RIPPLE-TRAC

CONCEPTS AND FACILITIES

April 28, 2016

UNRAVELING THE ENIGMA OF VSAM CATALOG CONCEPTS AND FACILITIES

CUSTOM LISTCAT allows you to re-organize your VSAM files without limiting you to a pre-determined structure (LISTCAT). In fact, with CUSTOM LISTCAT there are no limits. You can create a small project focused result set or grow your result set to hundreds if not thousands of VSAM files with selected components all of which are time stamped for trending analyst.

INSIDE CUSTOM LISTCAT UTILITY (THE METHODOLOGY)

This CUSTOM LISTCAT utility finds its greatest benefit in helping z/OS organizations distinguish high- from low-leverage features or behaviors (concerns) in highly complex systems. In effect, CUSTOM LISTCAT utility lies in seeing through complexity to the underlying structures and relationships.

CUSTOM LISTCAT utility organizes the complexity into a coherent story that illuminates the concerns/patterns and how the concerns/patterns can be decomposed into manageable and comprehensive parts without investing in new hardware or software.

The concept engaged by this CUSTOM LISTCAT utility follows the separation of concerns theory(MDSOC), it presents the information from the point of view of the concern/pattern – you only see information that is related to the concern/pattern and has control on what concern/patterns should be brought to the forefront or obscured into the background. The uniqueness of this approach is the targeted approach, which has the potential to support longer term refactoring of performance issues into highly tuned components and services.

Separation of concerns is a concept that is at the core of software/platform engineering. It refers to the ability to identify, encapsulate, and manipulate those parts of software/platform that are relevant to a particular concern (concept, goal, purpose, etc.).

Concerns are the primary motivation for organizing and decomposing software/platform into manageable and comprehensible parts. Many kinds of concerns may be relevant to different developers in different roles, or at different stages of the software/platform lifecycle. Appropriate separation of concerns has been hypothesized to reduce software/platform complexity and improve comprehensibility; promote traceability; facilitate reuse, non-invasive adaptation, customization, and evolution; and simplify component integration.

Multi-dimensional separation of concerns is an approach to separation of concerns, supporting construction, evolution and integration of software/platform. Its goals are to enable:

- Encapsulation of all kinds of concerns in a software/platform system, simultaneously.
- Overlapping and interacting concerns.

The term ***multi-dimensional separation of concerns*** (MDSOC) refers to flexible and incremental separation, modularization, and integration of software/platform artifacts based on any number of concerns/patterns. It overcomes limitations of existing mechanisms by permitting clean separation of multiple, potentially overlapping and interacting concerns simultaneously.

MDSOC promotes reuse, improves comprehension, reduces the impact of change, eases maintenance and evolution, improves traceability, and opens the door to system refactoring and reengineering.

MDSOC summary:

Involves decomposition of software/platform according to one or more dimensions of concern. A concern is any piece of interest or focus in a platform.

http://en.wikipedia.org/wiki/Separation_of_concerns

The separation allows:

- To allow people to work on individual pieces of the system in isolation;
- To facilitate reusability;
- To ensure the maintainability of a system;
- To add new features easily;
- To enable everyone to better understand the system;
- To allow support for multi-dimensional separation of concerns.

Remember, a dimension of concern is simply an approach to decomposing, organizing, and structuring software/platform according to concerns of a particular kind. This CUSTOM LISTCAT utility falls into the realm of multi-dimensional separation of concerns and can process up to 1200+ concerns/patterns at a time, more than enough to decompose all dimensions of the VSAM catalog.

This CUSTOM LISTCAT utility allows you to select from the concatenated IBM LISTCAT by dataset/group OR owner that gives you the information you want for performance issues or trending analysis without taking a performance hit in the production environment. All fields or just one can be queried for each type of catalog entry across any number of VSAM files.

THE VSAM CATALOG

Most of the information in a LISTCAT comes from the part of an ICF catalog called the VSAM Volume Dataset (VVDS). There is one VVDS for every (logical) DASD volume. The VVDS contains one or more VSAM Volume Records (VVR) for each VSAM dataset on that volume.

An ICF catalog consists of two separate kinds of data sets: a basic catalog structure (BCS); and a VSAM volume data set (VVDS). The BCS can be considered the catalog, whereas the VVDS can be considered an extension of the volume table of contents (VTOC).

The **basic catalog structure** is a VSAM key-sequenced data set. It uses the data set name of entries to store and retrieve data set information. For VSAM data sets, the BCS contains volume, security, ownership, and association information. For non-VSAM data sets, the BCS contains volume, ownership, and association information.

The **VSAM volume data set** is a VSAM entry-sequenced data set. A VVDS resides on every volume which contains a VSAM or SMS-managed data set cataloged in an ICF catalog. It contains the data set characteristics, extent information, and the volume-related information of the VSAM data sets cataloged in the BCS. If you are using the Storage Management Subsystem (SMS), the VVDS also contains data set characteristics and volume-related information for the non-VSAM, SMS-managed data sets on the volume.

The **Volume Table of Contents** and the VTOC index are system data sets which maintain extent and allocation information for a volume. The VTOC is used to find empty space for new allocations and to locate non-VSAM data sets. For all VSAM data sets, and for SMS-managed non-VSAM data sets, the VTOC is used to obtain information not kept in the VVDS.

THE LISTCAT (IBM VERSION)

http://www.ibm.com/support/knowledgecenter/SSLTBW_2.1.0/com.ibm.zos.v2r1.idai200/lcatall.htm

As you can see the IBM LISTCAT listing is extremely complex. Another problem with LISTCAT is timing. Some fields in the VVR are updated when the dataset takes an additional extent, while other information fields are updated only when the dataset is closed. A typical situation where a VSAM dataset is reorganized each night just before it is opened by CICS. If you run a LISTCAT sometime during the day, most of the counters (inserts, retrievals, splits, etc.) will be zero because the dataset has not yet been closed.

