A VSAM/ICF Catalog
White Paper

From Mainstar Software Corporation

MAINSTAR

ICF Catalog Reorg

– Should You or Shouldn't You?

By Ronald K. Ferguson

Preface: *I've always had the strong impression that many people reorg their ICF catalogs far too often. When a survey was conducted recently though, I was pleasantly surprised to find this not to be true. The reason is not because people don't want to reorg their catalogs, bur rather, that most catalogs simply can't be quiesced. Another reason we hear is, "if they aren't broken, don't mess with them", with the corollary statement that everyone is actually afraid to touch their catalogs for fear of uncovering damage that's under the surface. That last statement makes one wonder how so many people could also have the belief that "ICF catalogs don't break". I'm pleased with the conclusion that ICF catalogs aren't reorg'd very often, even if the reasons aren't very good, as the truth of the matter is, there really isn't a performance or DASD space utilization reason to reorg a catalog in most situations, and in fact, they'll perform better the less often you reorg.*

Introduction

Considering how critical ICF catalogs are to the success of your 390 system, you should be thoroughly comfortable with them, knowledgeable of how they work, and confident that you can repair them when problems occur. I can't stress enough that ICF catalogs have lots of 'silent' errors within them, and while they might be running OK right now, if you aren't keeping on top of these errors and fixing them, they'll come back to haunt you at the worst possible time.

This White Paper explains how to use the MAPBCS feature of **Mainstar: VSAM Manager** to see what is happening inside a typical BCS. (The MAPBCS command will soon become a standard feature within our **Catalog RecoveryPlus** product.) We think the MAPBCS feature, together with other diagnostic tools available in **Catalog RecoveryPlus**, will be very useful in helping you to understand more about how your ICF catalogs behave, and make you less afraid of 'touching' your ICF catalogs.

The following discussion will refer to the **VSAM Manager** reports provided at the end of this White Paper. It will be much easier to follow along if you print out the report examples to refer to as each part of the report is explained.

Each report is discussed under a separate heading and the example report is referred to by its report number, for example, Report <1>.

Note: You'll see that the catalog detailed in the reports is on a 3380 device, which many people will undoubtedly laugh at, but interestingly, there are still quite a few of these around, and besides, it makes no difference to the point of this White Paper. This report was produced on 06 April, 2001, so you can see that I'm not using some old dusty listing that's been dredged from the archives, but rather, quite a new listing.

The MAPBCS Command

The sample **VSAM Manager** reports provided at the end of this paper were produced with a MAPBCS command that is very easy to code, as shown below.

```
MAPBCS (NAME(ICFCAT.UCAT09))
```

A similar command, MAP, is used to map everyday, garden-variety KSDS files.

The Catalog Listing Report <1>

Report <1> shows the Catalog Listing, which intentionally looks very much like a standard IDCAMS LISTCAT, but with much improved LISTCAT field contents, which makes it easier to read and use.

For example:

- The DEVICE-TYPE field actually says 3380-3, rather than the unfathomable 3010200E that you see in a LISTCAT (which is the internal MVS UCB code for a 3380; whereas 3010200F is the UCB code for a 3390 geometry device).
- In the CYLS ALLOC field, rather than the excruciating individual tracks allocated that you had to add up yourself, you'll see that this has been gathered up in a single line of information that tells you how many total cylinders have been allocated, and if there are multiple extents, what the smallest and largest physically-allocated extent size is.
- Little things, such as comma separators in the 8-10 digit RBA values, make them much easier to read (although, back-end reporting tools, such as SAS or REXX won't like these as much).

None of these fields when taken alone are revolutionary, but as a whole, they produce a more easily readable and usable catalog listing, as well as one that is significantly shorter in length than your typical multi-extent and multi-volume dataset LISTCAT.

