



<http://twitter.com/IBMzVSE>




VSAM Hints & Tips, Limits, Fundamentals

Speaker Name: Stev Glodowski
Presentation by: Charles E. Olsen

WAVV 2010

Apr 2010



© IBM Corporation 2010

VSE/VSAM Components

- ❖ Virtual Storage Access Method (VSE/VSAM)
 - Two versions: MVS/ESA | S/390 | z/OS and z/VSE.
VM/ESA | z/VM uses VSE/VSAM.
 - Dasd Storage Control
 - Catalog Management
 - Utilities
 - Emulation for SAM and ISAM

- ❖ Catalog management
 - Maintains attributes of all files (clusters) defined to VSE/VSAM
 - Allocates dasd space.

- ❖ Open / Close
 - Connects and Disconnects a cluster with an application program.
 - Ensures access integrity

VSE/VSAM Components

❖ Record Management

- Performs all I/O access to clusters and catalogs
- Manages Buffer Pools
- Ensures SHR(4) integrity

❖ Space Management Feature

- Emulates Sequential Access Method (SAM) for non-VSAM applications.
- Open / Close activity performed by VSE/VSAM.
I/O performed by special VSE Basic Access Method (BAM) phase.
(\$IJGXSrv and its kith and kin)

VSE/VSAM Components

Utilities:

- IDCAMS
 - Catalog Access
 - Data manipulation (REPRO, PRINT)

- Backup / Restore
 - Archival data storage

- IKQVDU
 - VTOC maintenance

- IKQVEDA
 - Trace facility

- IKQVCHK
 - Catalog Corruption Checker

- IKQPRED
 - Compression prediction

VSE/VSAM Data Organization

❖ ESDS (Entry-Sequenced Data Set)

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

❖ KSDS (Key-Sequenced Data Set)

- Access by Sequential Browse, RBA, or Key
- Insert in key sequence.
- Contains data and index components.
- “Gobi desert” problem

❖ RRDS (Relative-Record Data Set)

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

VSE/VSAM Data Organization

- ❖ 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)

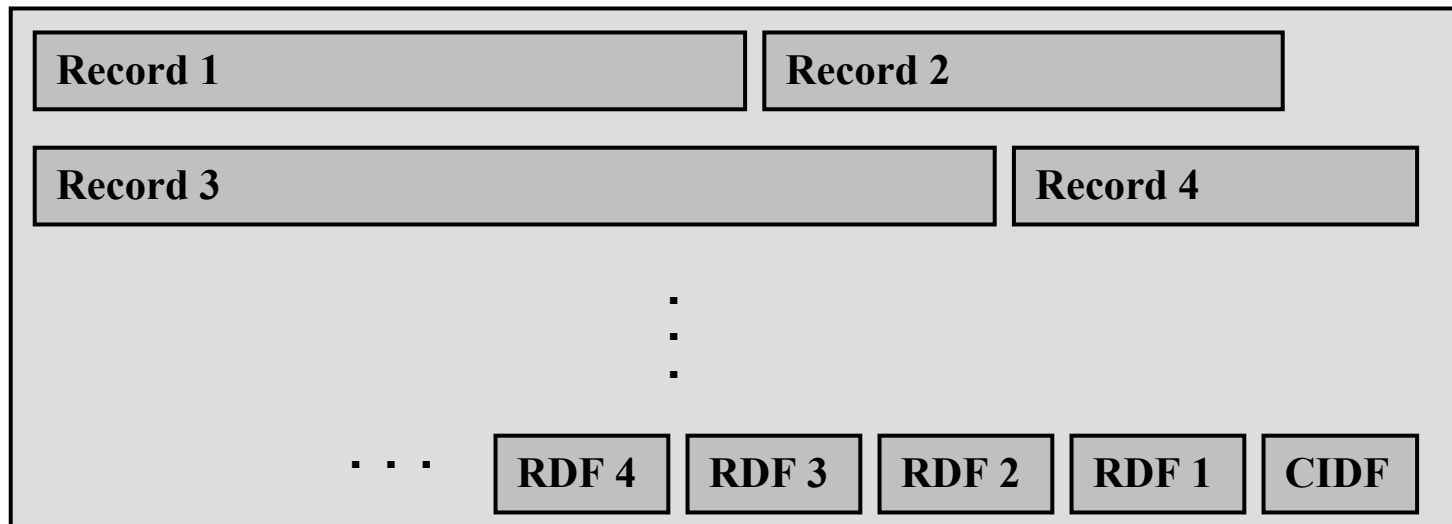
- ❖ Key Ranges
 - Archaic method to isolate I/O requests to different parts of a cluster (file).

- ❖ Cluster Definition:
 - Cluster Name
 - Data Name
 - Index Name (if KSDS or VRDS)
 - Meaningful Cluster names since z/VSE 4.2

VSE/VSAM Data Organization

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 freespace in control interval
- **Record Definition Field (RDF):** If records have different lengths, one (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.
- **Index:** One index record per CA, with a pointer for each control interval, containing the highest index value in CI.