Each catalog entry is identified by its type (for example: cluster, non-VSAM, data) and by its entryname.

An entry that has associated entries is immediately followed by the listing of each associated entry. That is, a cluster's data component (and, if the cluster is key-sequenced, its index component) is listed immediately following the cluster.

Each of these groups lists the field names associated with each type of the catalog entries with duplicate names between the groups.

Refer to below example for catalog groups and associated fields; there are approximately 379 fields that are used interchangeably across all groups.

Abbreviations AND Group Names

ALC ALLOCATION GROUP
ASN ASSOCIATIONS GROUP
ATT ATTRIBUTES GROUP
GDG GENERATION DATA GROUP BASE ENTRY, SPECIAL FIELDS
HIS HISTORY GROUP
NVS NON-VSAM ENTRY, SPECIAL FIELD
PRT PROTECTION GROUP
STA STATISTICS GROUP
VLS VOLUMES GROUP

http://publibfp.dhe.ibm.com/epubs/pdf/dgt2i220.pdf
Appendix B. For more information

CAVEATS

First caveat: The EXCP counts reported for files are added up for all tasks accessing that file, not just that individual CICS region. So the EXCP counts may include I/O's from other CICS regions or batch jobs sharing those files.

Second caveat: The reason for flagging statistics for VSAM data sets **(022097427*)** is that if the data set is not closed properly, we do not have a way to determine what the statistics should be. When a VSAM data set is opened the statistics reside in virtual storage and are updated in virtual storage. If the data set is not closed properly we do not know what the updates to the statistics were as the address space they resided in is no longer available. The data set may not have been updated, or updated 1 time, 10 times or 10,000 times, we do not know. This leads to the anomaly where a listcat of the data set shows 100 records, but a repro of the data set will copy 110 records. Also, other programs that copy a VSAM data set flagged as "not properly closed" will copy the correct number of records, but the record counts will not match what the data set looked like prior to the copy. You may determine this by looking to see if the LISTCAT of the data set, prior to the copy, had the "invalid statistics" flag set. (see DATSET FLAGS by DSN example)

SMF CONCERNS

IBM System Management Facility (SMF) is a component of IBM's z/OS for mainframe computers, providing a standardized method for writing out records of activity to a file (or data set to use a z/OS term). SMF provides full "instrumentation" of all baseline activities running on that IBM mainframe operating system, including I/O, network activity, software usage, error conditions, processor utilization, etc.

Using SMF type 64 records does not do the job either. A shortcoming of type 64 records is that they leave out important information about the dataset. For example, there is no information about allocated free space, share options, space allocation or other DEFINE options.

Both LISTCAT and type 64 records contain compression-related statistics for SMS-compressed datasets, but not for datasets compressed by third-party software/platform.

SMF type 42 records provide DASD performance statistics at the dataset level such as connect time, disconnect time and number of cache hits. This information can be useful in assessing the performance of IO to the dataset. However, the primary goal is to eliminate IO, not make it faster, isn't that true?

SMF Type 60 Records

These records are written whenever a VSAM dataset is closed. They essentially contain a snapshot of the datasets VVR at the time the dataset was closed. As such, they provide the advantages of using a VVDS processing program but eliminate the timing issue. For any VSAM dataset, the type 60 record with the largest number of EXCPs can be selected for processing and all other type 60 records for that dataset can be deleted. This provides the equivalent of taking a LISTCAT just before the dataset is reorganized. The drawback is that

type 60 records are not well documented. Neither the SMF manual nor the DFSMS Diagnosis Reference provides detail layouts of VVDS records. For these and other reasons, most installations do not use type 60 records and some installations suppress recording of these records.

Click this link to see….MVS System Management Facilities

http://publibfp.dhe.ibm.com/epubs/pdf/iea2g2c1.pdf

So, the bottom line is:

The ideal time to run CUSTOM LISTCAT is immediately before the dataset is reorganized. Datasets are reorganized at various times during the day or week, so there is never an "ideal" time to run this utility for multiple datasets. We need to determine when a reorg is scheduled and run CUSTOM LISTCAT statistics just before the reorg or when the CICS file is closed.

STRATEGY

The concept is that this analysis will allow data analysis over time, envision solutions, and perform transformation options, select and refine a strategy, and craft a roadmap to achieve that strategy.

How do we do this?

We issue a request to the catalog for a series of files to process. The series of files can be grouped by owner or any combination thereof.

The following graphic representation depicts the process from the VVDS by subsys through the listcat group selection and onto the result set.