The sample ***Report <1>*** Catalog Listing contains several unique things specifically about ICF catalogs that you should be aware of:

- All of the record counts for catalogs in the STATISTICS paragraph are useless, although there is always an unexplained garbage value in REC-DELETED - just ignore them.
- The CI and CA split counts are fairly accurate, as we'll see from our CA mapping report.
- The EXTENTS field indicates this BCS is currently into 18 total extents, and with an allocation of CYL(40 1), has a total allocated size of 57 cylinders.
- Both the data and index have a CISZ of 4096, which is perfect for a BCS - our survey analysis of key compression for the BCS' keys indicates that using 4096 for ***both*** the data and index CISZ never results in dead CIs.
- The FREESPACE(20 20) specification is essentially worthless, as we'll see from the CA mapping report, but it's also pretty benign. My choice for virtually all catalogs would be FREESPACE(0 0), and let CI and CA splits find their own level of free space as record inserts are made to the BCS - which is really the whole point of this White Paper.

Important note about IMBED & REPLICATE:

As you can see from the Data Component ATTRIBUTES paragraph in ***Report <1>***, this BCS has the IMBED attribute still in effect. Look down at the Index Component ATTRIBUTES paragraph, and you'll see that it also has the REPLICATE attribute in effect. Since this catalog resides on a 3380, odds are that the device is cached, and it's because of cache performance loss caused by record replication, that IBM has now removed both of these parameters from the DEFINE USERCATALOG command. Until a catalog is reorg'd though, both of these attributes will remain in effect, with the attendant performance loss to the cache on the DASD controller. For this reason alone, I would recommend that this catalog be reorganized as soon as possible, with the IMBED and REPLICATE keywords removed (they will be ignored if they are not removed). Refer to ***Report <6>***, the Recommendations Report, and you'll see a recommendation to remove the IMBED and REPLICATE attributes.

The Statistical Summary Report <5>

Report <5>, the Statistical Summary, is a very good extension of the Catalog Listing. It contains many of the same fields, although formatted differently, and contain quite a few analysis topics that might be of interest. To anyone who is familiar with a LISTCAT, the fields are, for the most part, self-explanatory, so they won't be discussed in any depth in this paper. A couple of fields though, might be of interest.

In the DATA COMPONENT CI FREESPACE ANALYSIS paragraph, there is a note that this is "based on a random sampling of the data records". Whenever possible (i.e., when the file is not extended format VSAM), the MAPBCS command 'drops down' into each CA of the file and, by default, samples three CIs within each CA to examine the internal contents. The purpose is to measure how well the file's records fit into the assigned CI size, and also, to determine record length statistics. ***Report <5>*** shows that in this catalog, the sampling was for a total of 2,587 records, finding a minimum record size of 73 bytes, an average record size of 106 bytes, a maximum record size of 1,708 bytes, and a standard deviation of 73 (meaning that 95% of the sampled records were within 73 bytes on either side of 106 bytes). That's nice to know, as you can see from the Catalog Listing (***Report <1>***) that the catalog was (and should be) defined with

the default RECORDSIZE(4086 32400), leaving you without any idea what the actual record lengths are. These are fairly standard record length numbers for a user catalog (particularly one without large GDGs in it), whereas master catalogs tend to have totally different record lengths.

The field titled TOTAL DEAD CIs is on the very last line of ***Report <5>***, with a note indicating that the DEAD CIs are caused by 'BAD INDEX CISZ', and the value shown is 0. This is very good, and is the result of this catalog having been defined with CISZ(4096) for both components. If your catalog(s) have any other CISZs, you should really run this MAPBCS report on each of your user catalogs, and you'll very likely find DEAD CIs. (This is the topic of a White Paper, *Managing ICF Catalog Record Activity* (found on the Mainstar website), which also covers some basic analysis of the MAPBCS report, but using a catalog that behaves very differently from the one in this paper. It's a good idea to read that discussion too.)

The Control Area Analysis Report <2>

Report <2>, the Control Area Analysis, is the 'heart' of **VSAM Manager**. This report produces a CA-by-CA 'map' of the entire dataset, with one line of information per CA. In the first column, you'll see a simple CA numbering sequence, starting from the file's CA #0, right through the last CA in the file, in this case, CA #56 - which means this catalog currently has a total of 57 CAs.

