

VSAM Hints & Tips Update

Stev Glodowski

Email: stev.glodowski@de.ibm.com

WAVV
Covington, Kentucky, USA
April 2012



Trademarks

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.

AGENDA

- ❖ VSAM restrictions overview
- ❖ Space utilization
- ❖ Backup/Restore
- ❖ Compression
- ❖ Do's & Don'ts
- ❖ Hints & Tips z/VSE 4.3 Update
- ❖ Migration and Recoverable Catalogs
- ❖ Service Update




VSAM limits overview



- ❖ 123 Extents per data component & 123 Extents per Index component
- ❖ 123 Volumes per data component & 123 Volumes per Index component
- ❖ UNIQUE, REUSEable files & catalogs only 16 Extents & 16 Volumes
- ❖ 4.3 Gig (X'FFFFFFFF') RBA for ESDS, SAMESDS, RRDS, VRDS and AIX
- ❖ 289 Gig for XXL Datasets but only KSDS data component !!
- ❖ 32K max CISIZE for data component = max Recordsize without SPANNED
- ❖ 8K max CISIZE for index component
- ❖ 16 million records per EXTENT & allocation in RECORDS (important for Backup/Restore)
- ❖ 4999 Cylinders or 8.388.096 blocks prim/sec allocation for CATALOGs (V4.1)

VSAM limits overview



- ❖ 16 Volumes used from DEFAULT-MODEL
- ❖ NO IMBED, NO RECOVERABLE CATALOGs since V3.1

FAT-DASD/BIG-DASD limitations

- ❖ Key length between 7-35 require min data CISIZE of 1024
- ❖ Key length between 36-55 require min data CISIZE of 2048
- ❖ Key length greater then 55 require min data CISIZE of 4096
- ❖ Define Catalog DEDICATE on a BIG- or FAT-DASD only on empty disk with VTOC at first or last Cylinder
- ❖ MIN allocation = 1 Cylinder
- ❖ MIN & MAX CA Size = 1 Cylinder -> SPANNED recordsize limit

DASD Overview

ECKD

MODEL	VSAM Capacity	Bytes/Cylinder	VSAM Classification
3380	max 2655 cylinders	712140	Small DASD
3390 mod 3	max 3339 cylinders	849960	Small DASD
3390 mod 9	max 10017 cylinders	849960	BIG- or FAT-DASD
3390 mod 27 / DS8000	max 65520 cylinders	849960	BIG- or FAT-DASD

SCSI-FBA

MODEL	Capacity in Blocks	Capacity in Bytes	VSAM Classification
Generic FBA	491.520 FBA blocks	2 Gigabyte	FBA
FBA-SCSI	33.553.920 FBA blocks	16 Gigabyte	FBA-SCSI

Space utilization

18K

❖ 18K data CISIZE provides best space utilization

- ❖ Its NOT about how much bytes are wasted inside a CI its about tracks and cylinders !!

SAMPLE: (Any KSDS, ESDS or AIX with nearly 4 gigabyte of data)

A) CISIZE of 512 bytes → space utilization of 24K per track

$4.294.967.295 / 25088$ bytes per track(CI 512) = 171196 tracks = 11413 cylinders

B) CISIZE of 18432(18k) → space utilization of 54K per track

$4.294.967.295 / 55296$ bytes per track(CI 18K) = 77672 tracks = 5178 cylinders

- ❖ same amount of VSAM data in both cases
- ❖ NO effect on the 4 gig limit
- ❖ 115% better space utilization starting with very first track allocated

Note: max 289 gig = 767963 cyl(512) and 348427 cyl(18k)
= (over 6 FATDASDs difference)

512

18K is bigger then 32K

❖ 18K provides max size for SPANNED RECORDsize

- ❖ Max recordsize = 1 CA = 1 Cylinder(ECKD)
or 30720 Blocks(SCSI)
- ❖ AVG/MAX recordsize, CISIZE, DEVICE Type together determine the max recordsize possible

18K CISIZE <<->> $(54k * 1024) * 15$ tracks per cylinder = **829440 bytes max rec**

32K CISIZE <<->> $(48k * 1024) * 15$ tracks per cylinder = **737280 bytes max rec**

