track)
- CISIZE vs track utilization (18K best VSAM space utilization, 512k worst case)

3390:	CI Size	CIs per Track	Track Capacity
	512 K	48	24 K
	2 K	21	42 K
	4 K	12	48 K
	8 K	6	48 K
	18 K	3	54 K

VSAM Data Organization – Clusters

Compression:

- Hardware or Software
- Compression Dictionary (created during initial load)
- Compression Control Dataset (CCDS) – one per Catalog.
- Cluster defined using “COMPRESSED” Attribute.

Advantages and Restrictions

- More data stored on dasd extent. Avoid 4 Giga-byte limit.
- For sequential access, more records per buffer (CI), so fewer I/Os.

- At least 40 bytes per record must be available for compression.
- Requires up to 1Meg additional 31-bit GETVIS per file for compression services.

VSAM Data Organization – Clusters

KSDS / VRDS Index Record Structure

Typical Tree Structure, with horizontal chain

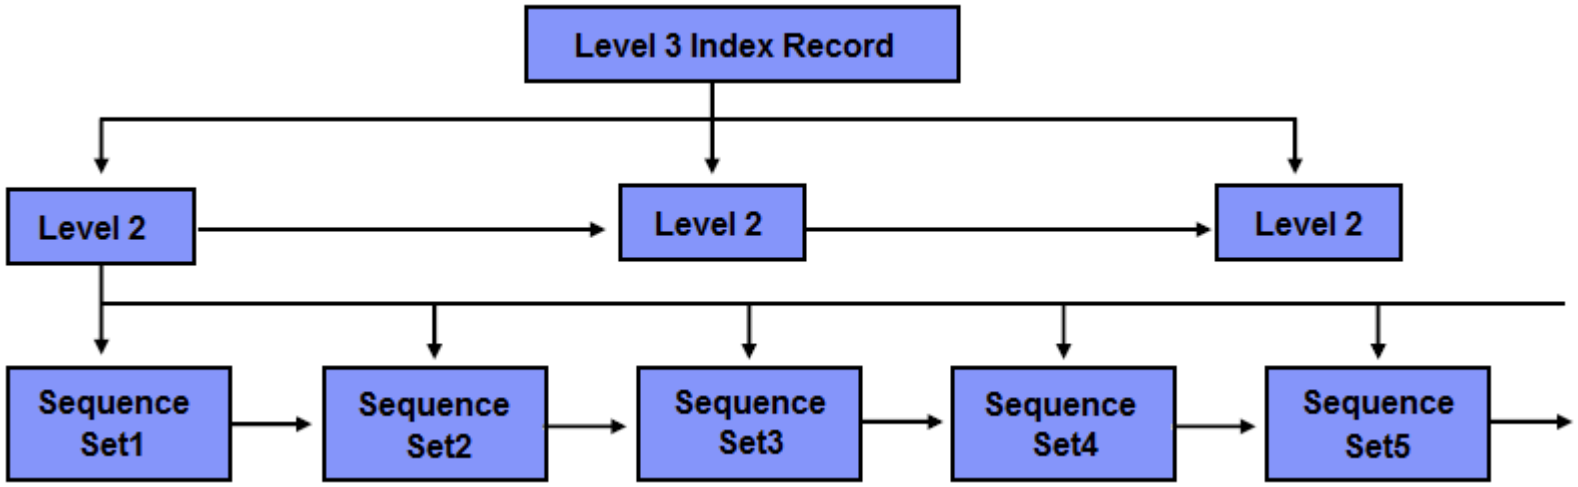

- Each Index Sequence Set record describes a data Control Area.
- Each entry within Sequence Set record describes a data Control Interval.