A primary feature of the program logic in this report is its ability to map CA splits, identifying each split with the CA-SPLIT note. In the second column, at the beginning of the report, notice that CA #0 is followed by #55, which means there has very recently been a CA split from that location. I would say the split was recent, as CA #55 is very close to the end of the file, and therefore, was the next-to-last CA split to occur on this file.

Now look to the right-half of the report at that same location, and you'll see that the approximate dataset name key for that CA is EUQS.BD1A.VSO2.W3W.TAPE.EP1. That most likely means a recent define (allocation) of a dataset with a name close to that value was inserted into this catalog, causing the CA split at this location.

Again looking at CAs #0 and #55, note that they have 69 and 70 free CI's, respectively, which is half of the total CI/CA count of 140. This denotes an SIS-mode 50-50 split, giving each of these CAs plenty of free space for record insertions as additional datasets are defined at this dataset name area in future.

Next, note that CA #11 is the third logical CA in the file (numbered logical CA #2 in the report because the report starts at CA #0), and there is a CA-SPLIT annotation at that location also. This is followed by CA #53, and after that, the CA # sequence gets back to #1. There's a notation out to the right, with a CA split count of 3. What this indicates is:

- Long ago on this catalog, CA #0 first split out to #11.
- Dataset define activity in CA #0 continued at a fairly high rate, causing CI splits until the CA filled up again, at which time it split out to CA #55 (this must have occurred over quite a long period of time, as all of the CAs between #11 and 55 were involved in 'other' CA splits in the interim period.
- Dataset activity was continuing at a high rate in the new CA #11 too, as it split out to CA #53 (actually, this split occurred prior to the split from CA #0 mentioned above, as #53 is before #55).

The conclusion from this analysis is: Dataset define activity for names starting with EUQS is continuing over a long time period, proving EUQS to be a fairly active high-level qualifier. The CA splits (caused by CI splits) are doing exactly what they are designed to do - namely, dynamically create free space at exactly the location it is required. For this little section of the catalog, it has now created lots of free space where there is obviously a lot of record insertion activity going on. If we were to reorg this catalog at this point, we'd likely have to pay for the splits in this area of the catalog to be done all over again.

A similar situation has occurred between physical CA #1 and #2,where two CA splits have occurred as a result of inserting records during define operations for dataset names that start with SBW. In this instance, CA #1 split out to CA #12, which in turn then split out to CA #34. Again, the fact that there is such a wide numbering difference between #12 and #34 indicates that this insertion activity is occurring over a long time period - the reason this is mentioned is that history often repeats itself, so areas that have had activity in the past are likely to have activity in the future - both dataset name areas we've looked at so far would seem to fit this category.

At CA #2, the record situation is a bit more 'intense'. Run your eye down the CA-SPLIT notes at that point, until you come to the total split count number, and you'll find that 22 (!!!) CA splits have occurred at this area of the file. Tracing the specific sequence of events here isn't of any particular value, so suffice it to say that this is an area of the dataset naming convention where there is much more activity with new datasets being defined (you might note that it's mostly SBW dataset names, which is the same as in the area described above, with a small number of SFN dataset records near the end of this sequence).

If you follow the report to its end, you'll see that there is more of the same pattern, with each original CA in the file having gone through multiple CA splits as new datasets are defined in it. What are the good and bad points about this? Here are some thoughts:

- It's impossible to know the specific times when these CA splits occurred, but it is known from the CREATION DATE (04 August, 1996) that the catalog is just under 5 years (or 56 months) since reorg at the time of this report. If we look at this simply from an averaging perspective, that's less than 1 CA split a month - costing maybe one elapsed second each - hardly what one could call a performance problem. At this point, the splits are done and paid for, and they have dynamically created lots and lots of free space - at the exact locations where it was needed. This would have been impossible any other way, as the FREESPACE parameter always creates evenly distributed space within the file to allow record insertions, and we've see that wasn't the case here at all.
- If we were to reorg this catalog right now - particularly if the purpose of the reorg was to 'clean it up', it would produce exactly the opposite effect. Record insertions would likely occur in many (maybe all) of the same locations in the future, and these same areas would have to be split again. If we reorg'd to reduce the DASD allocation of the catalog, we might scrunch it down to half its current size, and when the CA splits occur again, the catalog would just grow back to somewhat this same size - a self-defeating effort.
- The really nice thing about the Control Area Analysis report is that it helps to prove these points to you - without it, you're driving with your headlights off, making wild assumptions about what's happening inside a KSDS (and the catalog is a KSDS).