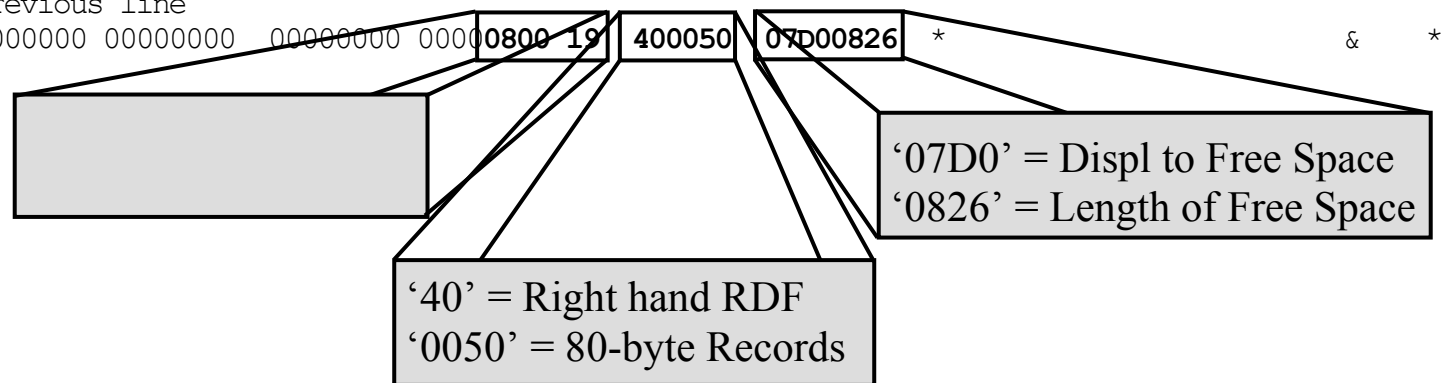
VSE/VSAM Data Organization

Data Control Interval with 25 fixed 80-byte records:

```

000000 F0F0F1F0 F0F1F4F5 F2F7F0F1 F9F4F7F0 F2F2F2F2 F4F3F6F2 F8F5F2F1 F7F24040 *001001452701947022224362852172 *
000020 40404040 40404040 40404040 40404040 40404040 40404040 40404040 40404040 * *
000040 40404040 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 *001071333896944462570802076386 *
0000C0 40404040 40404040 40404040 40404040 40404040 40404040 40404040 40404040 * *
0000E0 40404040 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 *001091677781739323916287100685 *
000160 40404040 40404040 40404040 40404040 40404040 40404040 40404040 40404040 * *
000180 40404040 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 00000000 00000000 00000000 * *

```




VSE/VSAM Data Organization

Data Control Interval with variable length records.

```

000000 F0F0F1F0 F0F1F4F5 F2F7F0F1 F9F4F7F0 F2F2F2F2 F4F3F6F2 F8F5F2F1 F7F24040 *001001452701947022224362852172 *
000020 40404040 40404040 40404040 40404040 40404040 4040F0F0 F1F0F6F9 F1F4F4F6 * 0010691446*
000040 F2F2F7F5 F2F0F1F1 F4F3F1F9 F6F4F9F7 F9F9F7F5 40404040 40404040 40404040 *22752011431964979975 *
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 *96944462570802076386 *
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 C0000036 0F830025 * 5 1 - U c *

```



VSE/VSAM Data Organization

Compression

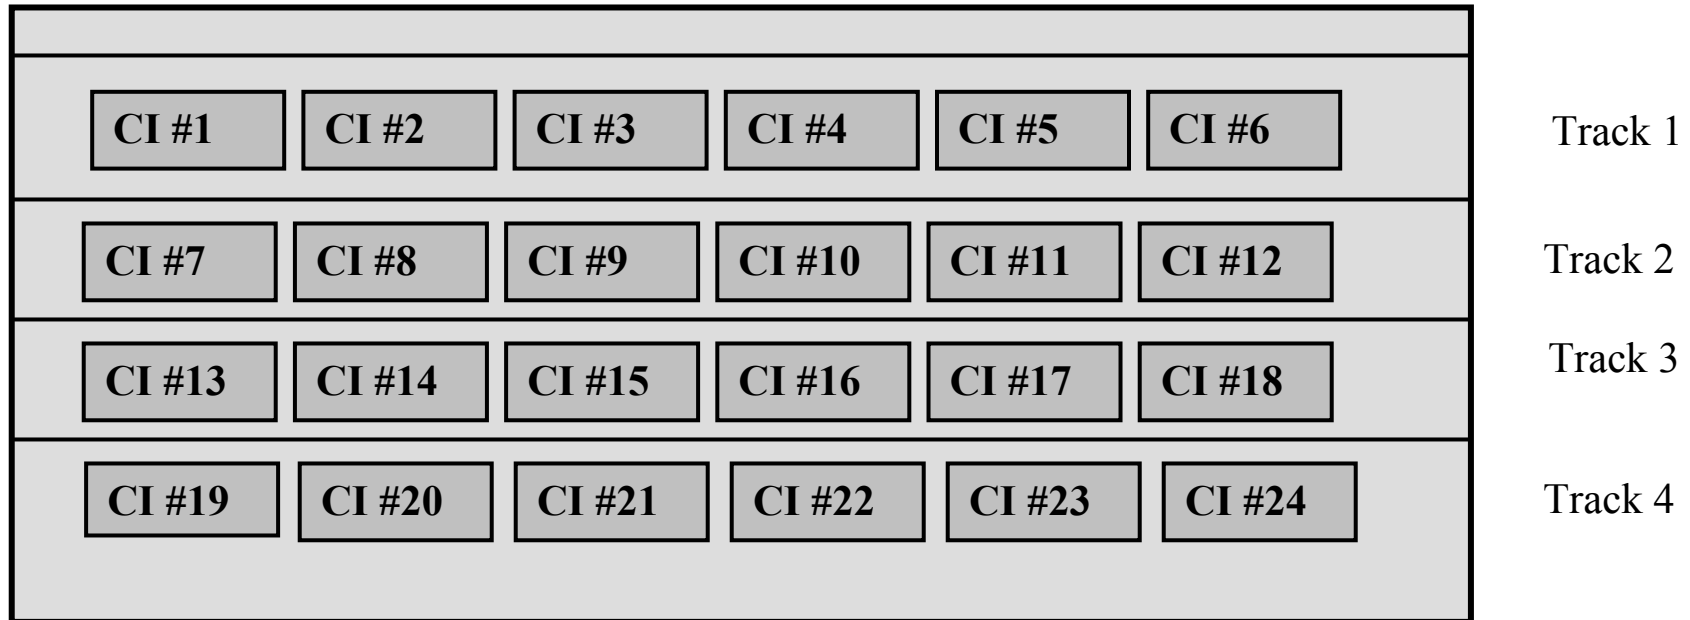
- **Hardware or Software**
- **Dictionary:**
- **Compression Control Dataset (CCDS)**
- **Cluster defined using “COMPRESSED” Attribute.**
- **Advantages:**
 - More data stored on dasd extent. Avoid 4 Giga-byte limit.
 - For sequential access, more records per buffer (CI), so fewer I/Os.
 - Some customers report substantial reductions in batch window.