VSAM Data Organization – Clusters

Format of Sequence Set (Level 1 Index) CI

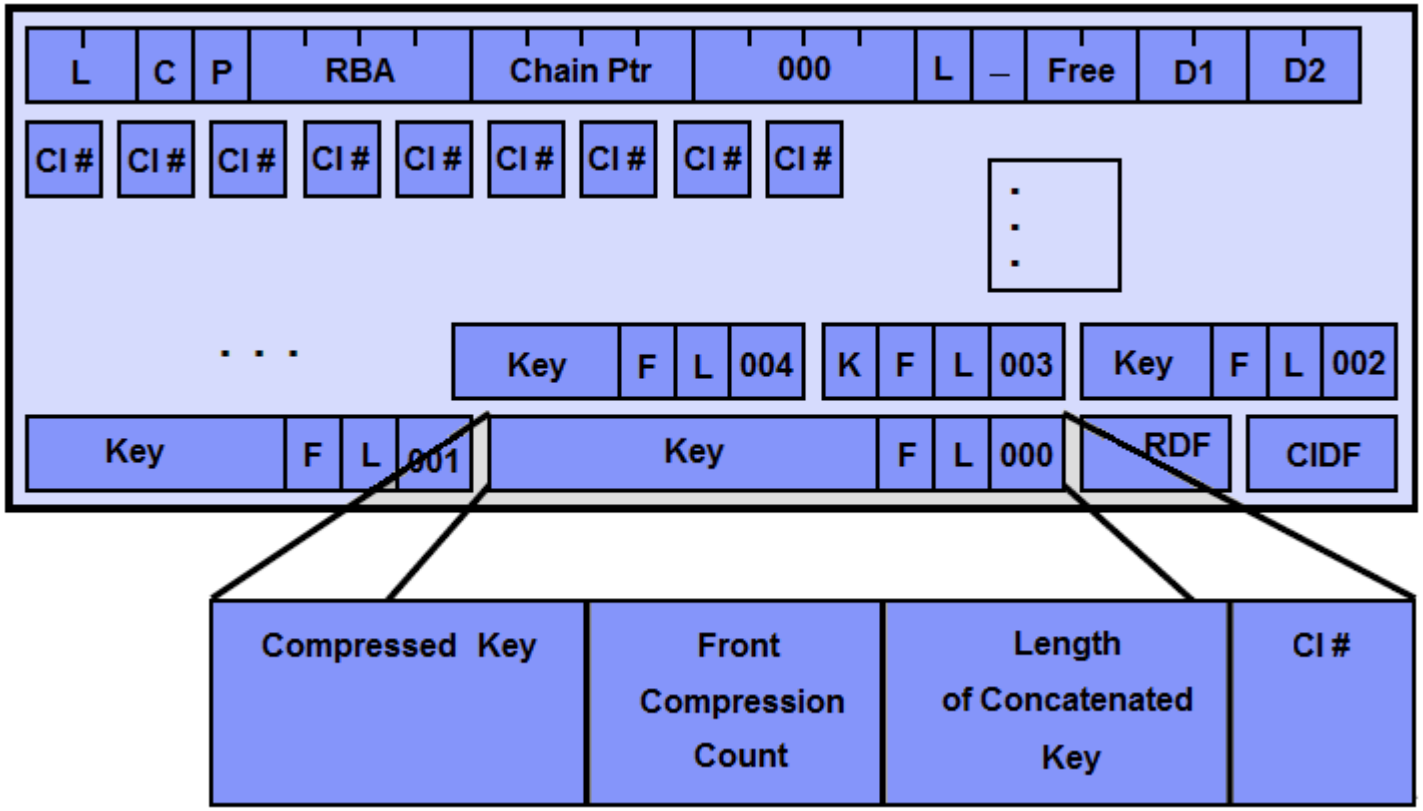

VSAM Data Organization – Clusters

Key Compression in Index records:

- Both front and rear compression
- Very efficient. 100,000 records in 750 tracks required 4 tracks index.

