MAINSTAR

A VSAM/ICF Catalog
White Paper

From Mainstar Software Corporation

Managing ICF Catalog Record Activity
FREESPACE, splits, reorgs, and dead CIs

By Ronald K. Ferguson

The BCS is a 'regular' VSAM KSDS, and as such, is tuned in a very similar manner to any other KSDS. One difference, though, is that you don't have statistics in the LISTCAT for a BCS as you would for other KSDSs. IBM decided years ago, since the BCS will often be cross-system shared, that maintaining valid statistics isn't practical. Therefore, the standard record total, record inserted, and record deleted values that you would rely on to determine record activity in a regular KSDS, simply isn't available for a BCS.

Interestingly, CA split counts (and the HURBA) are updated for a BCS that's been defined SHR(3 4), and the updates are made in real-time as they occur, and not at close-time as they do for 'regular' VSAM datasets. These updates do not occur for a BCS defined as SHR(3 3).

Recommendations For Catalogs

Here are some recommendations in a nutshell, with justifications to follow:

- Don't worry about splits
- FREESPACE(0 0)
- Reorg catalogs as <u>infrequently</u> as you can (i.e., almost never, if that's possible)

These recommendations require a considerable leap of faith for many people, as they fly in the face of conventional wisdom. The simple logic behind them is that most BCSs will have *unevenly* distributed record insertion activity, and therefore, *evenly* distributed free space is of very little value. Splits, both CI and CA, are the best technique available to handle this uneven distribution. To allow splits to work their magic (the magic being - to have free space at the location where it's needed), the best thing you can do is to leave them alone once they've occurred (and hence, don't reorg).

BCS Record Insertion Pattern

The key of the BCS is dataset name (at least, that's true for User Catalogs, which is what we're most concerned with here). For a User Catalog, multiple aliases are frequently assigned, where an alias equates to the high-level qualifier (*hlq*) for its application dataset name group. Since the number of new dataset allocations for one hlq is not likely to be the same across them all, it's almost certain that uneven numbers of record insertions will occur. Even within an hlq, there may be a lower level qualifier of the dataset name for a section of the application that gets many, many new datasets, again resulting in uneven record insertions within the BCS.

This concept, while easy to explain, is very difficult to 'see' inside the catalog. At most, all you can see from a LISTC, of the catalog is the total <u>number</u> of splits but not <u>where</u> they've occurred.

A 'map' report from Mainstar's VSAM Manager provides this visibility of free space inside the catalog, and illustrates exactly where splits have occurred. (See ***Figure 1***)

Here's how to interpret the report:

- The first column (on the left) simply numbers each CA within the catalog. Since a CA is equal to a cylinder, and this column is numbered from 0 to 19, this catalog has a total of 20 CA's (cylinders).
- The second column traces the logical sequence of the CA's, by following the horizontal chain pointers on the Sequence Set Index records (level 1 of the index).

By reading down this column, CA splits can be seen at each point where there is a 'jump' in the CA number. For example, CA #0 is followed by #13, which indicates that at some time in the past, there has been a CA split of the records from CA #0 out to CA #13.

Since CA #13 is then followed by #18, that indicates that a subsequent CA split occurred at this location.

For CA #18 to then be followed by #12 would seem to be an anomaly, but when explained, it's easy to see what has occurred.

Remember that this report is a point-in-time map of the CAs, and therefore, the sequence of events is not intuitively obvious. What's happened is, CA #0 originally split out to #12, at which time #12 would have shown up in the report column immediately after #0 if VSAM Manager were to have been run then. But after that split, heavy record insertion activity continued in CA #0, causing it to fill up once again and split out to #13. Following that, there were heavy insertions within the records in #13, causing it to fill and then split out to #18. Given sufficient time and inclination, you could actually map out the relative sequence of each CA split, but as a rule, it isn't important for our tuning purposes.

To conclude what has happened at this initial location of the file, we see that a total of six CA splits have occurred in this one location of the catalog, as there are six out-of-sequence CA numbers between #0 and #1. VSAM Manager has noted that in its analysis, and has indicated the number 6 after the string of six CA split notations.

- Continuing our analysis at that location of the catalog, go over to the Control Area High Key Value field on the right-half of the print line. Remember that the key of the catalog (BCS) is dataset name, so what we're seeing here is the key of the last (highest key) record of the CA. Since dataset records are assigned to a catalog based on the aliases defined for the catalog, we can see that this area of the file contains records for the hlq value P373. This must be a very active hlq (i.e., lots of new dataset defines), as that's what is causing the relatively high number of CA splits at this location in the catalog.