Flag	Prior to Key	Key	Available for Compression
-------------	---------------------	------------	----------------------------------

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

VSE/VSAM Data Organization

Control Area format



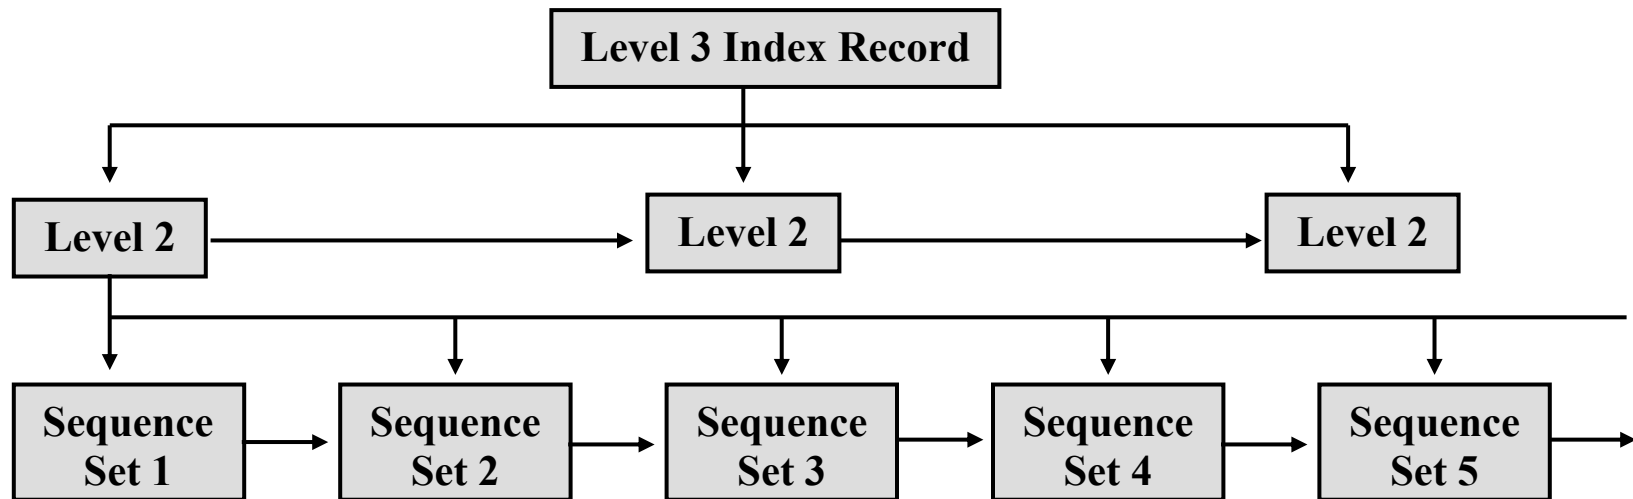
- SAM ESDS (non-CA format. CI can be split at end of track)
- CISIZE vs track utilization (18K best VSAM space utilization, 512k worst case)