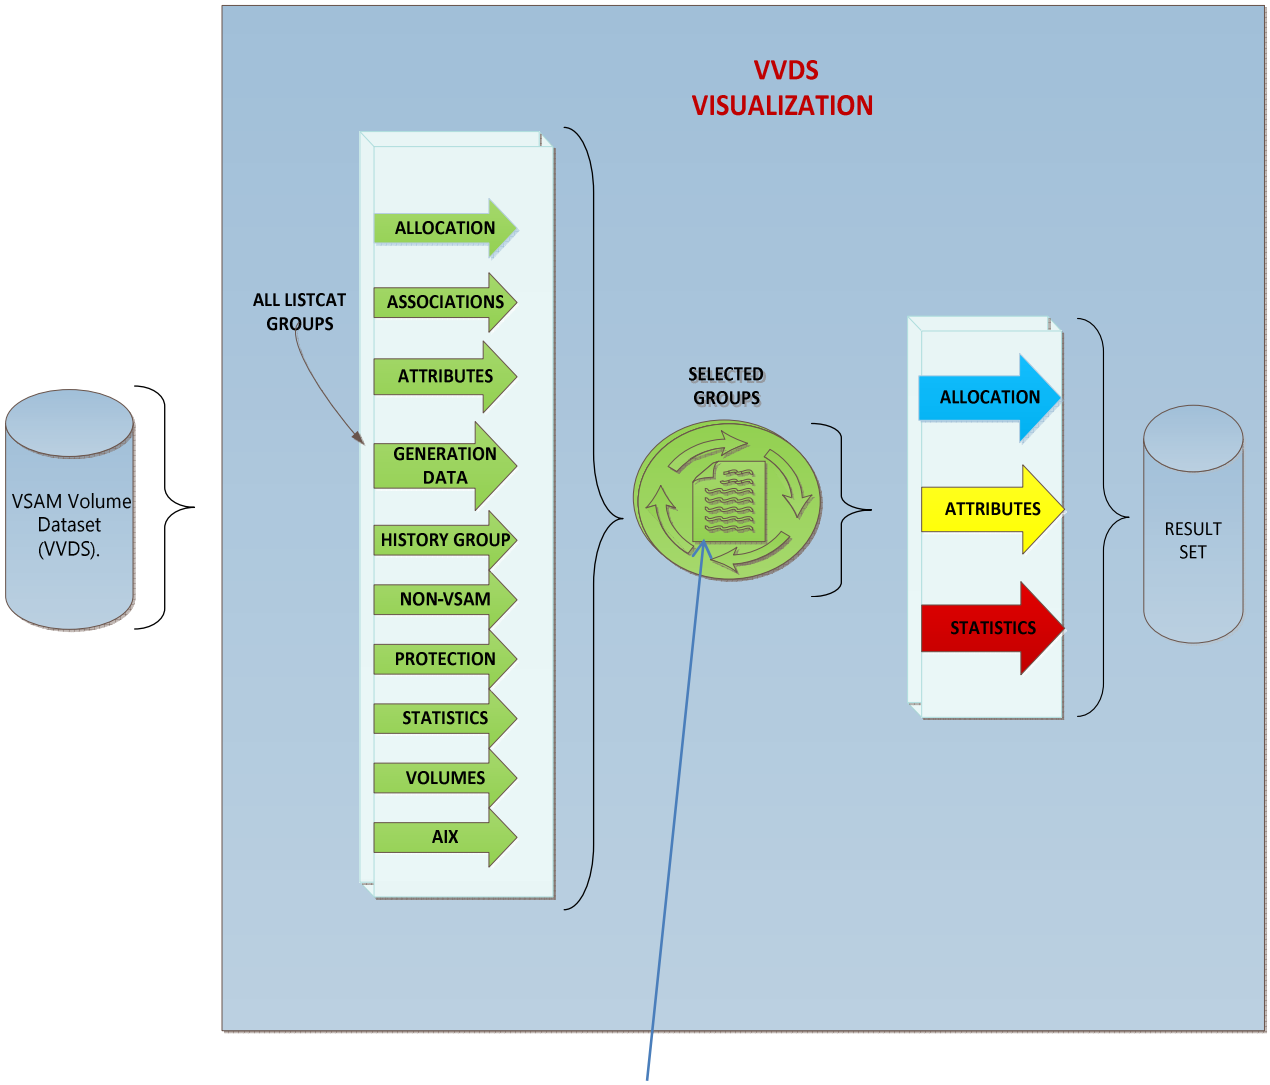

For clarification let's look down through the green sphere to see the progression. We will use the statistics group for this example. (next page) As you look down through the green sphere you can see the statistic details for one file.

Envision the following graphic stacked hundreds or thousands of times, multiple groups with a timestamp for the same file. Hopefully you can now see how each LISTCAT component for the same file compares to previous timelines across the same file or for that matter, across disparate files.