	Highest Key in Control Interval		F	L	CI
CI#1:	00130526303676931813118829736 3	'001305'	00	06	00
CI#2:	00156289438171131513884048010 0	'562'	03	03	01
CI#3:	00176022301356124073455511168 5	'76'	03	02	02
CI#4:	00194947353631991893406207061 0	'94'	03	02	03
CI#5:	00212423411365172552801461524 1	'212'	02	03	04
CI#6:	00232797165453973632874066503 1	'327'	03	03	05
CI#7:	00260150861706835064999945234 9	'60'	03	02	06
CI#8:	00286754647821749874846093380 0	'867'	03	03	07
CI#9:	00307422833414185790007610393 0	'30'	02	02	08
CI#10:	00339032186119283972629198752 9	'3903'	03	04	09
CI#11:	00362138767263491312810017562 3	'621'	03	03	0A
CI#12:	00384143477567355183148441483 2	'841'	03	03	0B

"F" Front Compression
"L" Length

VSAM Data Organization – Clusters

Sequence Set (Level 1 Index) record showing key compression:

```
00000  0FF90301 00000000  00001000 00000000  01000072 0DFE0EFA B3B2B1B0 AFAEADAC  * 9                                *
00020  ABAAA9A8 A7A6A5A4  A3A2A1A0 9F9E9D9C  9B9A9998 97969594 93929190 8F8E8D8C  *  zyxwvuts        rqponmlkj      *
00040  8B8A8988 87868584  83828180 7F7E7D7C  7B7A7978 77767574 73727170 6F6E6D6C  *  ihgfedcba "='@#:          ?>_%*
00060  6B6A6968 67666564  63626160 5F5E5D5C  5B5A0000 00000000 00000000 00000000  *,         /-*;)*$!               *
00080  00000000 00000000  00000000 00000000  00000000 00000000 00000000 00000000  *                                  *
                .
                .
                .
00DE0  00000000 00000000  00000000 00000000  00000000 00000000 0000F2F3 F7F70104  *                          2377    *
00E00  59F4F1F8 030358F3  F2F10203 57F9F703  0256F7F9 F3030355 F5F70302 54F3F3F2  * 418   321   97   793   57   332*
00E20  030353F1 F5030252  F2F0F1F5 020451F6  030150F4 F5F20303 4FF1F1F2 02034EF7  *   15   2015   6  &452  |112  +7*
00E40  03014DF6 F103024C  F4F1F003 034BF2F1  F203034A F2F0F0F0 010449F7 F9F10303  *  (61  <410  .212  ›2000   791   *
00E60  48F5F203 0247F9F1  F5020346 F8F80302  45F6F103 0244F303 0143F8F0 F1020342  * 52   915   88   61   3   801    *
00E80  F8F1F403 0341F5F2  030240F7 F2F00203  3FF8F303 023EF5F1 F503033D F3F30302  *814   52   720   83   515   33   *
00EA0  3CF6F1F1 02033BF8  F3F50303 3AF6F103  0239F3F1 030238F5 F0F10203 37F8F003  * 611   835   61   31   501   80 *
                .
                .
                .
00F80  F2020314 F8F30302  13F5F9F2 030312F3  F2030211 F5F0F402 0310F8F7 03020FF5  *2   83   592   32   504   87   5*
00FA0  F7F40303 0EF4F203  020DF4F1 F5F00204  0CF8F4F1 03030BF6 F2F10303 0AF3F9F0  *74   42   4150   841   621   390*
00FC0  F3030409 F3F00202  08F8F6F7 030307F6  F0030206 F3F2F703 0305F2F1 F2020304  *3   30   867   60   327   212   *
00FE0  F9F40302 03F7F603  0202  F5F6F2030301  F0F0F1F3F0F50006 00  000FF9  0FF90000  *94   76   562   001305     9 9  *
```

'562' = Concatenated Key
03 = Front compression
03 = Concatenated Key Len
01 = Data Control Interval 2

RDF

CIDF

'001305' = Concatenated Key
00 = No front compression
06 = Concatenated Key Len
00 = Data Control Interval 1