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

VSE/VSAM Data Organization

KSDS / VRDS Index Record Structure

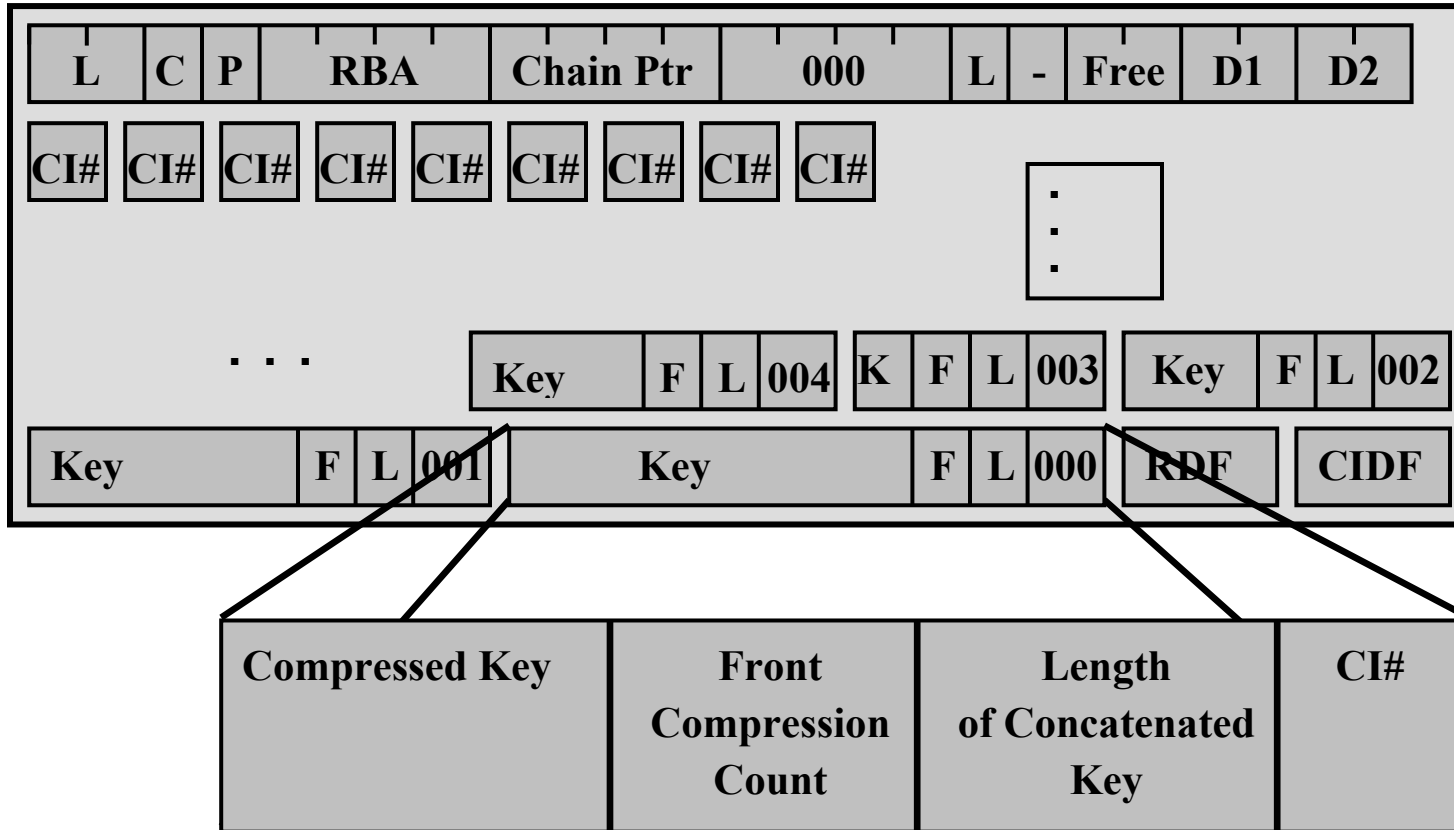
Typical Tree Structure, with horizontal chain:



Each Index Sequence Set describes a data control interval.

VSE/VSAM Data Organization

Format of Sequence Set (Level 1 Index) CI



VSE/VSAM Data Organization

Key Compression in Index Records

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

	<u>Highest Key in Control Interval</u>		<u>F</u>	<u>L</u>	<u>CI</u>
CI#1:	001305263036769318131188297363	'001305'	00	06	00
CI#2:	001562894381711315138840480100	'562'	03	03	01
CI#3:	001760223013561240734555111685	'76'	03	02	02
CI#4:	001949473536319918934062070610	'94'	03	02	03
CI#5:	002124234113651725528014615241	'212'	02	03	04
CI#6:	002327971654539736328740665031	'327'	03	03	05
CI#7:	002601508617068350649999452349	'60'	03	02	06
CI#8:	002867546478217498748460933800	'867'	03	03	07
CI#9:	003074228334141857900076103930	'30'	02	02	08
CI#10:	003390321861192839726291987529	'3903'	03	04	09
CI#11:	003621387672634913128100175623	'621'	03	03	0A
CI#12:	003841434775673551831484414832	'841'	03	03	0B