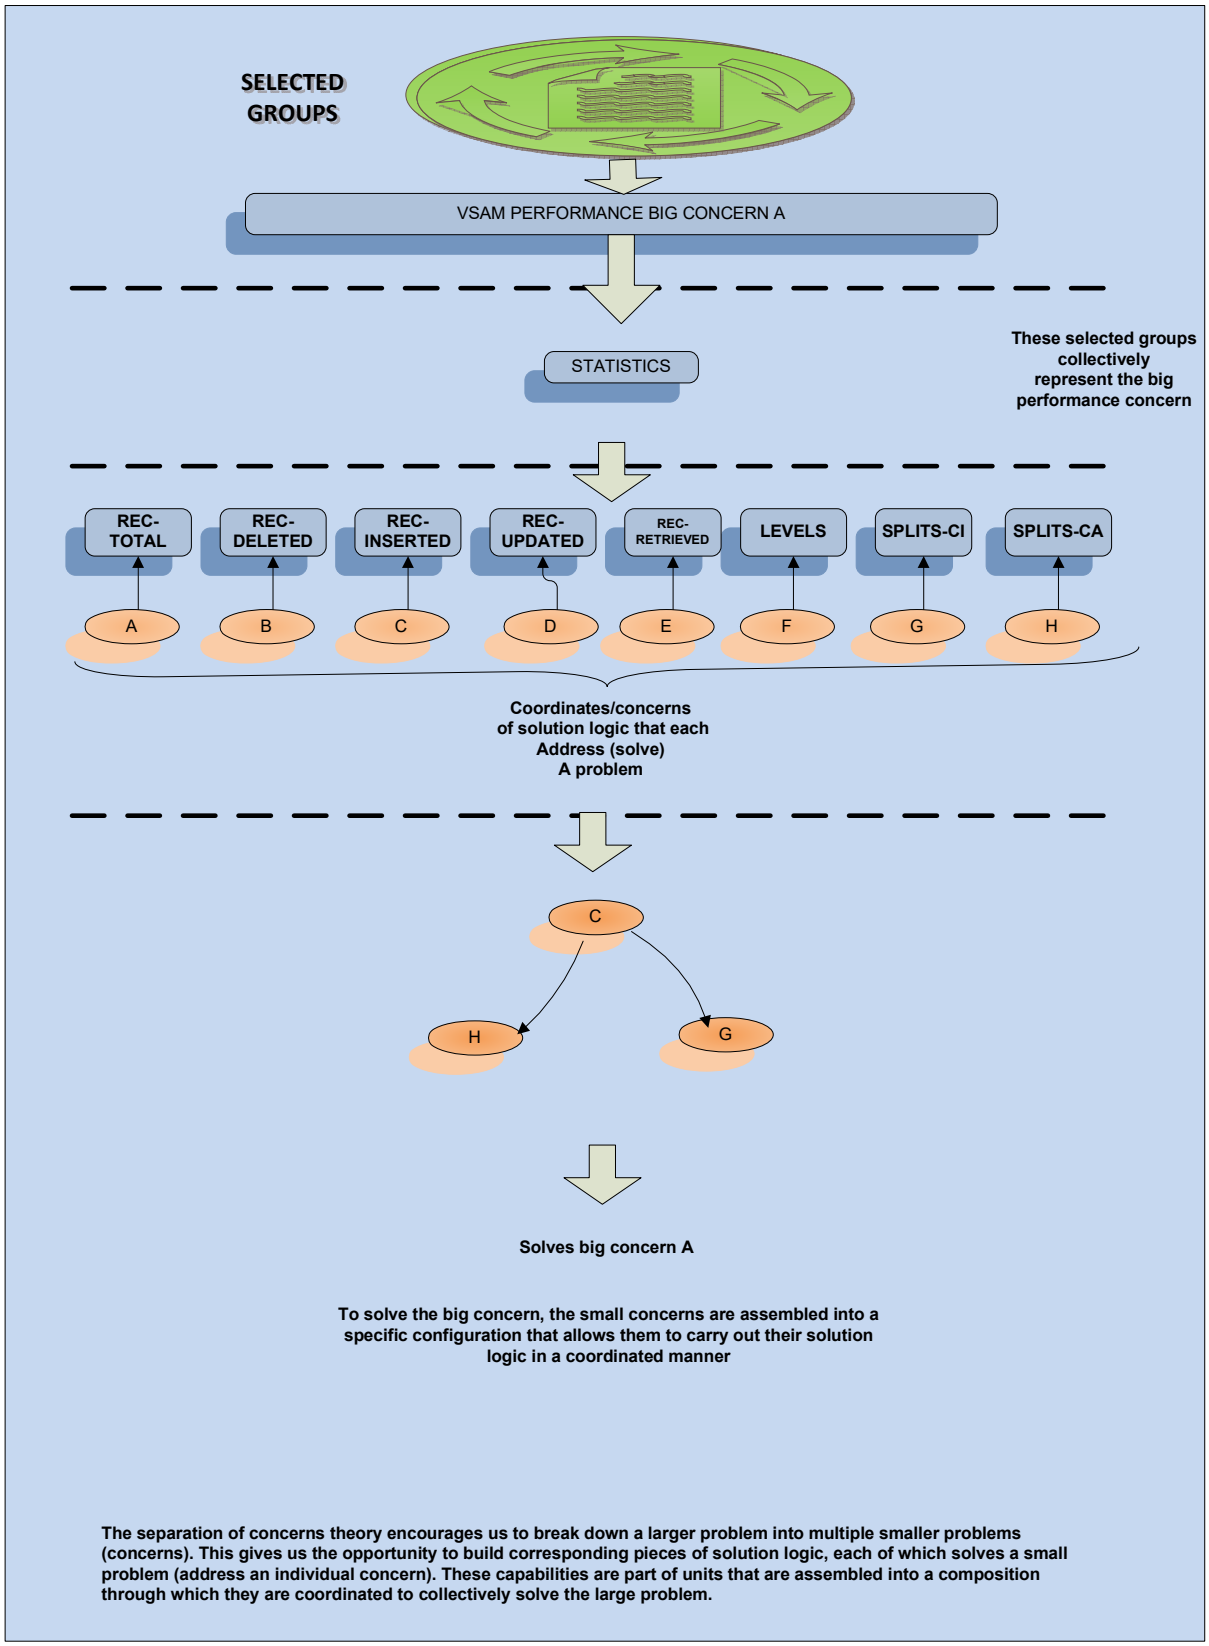
SELECTED GROUPS
VSAM PERFORMANCE BIG CONCERN A
STATISTICS
These selected groups collectively represent the big performance concern
REC-TOTAL
REC-DELETED
REC-INSERTED
REC-UPDATED
REC-RETRIEVED
LEVELS
SPLITS-CI
SPLITS-CA
A
B
C
D
E
F
G
H
Coordinates/concerns of solution logic that each Address (solve) A problem
C
H
G
Solves big concern A
To solve the big concern, the small concerns are assembled into a specific configuration that allows them to carry out their solution logic in a coordinated manner
The separation of concerns theory encourages us to break down a larger problem into multiple smaller problems (concerns). This gives us the opportunity to build corresponding pieces of solution logic, each of which solves a small problem (address an individual concern). These capabilities are part of units that are assembled into a composition through which they are coordinated to collectively solve the large problem.

The potential result: (summary view)

Each of the following examples has extensive use of EXCEL MACROS associated with each to ease the pain of producing the summary views manually.

This is what some of the raw data looks like:

```
AVGLRECL------------1350      BUFSPACE-----------18432      CISIZE
MAXLRECL------------1350      EXCPEXIT----------(NULL)      CI/CA-

UNIQUE              NOERASE   INDEXED        NOWRITECHK     UNORDE
EXT-ADDR

SPLITS-CI--------------0      EXCPS--------------26415
SPLITS-CA--------------0      EXTENTS----------------1
FREESPACE-%CI----------66     SYSTEM-TIMESTAMP:
FREESPACE-%CA----------70          X'CB8B084C2564BC8C'
FREESPC-------5425700864
                               TIME: 07:21:26         04/08/14

HI-A-RBA------7372800000
HI-U-RBA------6491013120
```

The concern: Find potential Buffer issues. View 'data_vs_index_excps'

More index EXCPS than data EXCPS. This generally indicates inadequate buffering for programs that perform random reads or writes to the dataset.