VSAM Data Organization – Clusters

Alternate Indexes

- Alternate way to access information in KSDS or ESDS file
- Path Name vs Alternate Index Name
  - Alternate Index is a KSDS file, Path is simply a logical connection between an alternate index and a base cluster.
  - Path can be defined directly over a base cluster (Alias).
- UNIQUEKEY vs NONUNIQUEKEY
  - Non-unique keys require definition of spanned records.
  - Logical Recordsize can be up to one CA in size (810K on 3390)
  - Major impact on 31-bit GETVIS.
- Feedback Code (Return Code):
  - x'08'(8) Duplicate AIX key
  - x'64'(8) (OPEN) empty alternate index
  - x'6C'(8) Too many alternate index keys for record size
  - x'90'(8) Mismatch between AIX and Base
  - x'94'(8) Max number of AIX pointers

VSAM Data Organization – Clusters

"Dummy" aix keys:

- IDCAMS **BLDINDEX** will assign even "dummy" keys to alternate index (blanks, nulls, high values).
- This often causes unnecessarily large alternate index records. (See discussion of "NONUNIQUE" keys on previous page)
- You can write your own build index routine.

Alternate Index Record Format:

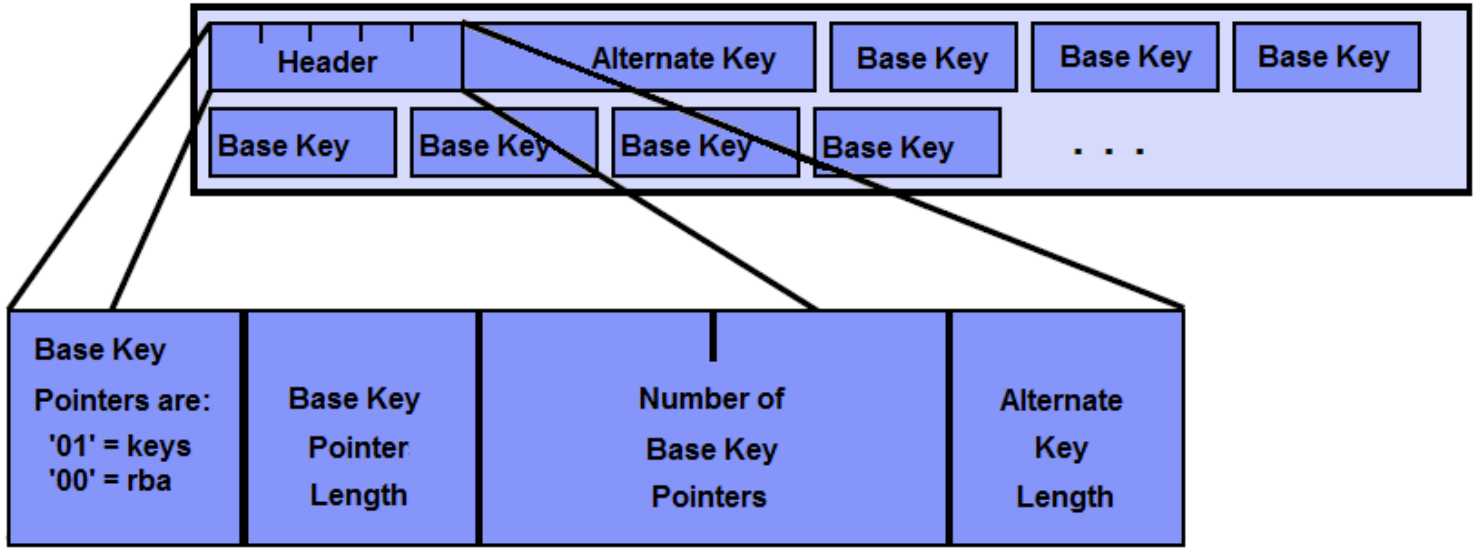

VSAM Data Organization – Clusters

Cluster Size Limits and Recommendations

Volumes considerations

- Up to 123 volumes per cluster component (data and index).
- Maximum 16 volumes can be inherited from Default Model (SAM-ESDS).