Conclusion

A good conclusion is that this catalog has found it's optimum level of free space, where it now has the appropriate amount at the dataset name locations where the most frequent defines are occurring. Put another way, from the standpoint of VSAM record processing, this catalog is performing at the absolute best performance that we could ask for. Put still another way, the overhead cost of everything that's been done inside this catalog is really very small, the CA splits are costing in the neighborhood of one elapsed second of overhead time each month, and the catalog is costing a modest 57 cylinders of DASD - if we desired to optimize catalog performance, an area that would be much more valuable is concentrating on Catalog Address Space (CAS) lookaside buffering (which will be the topic of an upcoming White Paper). Except for the IMBED and REPLICATE attributes mentioned earlier, it is really hard to find plausible reasons to reorg a catalog. So, if you aren't reorging them, well, it turns out you're doing the right thing - it might not be for all the right reasons, but at least it's the right thing.

Ronald K. Ferguson - is Founder, President & CEO of Mainstar Software Corporation

Ron Ferguson has a technical background in large-scale OS/390 systems. As a software instructor for 20+ years, he has presented over 600 courses on VSAM and ICF catalogs, and is recognized worldwide as an expert in these areas. Ferguson travels widely, meeting with customers and presenting at national and international conferences.

Report <1>: The Catalog Listing

```
*****************************************************************************************************
  MAPBCS (NAME(ICFCAT.UCAT09))
*****************************************************************************************************


         Copyright 1989, 1998 by Mainstar Software Corporation, Bellevue, Washington, USA.
 <V4R2.01>                              V S A M    M A N A G E R           06 Apr 2001  14.00  Page 1
                                     C A T A L O G   L I S T I N G  <1>
USERCATALOG---ICFCAT.UCAT09
  FROM CAT---ICFCAT.UCAT09
  HISTORY
     DATASET OWNER----(NULL)  CREATION----04 AUG 1996  EXPIRATION-------(NONE)
  COMPONENTS
     DATA---ICFCAT.UCAT09
     INDEX--ICFCAT.UCAT09.CATINDEX
 DATA---------ICFCAT.UCAT09
   ATTRIBUTES
     KEYLEN---------------45  AVGLRECL----------4,086  BUFSPACE---------12,291  CISIZE------------4,096
     RKP------------------9  MAXLRECL---------32,400  CA SIZE---------573,440  CI/CA---------------140
     SHR(3,3)          SPEED  NOERASE    NOWRITECHECK  IMBED           NOREUSE  SPANNED
   STATISTICS
     REC TOTAL-------------0  CI SPLITS--------20,728  EXCPS-----------------0  UPDATE/OUTPUT FLAG--OFF