As you can see, the CUSTOM LISTCAT result set summarizes this issue as opposed to the raw data above.

EXCPS by DSN: View 'EXCPS'

A	B	C	D	E	F	G	H	I	J	K
DSN	SE	OMM	TS	IN	CATALOG GROUPS		GROUP NAM		CATALOG COMPONENT	DATE OF EXTRACT
CLUST010	0	DAN		L	*****>SPHERE<*******		0CLUSTER -------		SYS2.LOGIC01.CLUST010	2016031214235185-040
CLUST010	12	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.DATA	2016031214235185-040
CLUST010	31	DAN		L	....STA EXCP		EXCPS-		26415	2016031214235185-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX -------		SYS2.LOGIC01.CLUST010.INDEX	2016031214235185-040
CLUST010	73	DAN		L	....STA EXCP		EXCPS-		53424	2016031214235185-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	2016031214235185-040
CLUST010	14	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.AIX.DAT	2016031214235185-040
CLUST010	36	DAN		L	....STA EXCP		EXCPS-		0	2016031214235185-040
CLUST010	58	DAN		L	*** CLUSTER INDX ***		0 INDEX -------		SYS2.LOGIC01.CLUST010.AIX.IND	2016031214235185-040
CLUST010	76	DAN		L	....STA EXCP		EXCPS-		0	2016031214235185-040
********	*****	*******	***	*	**********	****	**********	****	**********	**********

COL A: DSN IDENTIFIER, THIRD NODE IN DATA SET NAME.
COL B: INTERNAL
COL C: PERSON ASSIGNED TO OR COMMENTS
COL D: NOT USED
COL E: INTERNAL

COL F: CATALOG GROUP…DATA, INDX FOR BASE CLUSTER AND AIX CLUSTER
COL G:NOT USED
COL H:GROUP NAME
COL I: CATALOG COMPONENT…COMPONENT NAME
COL J: DATE OF EXTRACT…USED FOR TRENDING ANALYSIS BY COMPONENT

INDEX LEVEL by DSN: View 'INDEX_LEVEL'

When the KSDS Index is created it is compressed to fit into 1 Index record. As the data portion of the file grows, so does the index. When the index can no longer fit into 1 record a second level of compression is created and a second index record is created. By increasing the size of the index record, the index can remain at 1 level. Level 4 becomes Level 3; Level 3 becomes Level 2, etc. Fewer Index Levels equal less I/O, EXCPS, and CPU.

A	B	C	D	E	F	G	H	I	J	K
DSN	SE	OMM	TS	IN	CATALOG GROUPS		GROUP NAM		CATALOG COMPONENT	DATE OF EXTRACT
CLUST010	0	DAN		L	*****>SPHERE<*******		0CLUSTER --------		SYS2.LOGIC01.CLUST010	20160312142351850-040
CLUST010	29	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.INDEX	20160312142351850-040
CLUST010	71	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	75	DAN		L	....STA INDEX LEVELS		LEVELS-		3	20160312142351850-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	20160312142351850-040
CLUST010	34	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	58	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.AIX.IND	20160312142351850-040
CLUST010	74	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	78	DAN		L	....STA INDEX LEVELS		LEVELS-		2	20160312142351850-040
********	*****	*******	***	*	********************	****	**********	*****	******************************	********************
CLUST001	0	JOEL		L	*****>SPHERE<*******		0CLUSTER --------		SYS2.LOGIC01.CLUST001	20160312142351850-040
CLUST001	28	JOEL			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST001	54	JOEL		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST001.INDEX	20160312142351850-040
CLUST001	70	JOEL			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST001	74	JOEL		L	....STA INDEX LEVELS		LEVELS-		4	20160312142351850-040
********	*****	*******	***	*	********************	****	**********	*****	******************************	********************
CLUST002	0	MARSHALL		L	*****>SPHERE<*******		0CLUSTER --------		SYS2.LOGIC01.CLUST002	20160312142351850-040
CLUST002	28	MARSH			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST002	54	MARSH		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST002.INDEX	20160312142351850-040
CLUST002	70	MARSH			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST002	74	MARSH		L	....STA INDEX LEVELS		LEVELS-		4	20160312142351850-040
********	*****	*******	***	*	********************	****	**********	*****	******************************	********************

RECORD STATISTICS by DATA GROUP, INDEX GROUP BY DSN: view 'STATISTICS'

The fields in this group (STA) give numbers and percentages that tell how much activity has taken place in the processing of a data or index component. The statistics in the catalog are updated when the data set is closed. Therefore, if an error occurs during CLOSE, the statistics might not be valid.

REC-TOTAL—The total number of records actually in the data or index component. This statistic is not maintained when the data set is processed in control interval mode. For a variable-length RRDS, this is the count of slots in the data set.

SPLITS-CI—Control interval splits. Half the data records in a control interval were written into a new control interval and then were deleted from the old control interval.

EXCPS—EXCP (run channel program—SVC 0) macro instructions issued by VSAM against the data or index component.

REC-DELETED—The number of records that have been deleted from the data or index component. Statistics for records deleted are not maintained when the data set is processed in control interval mode.

SPLITS-CA—Control area splits. Half the data records in a control area were written into a new control area and then were deleted from the old control area.

EXTENTS—Extents in the data or index component.

REC-INSERTED—For a key-sequenced data set, the number of records that have been inserted into the data component before the last record; records originally loaded and records added to the end are not included in this statistic.