Extents considerations

- Clusters are limited to 123 extents per component (data and index).
  - Might become a problem if you specify a very small secondary extent (underallocation).
  - VSE/VSAM will sub-allocate an extent in several pieces (up to 5 segments) in case no continuous free dataspace available.
- Maximum 16 million records per extent (relevant when allocation is specified in RECORDS).
  - Compression can help - VSAM uses the uncompressed record length to calculate how much space to reserve.

4 Giga-byte limit

- 4 byte RBA (Relative Byte Address).
- 4500 – 11000 cylinders on 3390 depending on CI-Size and track utilization.
- Compression can be considered as a circumvention.
- 'ExtraLargeDataset' or 'XXL' - max 286 GB (No-RBA access).

VSAM Data Organization – Clusters

CI Split process

Half the records (those with the higher keys) in the CI are moved into the new CI. The new record is inserted (in key sequence) into the CI to which it belongs.

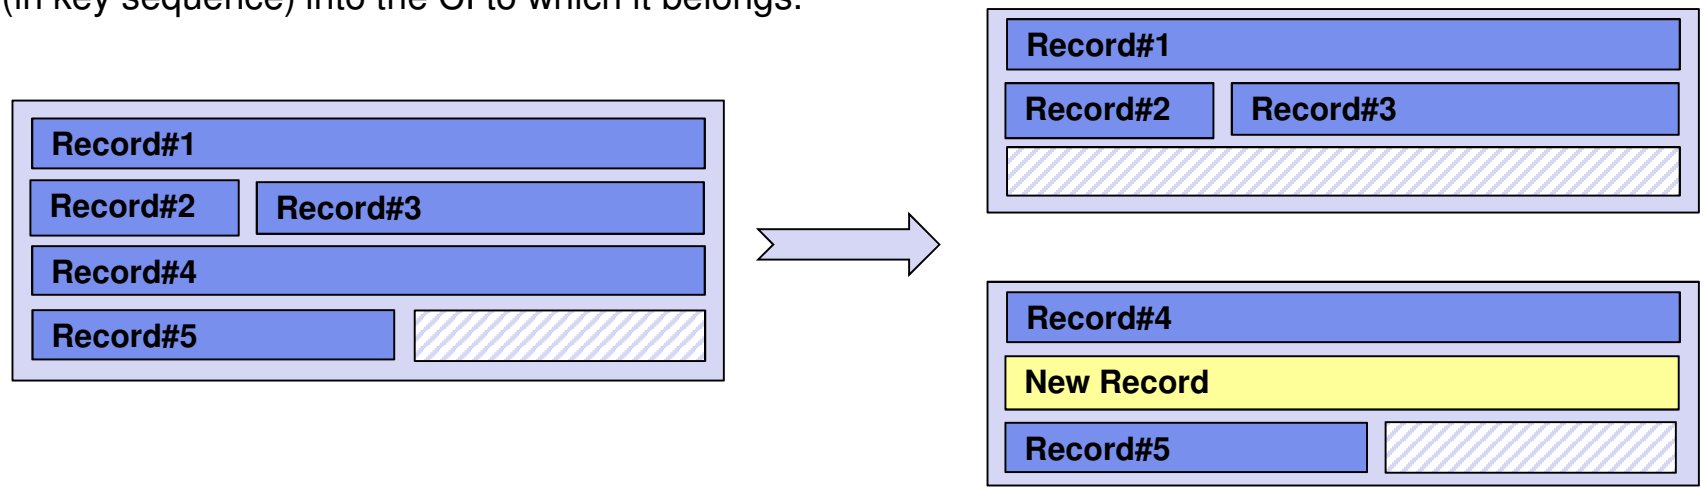

CI Split considerations and recommendations

- CI splits are not very costly in terms of system overhead (four I/Os, a bit of CPU processing overhead).
- Do not specify CI free space. Do not reorganize files just based on CI split numbers.
- If an error occurs (split is interrupted), original CI with Split-in-Process bit is kept.
  - Next time this CI is read, in keyed update mode, the split will be completed.
  - If the access is not keyed and not update, rc x'00' with feedback x'1C' is returned.
  - If the access is not keyed, but get-for-update, rc x'08' with feedback x'9C' is returned.