“F” Front Compression

“L” Length

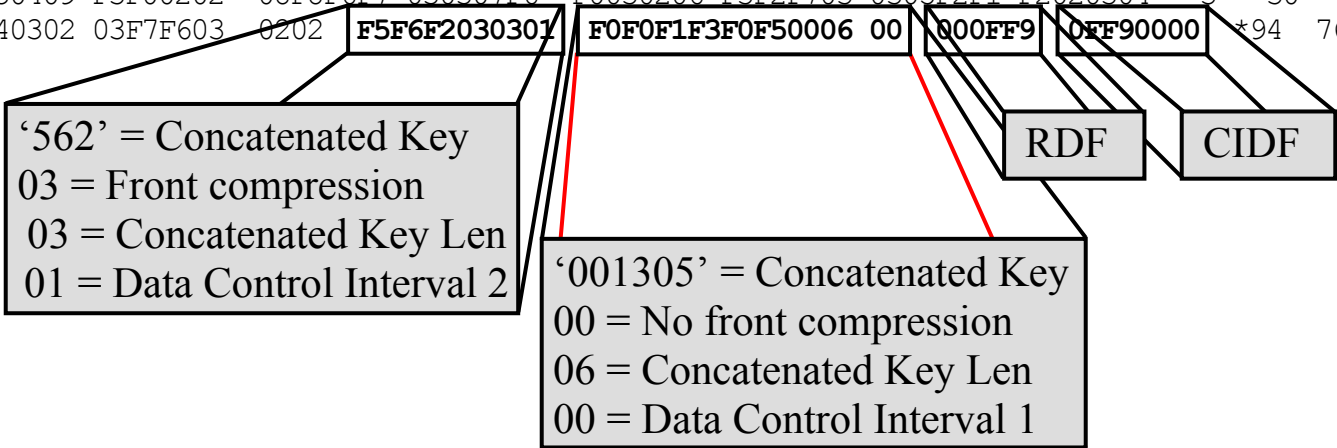
VSE/VSAM Data Organization

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 rqpnmklj *
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 *

```



VSE/VSAM Data Organization

“Gobi Desert” problem

- ❖ Can affect any KSDS file (including VSAM catalogs)
- ❖ Add at end, delete from beginning
- ❖ Index High Key not changed by delete
- ❖ Empty data CIs are never reused
- ❖ Impact:
 - Performance degradation
 - Cluster (catalog) growth
- ❖ Resolution:
 - Define keys (or cluster names) so that they are random
 - Frequent reorganization of file (or catalog).

VSE/VSAM Data Organization

Size Limits: (rc x'1C' "No more available extents")

- ❖ **123 volumes per cluster component (data and index).**
- ❖ **16 volumes from default model.**
- ❖ **4.3 Giga-byte**
 - **4 byte RBA (Relative Byte Address):**
 - **4500 – 11000 cylinders on D/T3390 depending on CI-Size and track utilization**
 - **Compression can help**
 - **'ExtraLargeDataset' or 'XXL' max 286 GB (No-RBA access)**
- ❖ **Clusters are limited to 123 extents per component (Data/Index).**
 - This is normally only a problem if you specify a very small secondary extent.
 - VSE/VSAM will sub-allocate an extent up to 5 times.
- ❖ **Catalogs, reusable files, and unique files are limited to 16 extents.**

VSE/VSAM Data Organization

Size Limits: (cont)

- ❖ **No single allocation over 16Meg records:**
- ❖ **64K control areas limitation for SHR(4)**
 - **Maximum SHR(4) file size is 51 GB**
- ❖ **16 million records per extent:** When defining files using “**RECORDS**”, you can specify up to 16 million records.
 - **“RECORDSIZE”:** You can request more data by specifying a larger maximum “**RECORDSIZE**” for the file.
 - **Compression can help:** Remember, if the file is compressed, VSAM uses the un-compressed maximum record length to calculate how much space to reserve for the file. This may give you more space than you actually need.

VSE/VSAM Data Organization

Large DASD Support

- ❖ BIG-DASD supports dasd with up to 10017 cylinders.
- ❖ FAT-DASD supports dasd with up to 65520 cylinders
- ❖ Implementation is transparent to existing applications and JCL.
Dasd is flagged in LISTCAT as “BIG-3390” or “FAT-3390”
- ❖ Automatic for BIG-DASD using FATDASD parm for DEFINE UCAT & SPACE
- ❖ Allocations converted to **CYLINDERS**
- ❖ Minimum data CISIZE increased to 1024 (depending on key size)
- ❖ BUFSPACE parameter may be increased
- ❖ Note: **RECOVERABLE** catalogs are not supported on “large dasd”.

VSE/VSAM Data Organization

Recommendations:

- ❖ Maximize size of Control Area
- ❖ Use reasonably large data Control Intervals
- ❖ Let index Control Interval size default.
- ❖ Compression will save I/Os, but will cost CPU
- ❖ Additional buffering will save I/Os
 - For sequential processing, use largest possible data CIs, and multiple data buffers.
 - For direct processing, use smallest possible index CIs, and multiple index buffers.

VSE/VSAM Data Organization

CI Split process:

1. Set Split-in-Progress bit in data CI to be split, and write it out.
 2. Move all records higher than record being inserted (sequential vs direct processing) into a new data CI, and write it out.
 3. Update index sequence set and write it out.
 4. Remove C-I-P bit from old CI, and write it out.
 5. If no room for another data CI in this CA, or if sequence set record is full, causes CA split.
- ❖ CI splits not very costly in terms of system overhead (four I/Os, a bit of CPU processing overhead)
 - ❖ Recommendation: Do not specify CI free space. Do not reorganize files just based on CI split numbers.
 - ❖ Between steps 2 and 4, duplicate records exist in the database. If an error occurs (split is interrupted) at this point, CI with s-i-p bit is already in database.
 - 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.

VSE/VSAM Data Organization

CA Split process:

1. Set Split-in-Progress bit in index sequence set, and write it out.
2. Format a new CA at high-used RBA in current extent. This may involve 150 I/Os for 4K data CIs.