     REC DELETED---3,130,144  CA SPLITS------------45  EXTENTS--------------18
     REC INSERTED----------0  FREESPACE CI%--------20  LAST UPDATED:
     REC UPDATED-----------0  FREESPACE CA%--------20           (DATE INVALID)
     REC RETRIEVED---------0  APPROX FREE CI'S--7,980
   ALLOCATION
     SPACE TYPE-----CYLINDER  HI ALLO RBA--32,686,080
     SPACE PRI------------40  HI USED RBA--32,112,640
     SPACE SEC-------------1  APPROX FREE CA'S------0
   VOLUME 1
     VOLSER-----------TESCC1  PHY BLOCK SIZE----4,096  HI ALLO RBA--32,686,080  EXTS THIS VOL--------18
     DEVTYPE----------3380-3  PHY BLOCKS/TRK-------10  HI USED RBA--32,112,640  EXTENT TYPE--------DATA
     VOLFLAG-----------PRIME  PHY CA SIZE----------15  HI KEY RBA---18,628,608  CI'S PER TRACK-------10
     CYLS ALLOC-----------57  SMALLEST EXTENT-------1  LARGEST EXTENT-------40  % UTIL PER TRACK-----85
 INDEX--------ICFCAT.UCAT09.CATINDEX
   ATTRIBUTES
     KEYLEN---------------45  RECORD SIZE-------4,089  CA SIZE----------40,960  CISIZE------------4,096
     AIX RKP--------------9                                                     CI/CA---------------10
     SHR(3,3)          IMBED  REPLICATE
   STATISTICS
     REC TOTAL-------------0  SEQ SET SPLITS--------0  EXCPS-----------------0  INDEX:
     REC DELETED---------N/A  IND SET SPLITS--------0  EXTENTS--------------19  LEVELS----------------2
     REC INSERTED--------N/A  APPROX FREE CI'S-----58  LAST UPDATED:            ENTRIES/SECT---------11
     REC UPDATED---------N/A                                    DATE INVALID)  SEQ SET RBA-------4,096
     REC RETRIEVED-------N/A                                                   HI LEVEL RBA----------0
   ALLOCATION
     SPACE TYPE--------TRACK  HI ALLO RBA-----237,568
     SPACE PRI-------------1  HI USED RBA-----233,472
     SPACE SEC-------------1
   VOLUME 1
     VOLSER-----------TESCC1  PHY BLOCK SIZE----4,096  HI ALLO RBA-------4,096  EXTS THIS VOL---------1
     DEVTYPE----------3380-3  PHY BLOCKS/TRK-------10  HI USED RBA-------4,096  EXTENT TYPE-------INDEX
     VOLFLAG-----------PRIME  PHY CA SIZE----------1                            CI'S PER TRACK-------10
     CYLS ALLOC---------0.07                                                    % UTIL PER TRACK-----85
   VOLUME 2
     VOLSER-----------TESCC1  PHY BLOCK SIZE----4,096  HI ALLO RBA-----237,568  EXTS THIS VOL--------18
     DEVTYPE----------3380-3  PHY BLOCKS/TRK-------10  HI USED RBA-----233,472  EXTENT TYPE-------IMBED
     VOLFLAG-----------PRIME  PHY CA SIZE----------15                           CI'S PER TRACK-------10
     CYLS ALLOC-----------57  SMALLEST EXTENT-------1  LARGEST EXTENT-------40  % UTIL PER TRACK-----85
*****************************************************************************************************
```