32K

18K

512

Backup/Restore

- ❖ Migration from 3390 mod3 to Large-/BIG-/FAT-DASD using BACKUP/RESTORE can cause problems, because keylength-CISIZE dependencies on the larger ECKD devices.
 - ❖ e.g. A cluster with CISIZE of 512 and a keylength of 15 cannot be RESTORED to any BIG-/FAT-DASD
 - ❖ Solution : restore to 3390 mod3 and use EXPORT/IMPORT for migration
- ❖ Restore to a different device type will always be slower than if BACKUP and RESTORE Volume are of the same device characteristics. (this is because remapping will operate on a CI instead of a CA basis)
- ❖ Scan tapes without restoring any cluster (since z/VSE 4.2 using XREF)
- ❖ For SCSI <<-> ECKD , be on the highest 4.3 service level (DY47309)
- ❖ Don't use Fastcopy or DDR or any similar product to migrate your CATALOGS --> *this will cause Catalog corruptions !!*

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

“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).

Don'ts !

- ❖ Never define a SAM-ESDS,UNIQUE or REUSable file with a big primary and a very small secondary allocation and having a long list of candidate volumes.
 - ❖ This will lead to EXTENTs wasted and those Cluster only have 16 Ext.max
- ❖ Don't try to save space with small Catalog secondary allocations
 - ❖ Catalog's can only allocate on one volume, so a big primary allocation will be used only once and never again
- ❖ Not to many AIX with UPGRADE over one BaseCluster
 - ❖ Upgrade-set is bottleneck for processing
- ❖ Do not reorganize after a certain number of CA splits.
- ❖ Do not specify CI free space & Do not reorganize files just based on CI split numbers.

Do's

- ❖ 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.
- ❖ Define CA free space (at least 20%) for on-line files. `freespace(0,20)`
- ❖ Place batch vs on-line files in separate catalogs

VSAM updates in Hints & Tips z/VSE 4.3

❖ **SHOWCB Enhancements**

- ❖ Described in Hints & Tips for z/VSE 4.3
- ❖ Described on z/VSE Homepage for z/VSE 5.1

❖ **CATLG API**

- ❖ Only described in Hints & Tips for z/VSE 4.3 and on z/VSE Homepage
- ❖ No reference in VSAM manuals

z/VSE 4.3 & 5.1 Migration Considerations

- ❖ Beginning with z/VSE 3.1, the ability to define catalogs as recoverable was removed.
- ❖ Beginning with z/VSE 4.3, the ability to extend files in a recoverable catalog or add additional space to a recoverable catalog was removed.
- ❖ You have recoverable catalogs on 4.2 and want to migrate to 4.3 or 5.1 ???
 - ❖ then please first migrate your existing recoverable catalogs to non-recoverable catalogs
- ❖ You migrated already to 4.3 and don't know if you have recoverable catalogs, then please use a LISTCAT and the MILA4VSAM tool to find out and PLEASE contact IBM Support

Recoverable LISTCAT

LISTCAT Attribute „**RECVABLE**” within the Catalog Data Entry

ATTRIBUTES

KEYLEN-----44	AVGLRECL-----505	BUFSPACE-----3072	CISIZE-
RKP-----0	MAXLRECL-----505	EXCPEXIT----- (NULL)	CI/CA-
SHROPTNS(3,3) RECOVERY	SUBALLOC NOERASE	NOCOMPRESS INDEXED	NOWRITECHK
NOREPLICAT UNORDERED	NOREUSE NONSPANNED	RECVABLE	

Recoverable Catalogs Conversion

Use the [APAR DY47322](#) to let VSAM do the conversion for you

- ❖ Catalog will be converted to non-recoverable catalog at the very first open
- ❖ No message is presented indicating the conversion
- ❖ No **RECVABLE** attribute anymore within LISTCAT
- ❖ No indication that this ever was a recoverable catalog except
 - ❖ 1 extra cylinder occupied within the VTOC
- ❖ No **LIMITATION** to work and extend this catalog in z/VSE 5.1

z/VSE 4.3 & 5.1 Migration Considerations

The migration steps are fairly straightforward and consist of:

1. IDCAMS BACKUP (*) a backup of all clusters from the catalog
2. Depending on the volume layouts, you might have to identify which extents on a volume belong to this catalog by using extent information from an IDCAMS LISTCAT SPACE.
 - ❖ If the volume only contains VSAM space belonging to this catalog, you can use IKQVDU to scratch the volume that contains the catalog and all volumes containing its space.
 - ❖ Otherwise, refer to the topic "Re-Build a Catalog" in this manual.
3. IDCAMS EXPORT DISCONNECT to remove the user catalog entry from the master catalog
4. IDCAMS DEFINE USERCATALOG.
5. IDCAMS DEFINE SPACE on all available volumes.
6. IDCAMS RESTORE OBJECTS(*) to restore all clusters to the catalog.