FREESPACE-%CI—Percentage of space to be left free in a control interval for subsequent processing.

REC-UPDATED—The number of records that have been retrieved for update and rewritten. This value does not reflect those records that were deleted, but a record that is updated and then deleted is counted in the update statistics.

FREESPACE-%CA—Percentage of control intervals to be left free in a control area for subsequent processing.

REC-RETRIEVED—The number of records that have been retrieved from the data or index component, whether for update or not for update. Statistics for records retrieved are not maintained when the data set is processed in control interval mode.

FREESPC—Actual number of bytes of free space in the total amount of space allocated to the data or index component. Free space in partially used control intervals is not included in this statistic. Some of this space may not be accessible due to the current amount of key compression that can be performed in the index.

LEVELS—The number of levels of records in the index. The number is 0 if no records have been loaded into the key-sequenced data set to which the index belongs.

Example of statistics:

A	B	C	D	E	F	G	H	I	J	K
CLUST010	12	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.DATA	20160312142351850-040
CLUST010	29	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	30	DAN		L	....STA REC TOTAL		REC-TOTAL-		0	20160312142351850-040
CLUST010	30	DAN		L	....STA CI SPLITS		SPLITS-CI-		0	20160312142351850-040
CLUST010	31	DAN		L	....STA EXCP		EXCPS-		26415	20160312142351850-040
CLUST010	32	DAN		L	....STA REC DELETED		REC-DELETED-		0	20160312142351850-040
CLUST010	32	DAN		L	....STA CA SPLITS		SPLITS-CA-		0	20160312142351850-040
CLUST010	33	DAN		L	....STA NBR EXTENTS		EXTENTS-		1	20160312142351850-040
CLUST010	34	DAN		L	....STA REC INSERTED		REC-INSERTED-		0	20160312142351850-040
CLUST010	34	DAN		L	....STA CI FREE SPA		FREESPACE-%CI-		10	20160312142351850-040
CLUST010	36	DAN		L	....STA REC UPDATED		REC-UPDATED-		0	20160312142351850-040
CLUST010	36	DAN		L	....STA CA FREE SPA		FREESPACE-%CA-		20	20160312142351850-040
CLUST010	38	DAN		L	....STA REC RETRIEVE		REC-RETRIEVED		0	20160312142351850-040
CLUST010	38	DAN		L	....STA FREESPACE		FREESPC-		737280	20160312142351850-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.INDEX	20160312142351850-040
CLUST010	71	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	72	DAN		L	....STA REC TOTAL		REC-TOTAL-		0	20160312142351850-040
CLUST010	72	DAN		L	....STA CI SPLITS		SPLITS-CI-		0	20160312142351850-040
CLUST010	73	DAN		L	....STA EXCP		EXCPS-		53424	20160312142351850-040
CLUST010	74	DAN		L	....STA REC DELETED		REC-DELETED-		0	20160312142351850-040
CLUST010	74	DAN		L	....STA CA SPLITS		SPLITS-CA-		0	20160312142351850-040
CLUST010	75	DAN		L	....STA INDEX LEVELS		LEVELS-		3	20160312142351850-040
CLUST010	75	DAN		L	....STA NBR EXTENTS		EXTENTS-		1	20160312142351850-040
CLUST010	76	DAN		L	....STA REC INSERTED		REC-INSERTED-		0	20160312142351850-040
CLUST010	76	DAN		L	....STA CI FREE SPA		FREESPACE-%CI-		0	20160312142351850-040
CLUST010	78	DAN		L	....STA REC UPDATED		REC-UPDATED-		0	20160312142351850-040
CLUST010	78	DAN		L	....STA CA FREE SPA		FREESPACE-%CA-		0	20160312142351850-040
CLUST010	80	DAN		L	....STA REC RETRIEVE		REC-RETRIEVED		0	20160312142351850-040
CLUST010	80	DAN		L	....STA FREESPACE		FREESPC-		43008	20160312142351850-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	20160312142351850-040
CLUST010	14	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.AIX.DAT	20160312142351850-040

STATISTICS continued.....

CA/CI SPLITS by DATA GROUP, INDEX GROUP by DSN: view 'CA_CI_SPLITS'

A	B	C	D	E	F	G	H	I	J	K
DSN	SE	OMM	TS	IN	CATALOG GROUPS		GROUP NAM		CATALOG COMPONENT	DATE OF EXTRAC
CLUST010	0	DAN		L	*****>SPHERE<*******		0CLUSTER -------		SYS2.LOGIC01.CLUST010	20160312142351850-040
CLUST010	12	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.DATA	20160312142351850-040
CLUST010	29	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	30	DAN		L	....STA CI SPLITS		SPLITS-CI-		0	20160312142351850-040
CLUST010	32	DAN		L	....STA CA SPLITS		SPLITS-CA-		0	20160312142351850-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.INDEX	20160312142351850-040
CLUST010	71	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	72	DAN		L	....STA CI SPLITS		SPLITS-CI-		0	20160312142351850-040
CLUST010	74	DAN		L	....STA CA SPLITS		SPLITS-CA-		0	20160312142351850-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	20160312142351850-040
CLUST010	14	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.AIX.DAT	20160312142351850-040
CLUST010	34	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	35	DAN		L	....STA CI SPLITS		SPLITS-CI-		0	20160312142351850-040
CLUST010	37	DAN		L	....STA CA SPLITS		SPLITS-CA-		0	20160312142351850-040
CLUST010	58	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.AIX.IND	20160312142351850-040
CLUST010	74	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	75	DAN		L	....STA CI SPLITS		SPLITS-CI-		0	20160312142351850-040
CLUST010	77	DAN		L	....STA CA SPLITS		SPLITS-CA-		0	20160312142351850-040
CLUST010	0	DAN		L	*** CLUSTER PATH ***		0PATH ----------		SYS2.LOGIC01.CLUST010.PATH	20160312142351850-040
********	*****	*******	***	*	**********	****	**********	*****	**********	**********

STATISTICS continued…..