3. Read all CIs from old CA. Turn on Split-in-Progress bit and write them back out (up to 300 I/Os).
4. Read all CIs higher than record being inserted (sequential vs direct processing), turn off c-i-p bit
5. Write CIs to new data CA. (up to another 300 I/Os).
6. Create new index sequence set and write it out.
7. Read and write higher-level index records (if required)
8. Read all CIs from old CA, turn off c-i-p bit. Clear, if required. Write back out to old CA (up to another 300 I/Os).
9. Update old index sequence set record, indicate free CIs, turn off c-i-p bit and write out.
10. If no room for another data CA in this extent, allocate a new extent. If no additional extents possible, or if catalog full, reject record update / insert request with feedback code x'1C'. If adding new extent would result in > 4Gig, set feedback code x'D8'

VSE/VSAM Data Organization

CA Splits 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.
- ❖ Consider preformatting file, when new record keys are predictable and are inserted in direct mode.
- ❖ Do not reorganize after a certain number of CA splits.
- ❖ Define CA free space (at least 20%) for on-line files. **FREESPACE(0,20)**
- ❖ Between steps 5 and 8, duplicate records exist in the database. If an error occurs (split is interrupted) at this point, CI with s-i-p bit is already in database.
 - 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 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,
- ❖ Definitions for further user catalogs,
- ❖ Definition for VSE/VSAM –managed libraries (PRD1, PRD2)

User Catalogs:

- ❖ Optional. As many as required
- ❖ Requires JCL specification
- ❖ May be shared by multiple z/VSE systems
- ❖ Only one catalog per volume
- ❖ Can own space on multiple volumes.
- ❖ Multiple catalogs can own space on one volume
(not recommended, except for DOSRES and SYSWK1)

VSAM 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)
- On-line Messages File (IESMSGs)
- VSE/VSAM Record Mapping Definitions (See e-business connectors)
- CICS REXX files (RFSDIR1, RFSPOL1, RFSDIR2, RFSPOL2)
- CICS Listener (EZACONF, EZACACH)

VSAM Catalogs

Catalog contents:


- ❖ Self-describing records (including cluster definitions for catalog itself)
- ❖ Volume (space) definitions
- ❖ Cluster definitions (including data, index, aix, and path)
- ❖ Compression information (CCDS dataset)
- ❖ Internal Format like VSAM KSDS key range file
- ❖ Can be shared with zVM in read-only mode.
Not shareable with z/OS.
- ❖ Three parts:
 - True-name (High-key) range. Contains index of 44-character names to internal catalog CI#
Subject to “Gobi Desert” problem .
 - Low-key range:
 - Volume descriptors
 - Cluster descriptors
 - Index (Used only for True-name records)

VSAM Catalogs

Catalog Recommendations:

- ❖ Name clusters (and catalogs, where applicable) to include application names.
- ❖ Name all cluster components (spec. data and index) explicitly. Volume list in LISTCAT includes data and index, not cluster name.
- ❖ Exploit partition and system independent naming (% or %%) (SAM ESDS files only)
- ❖ Place static (once defined, multi-access) and dynamic files (frequently redefined) in separate catalogs.
- ❖ Place batch vs on-line files in separate catalogs.
- ❖ Define catalog with “**DEDICATE**” at cluster level, and catalog allocation at “**DATA**” level.
- ❖ Do not put all your eggs in one basket.

Don't put all your eggs in one basket !!



Make your clusters happy by giving them space in separate Catalogs



Catalog Recommendations: (cont)

- ❖ Catalogs are limited to 16 extents, and can only expand on original volume.
 - Monitor this closely by checking self-descriptor cluster listing.
See “VSE/VSAM Commands” under “*Appendix A. Interpreting LISTCAT ... Output*”
 - 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.

- ❖ Allocating space to true name range
 - More space will be allocated to the true-name range than the low key range.
 - If the true-name range fills up, and a new allocation is acquired, the entire new allocation is assigned to the true-name range.