VSAM Data Organization – Clusters

CA Split process

Half the CIs (those with the higher keys) are moved to the new CA. Insertion then occurs through regular CI split processing, using the newly-created free space CIs

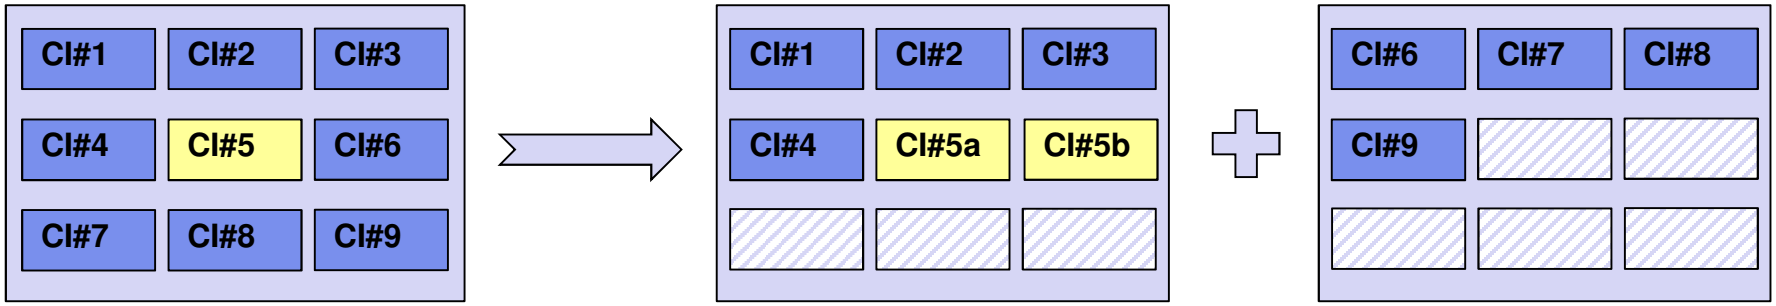

CA Split considerations and recommendations

- CA splits are quite expensive (up to a thousand I/Os) when they occur, but do not substantially impact future processing.
- If inserts are heavily clustered, CA splits may be more efficient than CA Free Space.
- Reorganization will consolidate the cluster, removing free space created by splits, which may have to be added back in.
- Do not reorganize dataset after a certain number of CA splits.
- Define CA free space (at least 20%) for on-line files. FREESPACE(0,20)

VSAM Data Processing

SHAREOPTIONS

SHR (1 x) – Single Open for Output OR Multiple for Input

- VSE/VSAM guarantees read and write integrity

SHR (2 x) – Single Open for Output AND Multiple for Input

- VSE/VSAM achieves fast response by keeping records in buffers and using read-ahead.
- VSE/VSAM guarantees write integrity, not read integrity

SHR (3 x) – No VSE/VSAM control over access

- VSE/VSAM guarantees neither write, nor read integrity.
- Assumes program has own locking mechanism.
- Used for VSE/VSAM catalogs and VSE libraries.
- Should not be used for regular VSAM files.

VSAM Data Processing

SHAREOPTIONS

SHR (4 x) – Multiple programs can access cluster for input and output

- Only within a single VSE system

SHR (4 4) – Multiple programs from multiple VSE systems can access cluster for input and output.

- VSE/VSAM guarantees write integrity, and read integrity(to some extent)
- Read integrity is provided by always reading a record from disk.
- Costs additional operating system overhead. Especially when updating the lock file for shared DASD.