FREESPACE by DATA GROUP, INDEX GROUP by DSN: view 'FREESPACE'

A	B	C	D	E	F	G	H	I	J	K
DSN	SE	OMM	TS	IN	CATALOG GROUPS		GROUP NAM		CATALOG COMPONENT	DATE OF EXTRACT
CLUST010	0	DAN		L	*****>SPHERE<*******		0CLUSTER -------		SYS2.LOGIC01.CLUST010	20160312142351850-040
CLUST010	12	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.DATA	20160312142351850-040
CLUST010	29	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	34	DAN		L	....STA CI FREE SPA		FREESPACE-%CI-		10	20160312142351850-040
CLUST010	36	DAN		L	....STA CA FREE SPA		FREESPACE-%CA-		20	20160312142351850-040
CLUST010	38	DAN		L	....STA FREESPACE		FREESPC-		737280	20160312142351850-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.INDEX	20160312142351850-040
CLUST010	71	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	76	DAN		L	....STA CI FREE SPA		FREESPACE-%CI-		0	20160312142351850-040
CLUST010	78	DAN		L	....STA CA FREE SPA		FREESPACE-%CA-		0	20160312142351850-040
CLUST010	80	DAN		L	....STA FREESPACE		FREESPC-		43008	20160312142351850-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	20160312142351850-040
CLUST010	14	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.AIX.DAT	20160312142351850-040
CLUST010	34	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	39	DAN		L	....STA CI FREE SPA		FREESPACE-%CI-		0	20160312142351850-040
CLUST010	41	DAN		L	....STA CA FREE SPA		FREESPACE-%CA-		0	20160312142351850-040
CLUST010	43	DAN		L	....STA FREESPACE		FREESPC-		55296	20160312142351850-040
CLUST010	58	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.AIX.IND	20160312142351850-040
CLUST010	74	DAN			*** STATS GRP ***				STATISTICS	20160312142351850-040
CLUST010	79	DAN		L	....STA CI FREE SPA		FREESPACE-%CI-		0	20160312142351850-040
CLUST010	81	DAN		L	....STA CA FREE SPA		FREESPACE-%CA-		0	20160312142351850-040
CLUST010	83	DAN		L	....STA FREESPACE		FREESPC-		25088	20160312142351850-040
********	*****	*******	***	*	**********	****	**********	*****	**********	**********

ATTRIBUTES by DSN: view 'ATTRIBUTES'

The fields in this group (ATT) describe the miscellaneous attributes of the entry.

KEYLEN—The length of the key field in a data record, in bytes.

AVGLRECL—The average length of data records, in bytes. AVGLRECL equals MAXLRECL when the records are fixed length. Do not, however, set AVGLRECL equal to MAXLRECL for variable-length relative records.

CISIZE—The size of a control interval, in bytes.

BUFSPACE—The minimum buffer space, in bytes, in virtual storage to be provided by a processing program.

MAXLRECL—The maximum length of data or index records, in bytes. MAXLRECL equals AVGLRECL when the records are fixed length. Do not, however, set MAXLRECL equal to AVGLRECL for variable-length relative records.

Note: For variable-length RRDSs, the MAXLRECL shown in the LISTCAT output is 4 greater than the user-specified length, reflecting the system-increased record size.

RKP—The relative key position:the displacement from the beginning of a data record to its key field.

EXCPEXIT—The name of the object's exception exit routine.

CI/CA—The number of control intervals per control area.

SHROPTNS—(n,m) The numbers n and m identify the types of sharing permitted. See SHAREOPTIONS in the DEFINE CLUSTER section for more details.

A	B	C	D	E	F	G	H	I	J	K
DSN	SE	OMM	TS	IN	CATALOG GROUPS		GROUP NAM		CATALOG COMPONENT	DATE OF EXTRAC
CLUST010	0	DAN		L	*****>SPHERE<*******		0CLUSTER -------		SYS2.LOGIC01.CLUST010	20160312142351850-040
CLUST010	12	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.DATA	20160312142351850-040
CLUST010	21	DAN			*** ATTRIB GRP ***				ATTRIBUTES	20160312142351850-040
CLUST010	22	DAN		L	....ATT KEY LENGHT		KEYLEN-		9	20160312142351850-040
CLUST010	22	DAN		L	....ATT AVG REC LNG		AVGLRECL-		80	20160312142351850-040
CLUST010	23	DAN		L	....ATT CI SIZE		CISIZE-		4096	20160312142351850-040
CLUST010	23	DAN		L	....ATT BUF GENED		BUFSPACE-		10240	20160312142351850-040
CLUST010	24	DAN		L	....ATT MAX REC LNG		MAXLRECL-		80	20160312142351850-040
CLUST010	24	DAN		L	....ATT REL KEY POS		RKP-		30	20160312142351850-040
CLUST010	25	DAN		L	....ATT EXCP EXIT		EXCPEXIT-		NULL	20160312142351850-040
CLUST010	25	DAN		L	....ATT CI'S IN CA		CI/CA--		180	20160312142351850-040
CLUST010	26	DAN		+	....ATT SHARE		SHROPTNS		3,3	20160312142351850-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.INDEX	20160312142351850-040
CLUST010	63	DAN			*** ATTRIB GRP ***				ATTRIBUTES	20160312142351850-040
CLUST010	64	DAN		L	....ATT KEY LENGHT		KEYLEN-		9	20160312142351850-040
CLUST010	64	DAN		L	....ATT AVG REC LNG		AVGLRECL-		0	20160312142351850-040
CLUST010	65	DAN		L	....ATT CI SIZE		CISIZE-		2048	20160312142351850-040
CLUST010	65	DAN		L	....ATT BUF GENED		BUFSPACE-		0	20160312142351850-040
CLUST010	66	DAN		L	....ATT MAX REC LNG		MAXLRECL-		2041	20160312142351850-040
CLUST010	66	DAN		L	....ATT REL KEY POS		RKP-		30	20160312142351850-040
CLUST010	67	DAN		L	....ATT EXCP EXIT		EXCPEXIT-		NULL	20160312142351850-040
CLUST010	67	DAN		L	....ATT CI'S IN CA		CI/CA--		21	20160312142351850-040
CLUST010	68	DAN		+	....ATT SHARE		SHROPTNS		3,3	20160312142351850-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	20160312142351850-040
CLUST010	12	DAN			*** ATTRIB GRP ***				ATTRIBUTES	20160312142351850-040
CLUST010	13	DAN			....ATT AIX UPDATE				UPGRADE	20160312142351850-040
CLUST010	14	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.AIX.DAT	20160312142351850-040
CLUST010	23	DAN			*** ATTRIB GRP ***				ATTRIBUTES	20160312142351850-040
CLUST010	24	DAN		L	....ATT KEY LENGHT		KEYLEN-		30	20160312142351850-040
CLUST010	24	DAN		L	....ATT AVG REC LNG		AVGLRECL-		53	20160312142351850-040
CLUST010	25	DAN		L	....ATT CI SIZE		CISIZE-		18432	20160312142351850-040
CLUST010	25	DAN		L	....ATT BUF GENED		BUFSPACE-		37376	20160312142351850-040
CLUST010	26	DAN		L	....ATT MAX REC LNG		MAXLRECL-		53	20160312142351850-040
CLUST010	26	DAN		L	....ATT REL KEY POS		RKP-		5	20160312142351850-040