Migration Considerations

Migration of VSAM files using IXFP Flashcopy or VSE/Fast Copy are 2 ways of moving data in a vse environment.

- ❖ All volumes for a particular catalog must be backed up at the same time.
- ❖ If another catalog shares space on the same volume(s) as the catalog being migrated, that data will not be usable on the target volume, unless the other catalog is also backed up and included within the list of SOURCEVOLUMES used by Flashcopy.
- ❖ All VSAM Clusters in the catalog(s) being backed up must have been closed.

Fast Copy Migration Considerations

Additional restriction for Fast Copy

- ❖ Do not attempt to only copy specific cylinder / RBA ranges.
The entire volume must be copied.
- ❖ The source and target volumes of Fast Copy need to be the same type and size.
- ❖ If Dump, Restore, Copy VOLUME is used, all files defined to the VTOC will be processed. The VTOC on the target volume must be at the same location as the source volume.
- ❖ If Dump, Restore, Copy ALL is used, the volume is copied physically, without regard to the VTOC. Therefore, the location of the VTOC is not important, since it will be restored with the rest of the volume.

PTFs z/VSE 4.3 VSAM (02C)

DY47373	UD53821	IDCAMS VERIFY recovery fix
DY47345	UD53794	PDUMP reduction
DY47335	UD53775	24-BIT GETVIS SHORTAGE
DY47322	UD53774	Recoverable Catalogs PTF
DY47290	UD53761	SHOWCB Fails in Case Fields=HALCRBA and MF=E
DY47262	UD53705	ABEND with Messages 4228I X'B4'(180) (OPNC1-15) CATALOG CLUSTER RECORD NOT FOUND(004,AH,002)
DY47309	UD53749	Restore VSAM Data with CISIZE Less than 8K to SCSI/FBA

PTFs z/VSE 5.1 VSAM (51C)

❖ PTF-1 available including several VSAM changes (PTF UD53714)

- RECOVERABLE CATALOGs automatic conversion (DY47322 same as 4.3)
- SHOWCB correction only regarding the new 5.1 fields (DY47290)

Is part of the current PSP BUCKET, available for ordering

❖ Additional APARs:

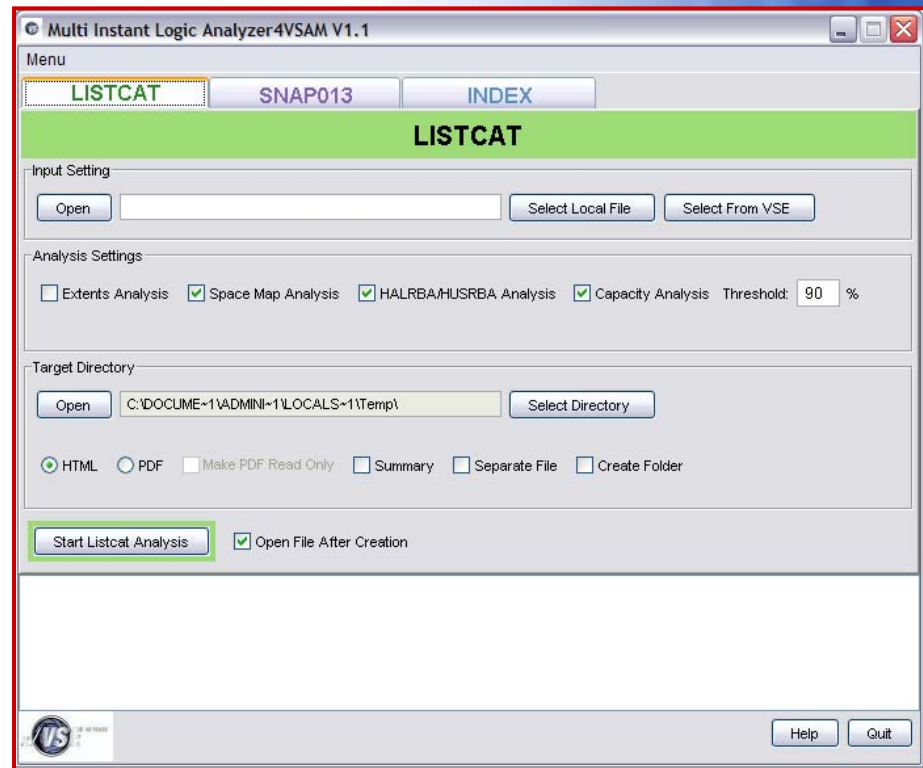
- EXCPAD 31-bit constraint relief
 - PTF UD53775 for z/VSE 4.3
 - PTF UD53719 for z/VSE 5.1
- PDUMP reduction, Extent error correction
 - PTF UD53755
- Backup/Restore SAMESDS
 - PTF UD53813
- VSAM Clusters created in z/VSE 4.x with extreme high numbers of extents
 - PTF UD53820
- IDCAMS VERIFY recovery
 - PTF UD53822