 - Control Area will always be two tracks
 - IDCAMS always adjusts your allocation as it sees fit.
 - Index Allocation

Backup & Restore



- ❖ IDCAMS Backup/Restore:
Fairly quick. Allows restoration of individual clusters. No data reorganization
- ❖ IDCAMS REPRO:
Slow. Use for compressed files. Reorganizes data.
- ❖ IDCAMS EXPORT / DISCONNECT: Slow.
- ❖ 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
 - See next page for a circumvention

- ❖ **Backup/Restore Cross-ReferenceLists with z/VSE 4.2**
 - Scan tapes without restoring any cluster

 - Provide better overview of type and organization of data on Tape

Backup & Restore

IDCAMS REPRO:

- ❖ Saves only 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 (COPY/NOCOPY/DDSR with z/VSE 4.2)
 - **IDCAMS IMPORT CONNECT**
 - **IDCAMS BACKUP SYNONYMLIST(..)**

Backup & Restore

Migration:

- ❖ 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



This page
was
deliberately
left blank

Alternate Indexes

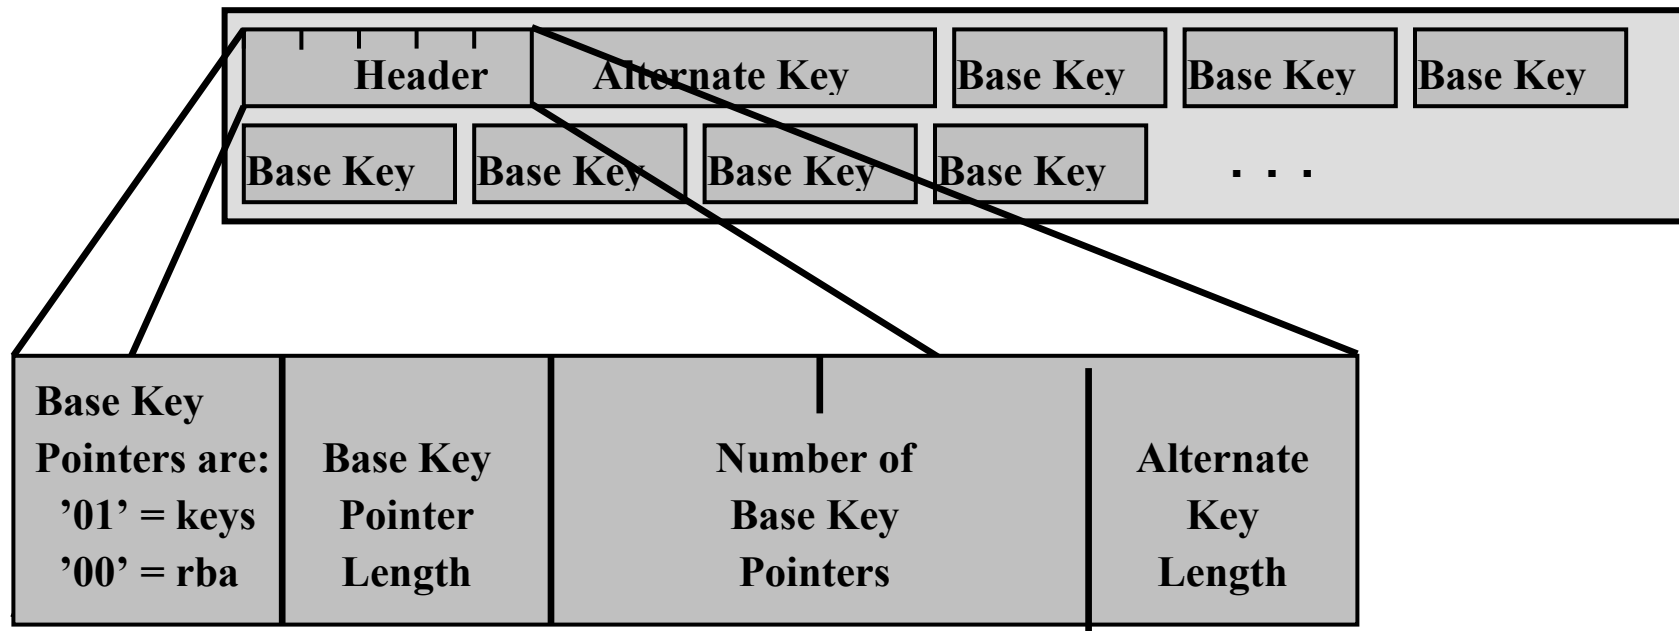
- ❖ Alternate way to access information in KSDS or ESDS file
- ❖ Path Name vs Alternate Index Name
 - Alternate Index is a KSDS file, Path is simply a logical connection between an alternate index and a base cluster.
- ❖ **UNIQUEKEY vs NONUNIQUEKEY**
 - Non-unique keys require definition of spanned records.
 - 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'(0) Duplicate aix key
 - x'28'(8) GETVIS error
 - x'94'(8) Max number of AIX pointers
 - 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

Alternate Indexes

“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)
- ❖ Solution: Write your own build index routine.

Alternate Index Record Format:



SHAREOPTIONS

SHR(1,x) One program opens file for output

-or-

Multiple programs open file for input

- VSE/VSAM guarantees read and write integrity

SHR(2,x) One program opens file for output

-and-

Multiple programs open file 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 read nor write integrity
- Assumes program has own locking mechanism.
- Used for VSE/VSAM catalogs, VSE libraries
- Should not be used for "normal" VSAM files

SHAREOPTIONs

SHR(4,x) Multiple programs can access cluster for either input or output, but only from a single VSE system.

SHR(4,4) Multiple programs from multiple VSE systems can access cluster for either input or output.

- VSAM guarantees write integrity, and some read integrity
- VSAM enhances read integrity by always reading a record from disk
- This additional operating system overhead can be costly
SHR(4,4) is especially costly, because each lock must be reflected to dasd lock file.

In addition, VSE/VSAM controls access from different strings (programs) within the same VSE task (i.e. CICS) using The Hold Block (THB).

SHAREOPTIONs

File Access Contention:

- ❖ SHR(1) or SHR(2): Second program receives x'A8' during OPEN
 - Since z/VSE 4.3 the x'A8' error message will include the task id of the task holding the lock

- ❖ SHR(4) Second program is allowed to open file. When a record update is requested:
 - VSE/VSAM will attempt to lock record
 - If not available, four possibilities:
 - a. VSE sub-task will wait until record is available.

Online:

- b. Return to CICS via Exit List
 - c. Return code x'08..D0'
 - d. Return code x'08..34'
- Return code x'08..14'

Default Models

NOALLOCATION files (only a catalog entry)

❖ Reserved name:

DEFAULT.MODEL.ESDS

DEFAULT.MODEL.ESDS.SAM

DEFAULT.MODEL.KSDS

DEFAULT.MODEL.RRDS

DEFAULT.MODEL.VRDS

DEFAULT.MODEL.AIX

❖ Over-rides system defaults

❖ Useful for defining default volumes

❖ Any file can be used as model for a subsequent IDCAMS DEFINE by coding the MODEL parameter

❖ See write-up in "*VSE/VSAM User's Guide*"

Default Models

Missing DEFAULT.MODEL

The following series of messages appears rather intimidating, but in actuality only mean that you tried to open a SAM ESDS file (new file for output) without a DEFAULT.MODEL defined in the catalog (or perhaps defined a valid one on the // EXTENT card which did not match the candidate volumes in the default model).