VSAM Data Processing

Local Shared Pool concept

- Shares a buffer pool with all other clusters assigned to the same LSR pool.
- Buffer sub-pools of various size to match the CI Size of the clusters assigned to the pool.
- Possible to define separate LSR pools for Data and Index
- Buffers within a sub-pool are reused on a “least recently used” basis.
- Automatic Pool Definition (CICS).

Better Storage Management

➢ Sharing resources between multiple clusters provides more resources during periods of intense activity for a specific cluster, yet fewer total resources.

Data in Memory

➢ With a large buffer pool, CIs are retained in memory once referenced, and save I/O if they are referenced again.

LSR Pools Monitoring and Tuning

➢ See zVSE LVC – ’*Tuning VSAM file performance under CICS TS*’.

VSAM Data Processing

Dataset Name Sharing (DSN)

- Allows multiple output OPENs for SHR(2) file.
- Shares more than just buffers.
- More efficient usage of storage.
- Buffer sharing, reduce chance of read integrity problems.

Restrictions:

- First file opened determines access mode.
- First file opened must specify sufficient resources (strings, buffers).
- Reusable files are not supported.

Recommendations:

- Use both Dataset Name Sharing and LSR

VSAM Utilities

IDCAMS Utility

IDCAMS is a utility program that is part of VSE/VSAM. It serves to create and maintain VSAM Catalogs and Clusters.

Functions of IDCAMS

- DEFINE - create VSAM objects (Catalog, Space, Cluster, AIX, Path)
- ALTER - modify VSAM objects
- DELETE - delete VSAM objects
- REPRO - copy Cluster data
- BACKUP/RESTORE - backup/restore VSAM objects
- EXPORT/IMPORT - create portable copy of the Cluster
- BLDINDEX - build alternate Index
- PRINT - print VSAM data
- LISTCAT - list Catalog contents

Backup & Restore

Native Backup & Restore facilities:

IDCAMS Backup/Restore:

Fairly quick. Allows restoration of individual clusters. No data reorganization.

IDCAMS REPRO:

Slow. Use for compressed files. Reorganizes data.

IDCAMS EXPORT / IMPORT:

Slow. Compatible with OS/390 | z/OS.

FASTCOPY

Fast. Cannot restore individual clusters. No data reorganization. Must backup all volumes for catalog.

IXFP / Flashcopy:

Extremely fast, Cannot restore individual clusters. No data reorganization. Must backup all volumes for catalog.

Backup & Restore

IDCAMS Backup/Restore

- Saves file contents, catalog definitions, and compression tokens
- Device-independent: Backup to either tape or DASD.
- High-speed backup: faster than REPRO or EXPORT.
- Backup all files from catalog, or selectively via generic list
- Allows files to be selectively restored, or restored to a different catalog.
- Operates at CI-level
- Not compatible with z/OS.
- Compaction (“COMPACT” option)
  - Software compaction of backup data via “COMPACT” option during Backup.
  - More efficient to use hardware
  - Do not use to backup compressed data
- Performance:
  - “BLOCKSIZE(65535)”
  - “BUFFERS(4)” (maximum 8)
  - Multiple concurrent backups very efficient.

Backup & Restore

IDCAMS BACKUP/RESTORE (cont)

- **Return Code x'29':** Warns of files open for update during backup.
  - Identifies potential integrity problems in backup copies.
  - Console message does not identify file. See SYSLST.
  - SHR(4) files use same lock for read and write open. Message says “might”.
- **Multiple Catalog Backup**
  - **“NOREWIND”** parameter
  - Valid only for Standard Labeled tapes.
  - One jobstep for each catalog backed up.
  - User Positioning Required
- **Backup/Restore Cross-ReferenceLists**
  - Scan tapes without restoring any cluster
  - Provide better overview of type and organization of data on tape