Watch out for future VSAM PTFs at:

<http://www.ibm.com/zvse/support/vsam.html>

Multi Instant Logic Analyzer4VSAM v1.2+

- ❖ What is the **Multi Instant Logic Analyzer4VSAM** ?
 - ❖ A collection of multiple tools to analyze VSAM data instantly
 - ❖ LISTCAT, SNAP013, INDEX and Capacity analysis included
 - ❖ VSE Connector integration
 - ❖ Helps identifying & solving potential problems early
 - ❖ HTML / PDF output



Hints and Tips for z/VSE 4.3



- ❖ **Several updates on VSAM regarding**
 - ❖ Migration
 - ❖ Recoverable Catalogs
 - ❖ SHOWCB
 - ❖ and more

Hints and Tips for z/VSE 4.3

<http://www.ibm.com/zvse/documentation/#hints>



XX03-6006-00

z/VSE Live Virtual Classes (LVC)

- ❖ Join in on z/VSE Online Training
- ❖ Follow IBMzVSE Twitter account and join the LVC mailing list

The screenshot displays a web-based interface for a z/VSE Live Virtual Class. On the left, there is a sidebar with two sections: 'Attendee List (32)' and 'Chat (Everyone)'. The 'Attendee List' section shows a list of participants under 'Hosts (3)' (Ingo Franzki, Siegfried Langer, Stev Glodowski) and 'Participants (29)' (Anson Ngai, Aubrey Hayes, Brandon Richardson, Colin Smith). The 'Chat' section shows messages from Darla Erdmann and Stev Glodowski. The main content area is titled 'zVSE Monitoring - 0322' and features a presentation slide. The slide has the IBM logo in the top right corner and the text 'IBM System z – z/VSE – Live Virtual Class'. The main title of the slide is 'Monitoring Principles & z/VSE Monitoring Options'. Below the title, the names and email addresses of the presenters are listed: Wilhelm Mild (mildw@de.ibm.com) and Ingo Franzki (ifranzki@de.ibm.com). A circular logo with 'z/VSE 40 YEARS' is also present. The slide background is a blue-tinted image of a globe. At the bottom right of the slide, there is a copyright notice: '© 2012 IBM Corporation'.

z/VSE LVCs

<http://www.ibm.com/zvse/education/#completed>

z/VSE & Linux LVCs <http://www.vm.ibm.com/education/lvc/>

z/VSE Requirements

United States [change]

IBM

Home Solutions ▾ Services ▾ Products ▾ Support & downloads ▾ My IBM ▾

Welcome [IBM Sign in] [Register]

IBM Systems > Mainframe servers > Operating systems > z/VSE >

Contact z/VSE

[Send questions or comments](#) [Submit a requirement](#)

Send us your requirement

If you think that a function or feature is missing in VSE, VSE related products, or on this web page, please fill out the form to submit your requirement.


The fields indicated with an asterisk (*) are required to complete this transaction; other fields are optional. If you do not want to provide us with the required information, please use the "Back" button on your browser to return to the previous page, or close the window or browser session that is displaying this page.


Salutation: * (eg:Mr., Ms)

First name: *


Last name: *

We're here to help

 Easy ways to get the answers you need.

 [E-mail us](#)

Stay informed

 [Get the latest news about z/VSE through Twitter](#)

Need help?

[→ Contact IBM](#)

[→ IBM System z frequently asked questions](#)

Thank You

Questions



Please forward your questions or remarks to: stev.glodowski@de.ibm.com