Report <2>: The Control Area Analysis

CATALOG: ICFCAT.UCAT09

Control Area Analysis Report <2>

-CA Log	Nber- Phy	-----Notes------ #1	#2	Free CI'S	CA %	Dead CI'S	CI Util	Cmpresd Key Avg	-------Control Area High Key Value---------- 0 1 2 3 4 4 1--------0---------0---------0---------0----5
0	0	CA-SPLIT	1/4 USED	69	>	44	11		*EUQS.BD1A.VSO2.W3W.TAPE.EP1.....*
1	55	CA-SPLIT		70	>	54	12		*EUQS.IT.....*
2	11	CA-SPLIT		49	>	54	13		*EUQS.ST1A.T01EC.MCCS.TQRMHIST .....*
3	53	3		34	>	47	15		*SBW.BLQSAM.PCOF0030.CONVRENS.P.....*
4	1	CA-SPLIT	LOW USE	122	>	15	15		*SBW.BLVSAM.FRONTEND.FIAVRPT.D.....*
5	12	CA-SPLIT	LOW USE	82	>	36	12		*SBW.CTRB0114.PLM034 .....*
6	34	2	LOW USE	67	>	15	14		*SBW.PDAST.LOAD.OL-DATA.AUG28.DB907.....*
7	2	CA-SPLIT		1	<	31	14		*SBW.PCOST.UNLOAD.OL-DATA.HCM.JAN0998.DB004...*
8	23	CA-SPLIT		21	<	34	15		*SBW.PCP.UT.B0111.F0001.ERR5B.....*
9	13	CA-SPLIT	LOW USE	67	>	27	16		*SBW.PCPST.TST20.PCPF003I.VALREQ06.....*
10	33	CA-SPLIT	LOW USE	70	>	38	15		*SBW.PCPUT.B0301.F0021.TE.....*
11	16	CA-SPLIT	LOW USE	116	>	34	14		*SBW.PLM.PPR.T930-005.RUN001CO.PLM937.....*
12	20	CA-SPLIT	LOW USE	41	>	23	12		*SBW.PLMISA.SYSTEST.RTGR01.D091.....*
13	46	CA-SPLIT		9	<	58	10		*SBW.PLMTT.PPR2.T930-650.D.....*
14	29	CA-SPLIT	LOW USE	53	>	40	11		*SBW.PLMTQRLO.SYSTEST.RTCR02.D983XXXX.E.....*
15	41	CA-SPLIT	LOW USE	103	>	25	12		*SBW.PLMT002.PLMFPASI.RUN001.....*
16	50	CA-SPLIT		18	<	39	9		*SBW.PLMWT.B0103.PCCF.ERR10.....*
17	24	CA-SPLIT		39	>	38	12		*SBW.PLMXPR.SYSTEST.RTCR01D1.D81.....*
18	39	CA-SPLIT	LOW USE	97	>	35	11		*SBW.PLMXPR.SYSTEST.RTFR03D1.F012.....*
19	26	CA-SPLIT	MAX-FSPC	139	>	69			*SBW.PLMXPR.SYSTEST.RTFR13D3.DINFXXXX.E.....*
20	40	CA-SPLIT	LOW USE	110	>	69	11		*SBW.PLMXPR.SYSTEST.RTGR03D1.D986.....*
21	38	CA-SPLIT	LOW USE	87	>	24	10		*SBW.PLMXPR.SYSTEST.RTJR01D1.D792.....*
22	44	CA-SPLIT		73	>	54	10		*SBW.RDUST.FO.NWRD.#233001.....*
23	8	CA-SPLIT	LOW USE	72	>	39	15		*SBW.RWD.S002.NOV0298.BACKUP.PCO001.....*
24	52	CA-SPLIT		66	>	64	13		*SBW.TSS.BLQSAM.PRPLFILE.M11.....*