SPACE ALLOCATION by DATA GROUP, INDEX GROUP by DSN: view 'ALLOCATION'

The fields in this group (ALC) describe the space allocated to the data or index component defined by the entry.

SPACE-PRI—Gives the number of units (indicated under TYPE) of space allocated to the data or index component when the cluster was defined. This amount of space is to be allocated whenever a data component, a key range within the data component.

SPACE-SEC—Gives the number of units (indicated under TYPE) of space to be allocated whenever a data set (or key range within it) is extended on the same volume.

SPACE-TYPE—Indicates the unit of space allocation:

CYLINDER—Cylinders
KILOBYTE—Kilobytes
MEGABYTE—Megabytes
TRACK—Tracks

A	B	C	D	E	F	G	H	I	J	K
DSN	SE	OMM	TS	IN	CATALOG GROUPS		GROUP NAM		CATALOG COMPONENT	DATE OF EXTRAC
CLUST010	0	DAN		L	*****>SPHERE<*******		0CLUSTER -------		SYS2.LOGIC01.CLUST010	2016031214235185-040
CLUST010	12	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.DATA	2016031214235185-040
CLUST010	39	DAN			*** ALLOC GRP ***				ALLOCATION	2016031214235185-040
CLUST010	40	DAN		L	....ALC SPACE TYPE		SPACE-TYPE-		CYLINDER	2016031214235185-040
CLUST010	41	DAN		L	....ALC SPACE PRI		SPACE-PRI--		1	2016031214235185-040
CLUST010	42	DAN		L	....ALC SPACE SEC		SPACE-SEC--		1	2016031214235185-040
CLUST010	53	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.INDEX	2016031214235185-040
CLUST010	82	DAN			*** ALLOC GRP ***				ALLOCATION	2016031214235185-040
CLUST010	83	DAN		L	....ALC SPACE TYPE		SPACE-TYPE-		TRACK	2016031214235185-040
CLUST010	84	DAN		L	....ALC SPACE PRI		SPACE-PRI--		1	2016031214235185-040
CLUST010	85	DAN		L	....ALC SPACE SEC		SPACE-SEC--		1	2016031214235185-040
CLUST010	0	DAN		L	*** CLUSTER AIX ***		0AIX ----------		SYS2.LOGIC01.CLUST010.AIX	2016031214235185-040
CLUST010	14	DAN		L	*** CLUSTER DATA ***		0 DATA -------		SYS2.LOGIC01.CLUST010.AIX.DAT	2016031214235185-040
CLUST010	44	DAN			*** ALLOC GRP ***				ALLOCATION	2016031214235185-040
CLUST010	45	DAN		L	....ALC SPACE TYPE		SPACE-TYPE-		TRACK	2016031214235185-040
CLUST010	46	DAN		L	....ALC SPACE PRI		SPACE-PRI--		1	2016031214235185-040
CLUST010	47	DAN		L	....ALC SPACE SEC		SPACE-SEC--		1	2016031214235185-040
CLUST010	58	DAN		L	*** CLUSTER INDX ***		0 INDEX ------		SYS2.LOGIC01.CLUST010.AIX.IND	2016031214235185-040
CLUST010	85	DAN			*** ALLOC GRP ***				ALLOCATION	2016031214235185-040
CLUST010	86	DAN		L	....ALC SPACE TYPE		SPACE-TYPE-		TRACK	2016031214235185-040
CLUST010	87	DAN		L	....ALC SPACE PRI		SPACE-PRI--		1	2016031214235185-040
CLUST010	88	DAN		L	....ALC SPACE SEC		SPACE-SEC--		1	2016031214235185-040
********	*****	*******	***	*	********************	****	**********	****	******************************	********************

A possible performance issue based on I/O rate contention with DATA and INDEX on the same volume.

Decrease the I/O rate against the device by avoiding placement of several active data sets on the same volume, (mainly index and data from the same KSDS). If this happens, verify your ACS routines, perhaps by using guaranteed space to force the index in a specific volume. Use different storage groups or the new function DFSMS Data Set Separation announced in DFSMS z/OS 1.3, where you can separate data sets from each other in different physical DASD controllers.