```
4A37I FILE SAMFILE CATALOG ERROR DURING IMPLICIT DEFINE - 248, BX, 000
```

(248 indicates "Volume Record not found", "BX" is IGG0CLBX, the VSAM Catalog Mgmt routine that caught the error)

```
4228I FILE SAMFILE OPEN ERROR X'4F' (079) CAT=UCAT ( 8, CG, 6)
      (IKQOPN-3) RC X'000004F' ON CALL TO IKQOPNHC
```

('4F' indicates an error during implicit define. "08..06' indicates a no-record-found condition, in this case, the search for the default model. "CG" is, of course, IGG0CLCG)

```
OV15I REQUEST FROM SYSTEM SERVICE ROUTINE
```

```
OS08I LOG. TRANS. AREA CANCELED, PHASE = $$BOSMXT
```

(\$\$BOSMXT is the SAM interface to VSAM, so he is the one that catches it when something goes bump in the night)

```
OS00I JOB RUN CANCELED.
```

```
1151I DUMP COMPLETE
```

```
1S78I JOB TERMINATED DUE TO PROGRAM ABEND
```

```
E0J RUN
```

SAM-ESDS

- ❖ Expedites migration from dasd sequential files (DTFSD) to native VSAM.
- ❖ Allows programs to access VSAM files without re-write.
- ❖ Allows dynamic allocation of files (i.e. compiler / sort work files)
- ❖ No comparable function under MVS | OS/390 | z/OS.
- ❖ VSAM SAM ESDS files are slightly different (non-CA) from ESDS files.

```
// DLBL SAMFILE, ' THIS. IS. A. SAM. FILE' , , VSAM, CAT=<catalog fname>,      X
      RECORDS=(<primary>, <secondary>), RECSIZE=nn                          X
      DISP=(<open disp>, <close disp>, <error disp>)                          X
```

Note: instead of Records and Recordsize starting with z/VSE 4.3 you can also specify Cylinders(<primary>,<secondary>) or Blocks(<primary>,<secondary>)

SAM-ESDS

OPEN:

- Space Management intercepts OPEN request: Points DTF at location of VSAM cluster. If file does not already exist, it will be dynamically defined. This requires a default model (for volume information) or an EXTENT card following the DLBL, and **RECORDS** and **RECSIZE** parameters on “// DLBL”.
- Some programs require extent to be contiguous
File-id = 'DOS. WORKFILE. SYS . . .'
Special handling for 'DOS.WORKFILE.SYSLNK'
- Disposition Processing: “DISP=(OLD | NEW,...)”
 - If file already exists, “DISP=(NEW,...)” will only reset file, not delete and re-define.
- Place pointer to SSR Service Routine (\$IJGXSrv) into DTF Extension (DTFX)

SAM-ESDS

Record Processing:

- I/O requests branch-and-link (BAL) directly to special BAM module (\$IJGXSrv and its relatives).
- BAM handles I/O: executes EXCP for I/O, reads, deblocks, returns –or- blocks and writes user records into a VSE/VSAM-format Control Interval
- If this extent fills up: SAM sets “Extend mode” into DTF, and issues a new OPEN. VSE/VSAM adds an additional extent (if possible), updates DTF, and returns to BAM.
- If no extension is possible, job is cancelled with: “4250I NO MORE AVAILABLE EXTENTS”

CLOSE:

- Space Management intercepts CLOSE request and resets DTF fields.
- Disposition Processing: DISP=(...,DELETE | KEEP | <date>,...)
-or, if job cancelled: DISP=(..., ..., DELETE | KEEP | <date>)
- No record statistics kept in catalog.

Additional VSAM Programs

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

- Checks catalog for internal consistency, primarily between volume, cluster and true-name descriptor records.
- Customer's should run as part of normal monthly maintenance

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

- Internal trace points (SNAP TRACE) within VSAM
- Output to SYSLST (except SNAP001 and SNAP013 in addition in z/VSE 4.3 SNAP014-016)
- Only use under advisement from Level2

```
// ASSGN SYS000, DISK, VOL=<vol id>, SHR
```

```
// UPSI nn
```

```
// EXEC IKQVDU, SIZE=AUTO
```

- Manipulate VTOC (delete, define, reset ...)
- Only use under advisement from Level2

Additional VSAM Programs

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

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

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

VSAMIO (REXX)

- Allows access of VSE/VSAM data from REXX/VSE.
- Belongs to the host command environment ADDRESS VSE.
- Supports KSDS, ESDS, and RRDS data sets.
- Supported operations are READ, WRITE, DELETE, and UPDATE.
- Can read or write data from/to REXX stem variables.

Trademarks

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

CICS*	FlashCopy	Parallel Sysplex*	WebSphere*
DB2*	GDPS*	System Storage	z/OS*
DFSORT	HyperSwap	System z	z/VM*
DFSMS	IBM*	System z9	z/VSE
DS6000	IBM eServer	System z10	zSeries*
DS8000	IBM logo*	System z10 Business Class	z9
Enterprise Storage Server*	IMS	Tivoli	z10
ESCON*	MQSeries*	TotalStorage*	z10 BC
FICON*	OMEGAMON*	z/VSE / VSE/ESA	z10 EC

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.

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

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

* 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.

Stev Glodowski
stev.glodowski@de.ibm.com



<http://twitter.com/IBMzVSE>