25	18	CA-SPLIT	LOW USE	90	>	44	16		*SFN.BLQSAM.....*
26	21	CA-SPLIT	LOW USE	34	>	23	11		*SFN.BLVSAM.S001.FRON.FI001SCN .....*
27	25	CA-SPLIT	LOW USE	34	>	23	11		*SFN.JH4.BLVSAM.PRELIVE.FIAVAUD .....*
28	19	CA-SPLIT	LOW USE	30	>	31	13		*SFN.ST024.BLQSAM.FPRFFILE.G0006.....*
29	43	22		59	>	57	15		*SJN.DSBS001.#006.REA.....*
30	3	CA-SPLIT	LOW USE	68	>	41	17		*SJN.ENCPAP.LEPRE001.L.....*
31	56	CA-SPLIT		71	>	56	16		*SJN.MLSC130.D.....*
32	35	CA-SPLIT	LOW USE	59	>	29	13		*SJN.MLS025N.LABEL.C222.J049.SY.W.....*
33	14	CA-SPLIT	LOW USE	60	>	39	15		*SJN.PAPO.TABANTE.PPSB91.....*
34	48	4		59	>	51	15		*SJN.PPSB701A.PPSBF201.V01.N.....*
35	4	CA-SPLIT	LOW USE	124	>	24	14		*SJN.RDUUT.FO.PSLOC.#220.....*
36	28	1		38	>	37	15		*SQO.BLVSAM.N56.D960401 .....*
37	5	CA-SPLIT	LOW USE	17	<	24	11		*SQO.IT6A.P0202.....*
38	10	CA-SPLIT	LOW USE	54	>	17	10		*SQO.JWP.P0628BAT.A.....*
39	51	CA-SPLIT	LOW USE	55	>	38	9		*SQO.MPR.LETTERS.F.....*
40	22	CA-SPLIT		27	<	32	14		*SQO.ST6A.M0061A.....*
41	54	4	LOW USE	121	>	38	13		*SQO.ST6B.L.....*
42	6	CA-SPLIT	LOW USE	57	>	22	11		*SQO.UT6A.RCONTRIT.D.....*
43	30	CA-SPLIT	LOW USE	121	>	12	18		*SWN.BLQSAM.MENOA.L1002677.H.....*
44	37	CA-SPLIT	LOW USE	89	>	40	13		*SWN.TTS006.C.....*
45	9	CA-SPLIT		60	>	55	13		*SWN.TTS030.L001R01U.DAT.C.....*
46	45	CA-SPLIT	LOW USE	51	>	27	10		*SWN.TTS032.L001R01A.....*
47	36	CA-SPLIT	LOW USE	42	>	29	10		*SWN.TTS033.L001R01.IVITF5.....*
48	49	CA-SPLIT		44	>	51	12		*SWN.TTS040.L001.DAT.VMS405.....*
49	42	CA-SPLIT	LOW USE	59	>	43	11		*SWN.TTS047.L001R02.R.....*
50	47	CA-SPLIT		5	<	43	11		*SWN.TTS101.L005.DAT.VMS021.....*
51	31	CA-SPLIT	1/4 USED	69	>	51	13		*SWN.TTS121.L004.DAT.VMS018I.....*
52	17	CA-SPLIT		56	>	49	12		*SWN.TTS131.L002R01.OUT.VDFDR01.....*
53	27	11		28	=	48	13		*SWN.TTS134.L002R01.OUT.PSCDFA0.....*
54	7	CA-SPLIT		4	<	34	13		*SWN.UT.VBCB071.O553BCA2.....*
55	15	CA-SPLIT	LOW USE	66	>	41	13		*SWN.UUU716.L001.DAT.O.....*
56	32	2		23	<	47	14		*SWN.VMSRCS.TSHPPT.LOGS *