Backup & Restore

IDCAMS REPRO

- Process file contents in uncompressed (record) format.
- Restore requires separate cluster definition step.
- Output can be a sequential file (tape or disk), or another VSAM cluster
- Can be used to reorganize cluster
- All compressed files should be backed up in non-compressed format.
- Can also be used to reorganize catalog suffering from “gobi desert” problem.
- Not recommended for catalog-only backup

IDCAMS EXPORT / IMPORT

- Saves file contents in uncompressed (record) format or compressed CI format.
- Also saves catalog definition information
- “SOURCEINHIBIT”
- Can be used to migrate files to/from MVS | OS/390 | z/OS.
  - Use **“BLKSIZE(32767)”** for compatibility

Backup & Restore

IXFP/SnapShot and Flashcopy Support by IDCAMS SNAP

- Adds VSE/VSAM access to IXFP/FlashCopy.
- Fast backup.
- Off-line backup.
- Significantly reduces the time when datasets may not be available to on-line processing.
- Synchronized backup.
- Duplicate volids.
- Steps:
  - **IDCAMS SNAP calls IXFP FlashCopy**
  - **IDCAMS IMPORT CONNECT**
  - **IDCAMS BACKUP SYNONYMLIST(..)**

Backup & Restore

Migration – general scenario

- Backup files using **IDCAMS BACKUP**
- Make a list of all volumes (and extents) managed by old catalog.
- Remove old user catalog from Master Catalog using **IDCAMS EXPORT DISCONNECT**
- Remove catalog extents from VTOC on primary and managed volumes. Use **IKQVDU** or **VSE/DITTO**.
- Define new catalog and space on managed volumes.
- Restore files from backup tapes.
- Do not copy catalogs using **FASTCOPY**, unless old and new volume are absolutely identical.

Backup & Restore

Migration from ECKD to SCSI/FBA

Due to different DASD structure, the Restore can calculate a CISIZE different from the one of the Backup (during dataset definition) which causes a Restore failure :

- **`IDC31337I CANNOT RESTORE <dataset name> WITH SPECIFIED MODIFICATIONS`**

For example Restore of KSDS datasets might fail when they are defined with:

- KEY length > 56, CISIZE less than 4096
- KEY length > 36, CISIZE less than 2048
- KEY length > 6, CISIZE less than 1024

➢ Use REPRO for migrating individual datasets.

➢ Always use REPRO for migrating to SCSI devices from Clusters with SPANNED attribute.

Other VSAM Utilities

IKQVCHK – VSAM Catalog Checker

```
// EXEC IKQVCHK,SIZE=AUTO,PARM='<catalog.name>'
```

- Verifies User Catalog for internal consistency.
- Should be run as part of regularly scheduled maintenance.

IKQVEDA – VSAM SNAP Trace utility

```
// EXEC IKQVEDA,PARM='SYSIPT`
```

See "*VSE/VSAM User's Guide and Application Programming*" for details

- Internal VSAM trace points (SNAP Traces)
- Output to SYSLST (except SNAP001 and SNAP013-15)

IKQVDU – VTOC Management

```
// ASSGN SYS000,DISK,VOL=<volid>,SHR
// UPSI nn
// EXEC IKQVDU,SIZE=AUTO
```

- Manipulates VTOC (delete, define, reset ...)

IKQPRED- Compression Prediction tool

```
// EXEC IKQPRED,PARM='<catalog.name> / <cluster.name>`
```

See "*VSE/VSAM User's Guide and Application Programming*" for details

- Checks an entire catalog or a series of files (generic specification supported).

Thank You

Please forward your questions or remarks to
zvse@de.ibm.com
zaslonko.mikhail@ru.ibm.com

z/VSE Live Virtual Classes

z/VSE @ http://www.ibm.com/zvse/education/

LINUX + z/VM + z/VSE @ http://www.vm.ibm.com/education/lvc/

Read about upcoming LVCs on @ http://twitter.com/IBMzVSE

Join the LVC distribution list by sending a short mail to zvse@de.ibm.com