****

Report <5>: The Statistical Summary

```
         Copyright 1989, 1998 by Mainstar Software Corporation, Bellevue, Washington, USA.

 <V4R2.01>                              V S A M    M A N A G E R            06 Apr 2001  14.00  Page 3
CATALOG: ICFCAT.UCAT09
                                      Statistical Summary Report <5>
  CREATION DATE-----04 AUG 1996       LAST UPDATE TIME----------N/A       LAST UPDATE DATE----------N/A
  DAYS BETWEEN CRE/UPD------N/A
DATA COMPONENT:                                                       INDEX COMPONENT:
  RECSZ------------(4086 32400)       CISZ--------------------4,096       CISZ--------------------4,096
  KEYS-------------------(45 9)       FREESPACE------------(20 20)       IMBED--------------REPLICATE
  ALLOCATION----------CYL(40 1)       SPANNED--------------NOREUSE       ALLOCATION-----------TRK(1 1)
  CURRENT ALLOC SIZE-----57 CYL       EXTENTS--------------------18       CURRENT ALLOC SIZE------1 TRK
  DEVICE TYPE------------3380-3       SHROPTIONS(3,3)---------SPEED       LEVELS-----------------------2
  VOL-TESCC1  CYL---57  EXT--18                                           VOL-TESCC1  TRK----1  EXT---1
DATA COMPONENT CI FREESPACE ANALYSIS (BASED ON A RANDOM SAMPLING OF THE DATA RECORDS):
  RECORDS SAMPLED---------2,587       COUNT OF SIZES-------------78       STD DEVIATION--------------73
  MINIMUM REC SIZE-----------73       AVERAGE REC SIZE----------106       MAXIMUM REC SIZE--------1,708
  MAX RECORDS PER CI---------37       REQUESTED CI FSPC %--------20       CI FSPC COST IN CYL-------1.4
  % UTILIZATION OF CI--------96       ACHIEVED CI FSPC %---------19
  RECORDS LOADED W/FSPC------30       WASTED CI FSPC %------------4
DATA COMPONENT CA FREESPACE ANALYSIS:
  CI/CA--------------------140       CA SIZE (TRK)--------------15       CA FSPC COST IN CYL-------1.6
  INITIALLY LOADED CI/CA----112       INITIALLY FREE CI/CA-------28       TOTAL FSPC COST IN CYL----3.0
DATA COMPONENT CI/CA SPLIT ANALYSIS:
  LISTCAT-CI SPLITS------20,728       NIS-MODE CA SPLITS---------49       RATIO CI:CA SPLITS------461:1
  LISTCAT-CA SPLITS----------45       SIS-MODE CA SPLITS----------0       CA SPLITS - % OF FILE-----613
LOGICAL RECORD ANALYSIS:
DATA COMPONENT:                                                       INDEX COMPONENT
  REC-TOTAL------------------0       INITIAL RECORDS-----3,130,144       REC-TOTAL------------------0
  REC-DELETED---------3,130,144       PERCENT DELETED-----------100       REC-DELETED----------------0
  REC-INSERTED---------------0       PERCENT INSERTED------------0       REC-INSERTED---------------0
  REC-UPDATED----------------0       NET GAIN (LOSS)---(3,130,144)       REC-UPDATED----------------0
  REC-RETRIEVED--------------0                                           REC-RETRIEVED--------------0
  LOGICAL I/O REQ-----6,260,288                                          LOGICAL I/O REQ------------0
  EXCP'S---------------------0                                           EXCP'S---------------------0
  RATIO: LOGICAL:EXCP-------N/A                                          RATIO: LOGICAL:EXCP-------N/A
DATA COMPONENT SPACE ALLOCATION ANALYSIS:
  TOTAL ALLOCATED CA'S IN THE FILE-------57  TOTAL ALLOCATED CYL-------57  PERCENT OF 4 GB----------0.8
  TOTAL USED CA'S IN THE FILE------------57  PERCENT USED-------------100  EQUAL TO HURBA/CASZ?------NO
  TOTAL FREE CA'S IN THE FILE-------------0  PERCENT FREE---------------0  OVERALL FSPC IN CYL-----41.5
  TOTAL DEAD CA'S (MASSIVE DELETIONS)-----0  PERCENT OF USED CA'S-------0  EXPRESSED IN CYL---------0.0
  TOTAL CI'S IN THE FILE------------------0  (INCLUDING FREE & DEAD CA'S)
  TOTAL CI'S TO 'HURBA'--------------7,980  (IN USED CA'S ONLY)
  TOTAL USED CI'S TO 'HURBA'----------4,531  PERCENT USED--------------57  EXPRESSED IN CYL--------30.2
  TOTAL FREE CI'S TO 'HURBA'----------3,449  PERCENT FREE--------------43  EXPRESSED IN CYL--------23.0
  TOTAL FREE CI'S TO HARBA------------3,449  PERCENT OF TOTAL FILE------0  EXPRESSED IN CYL--------23.0
  ORIGINAL FREE CI'S PER CA-------------28  AVERAGE FREE CI'S PER CA--61  PERCENT OF ORIGINAL------218
  TOTAL DEAD CI'S (BAD INDEX CISZ)--------0  PERCENT OF TOTAL FILE------0  EXPRESSED IN CYL---------0.0

*******************************************************************************************************
```

Report <6>: Recommendations

```
                                    R E C O M M E N D A T I O N S  <6>
VMP1604-Do not use IMBED with a KSDS/VRRDS on a 3380/3390 disk, especially when 'cached'.
        DATA allocation could be reduced by      3 CYLINDER(S).
VMP1603-Do not use REPLICATE with a KSDS/VRRDS on a 3380/3390 disk, especially when 'cached'.
        INDEX allocation could be reduced to       1 track(s), saving       0 track(s).
*******************************************************************************************************
```