  Now, consider for a moment, the overall amount of clustered record insertions that would be necessary to produce CA splits at a single location in the file. On the very first CA split at this location of the catalog, half of the CIs from CA #0 would have been moved out to the new CA. As you can see from the abbreviated Catalog Listing report, there are 180 CIs/CA, so each CA split moves 90 CIs of data, leaving 90 free CIs in its place. In order to once again fill up the just-split CA and cause another CA split, there would have to be enough record insertions to fill and split 90 CIs. This could easily have been several thousand catalog records, all inserted in the same location, just to cause the second CA split. For this catalog to have 6 CA splits at this location indicates a huge number of records insertions, all for the dataset name value of P373 on the hlq.
- Now move down to CAs #1-5 (in the PHYSICAL CA column). This is an area of the catalog where there haven't been any inserts to speak of, for hlq values that range from P384 to P515. To put it another way, this must be a very inactive range of dataset name values, where there aren't any dataset defines or deletes going on.
- Further down, we see that hlq P568 is another highly active dataset name hlq, as there have been three CA splits from the original CA #6 that contained dataset name records of this hlq.

The conclusion from this analysis is: the pattern of record insertion activity for this catalog is so varied, between no insertions at some locations, and heavy insertions at others, that *evenly* distributed FREESPACE would be of no value. In fact, at locations where there are no insertions, it would simply waste space and at locations where there are heavy insertions, it would be far too little free space, resulting in CA splits anyway.

This should leave us with the conclusion that the CA splits are the only way for this file to achieve its necessary free space, at the locations where it is really needed.

Put another way, the CA splits shown for this catalog are doing exactly what they were designed to do - namely, to dynamically allocate free space exactly where it is needed most.

```
                           C A T A L O G   L I S T I N G
USERCATALOG---CATALOG.VSYS002.ICF
 ATTRIBUTES
  KEYLEN-------------45   AVGLRECL--------4,086   BUFSPACE--------9,728   CISIZE----------4,096
  RKP-----------------9   MAXLRECL-------32,400   CA SIZE-------737,280   CI/CA-------------180
STATISTICS
  REC TOTAL-----------0   CI SPLITS------55,499   EXCPS---------------0
  REC DELETED---129,263   CA SPLITS----------10   EXTENTS-------------1
  REC INSERTED--------0   FREESPACE CI%-------0
  REC UPDATED---------0   FREESPACE CA%-------0
  REC RETRIEVED-------0   FREE CIS-------1,842
INDEX--------CATALOG.VSYS002.ICF.CATINDEX
 ATTRIBUTES
  KEYLEN-------------45   RECORD SIZE-----2,041   CA SIZE--------45,056   CISIZE----------2,048
  AIX RKP-------------9                                                   CI/CA--------------22
  SHR(3,4)        IMBED   NOREPLICATE


                        Control Area Analysis Report
                                        ----------Control Area High Key Value----------
CA Number  -Notes-  Free  CA    CI   0         1         2         3         4
Log  Phy            CIS    %   Util  1---------0---------0---------0---------0------
-----------------------------------------+---------+---------+---------+---------+------

  0    0  CA-SPLIT   72    >    46  *P373.DSNDBD.OESOADB.POA01X1.I0001.A042.......*
  1   13  CA-SPLIT   53    >    65  *P373.P257....................................*
  2   18  CA-SPLIT   38    >    53  *P373.P733000E.D373301........................*
  3   12  CA-SPLIT   62    >    54  *P373.P900003E.TSORTORD ......................*
  4   14  CA-SPLIT   26    >    65  *P373.R130004E.TRTINDA........................*
  5   17  CA-SPLIT   47    >    53  *P373.R733000E.D373301........................*
  6    1     6       83    >    54  *P384.R788003E.ENDSGATD ......................*
  7    2              0    =    65  *P386.T504000K.RUBFLGSC.......................*
  8    3              0    =    53  *P414.P844000E.D4942257.......................*
  9    4              0    =    54  *P438.X914523E.FS14DRDA ......................*
 10    5              0    =    65  *P515.P373....................................*
 11    6  CA-SPLIT   56    >    53  *P568.G845000E.E217701........................*
 12   15  CA-SPLIT   68    >    54  *P568.H113561E.DEFDDIND ......................*
 13   19  CA-SPLIT   38    >    53  *P568.N768912D.VSRTOUT1.......................*
 14    7     3       81    >    54  *P574.A268005E.TSORTORD ......................*
 15   10  CA-SPLIT   83    >    61  *P574.F944112E.VAPP0015.......................*
 16   16  CA-SPLIT   42    >    64  *P589.PO......................................*
 17    8     1       89    >    77  *SYS1.VVDS.VONLG22                          ..*
 18    9              0    =    53  *P770.EBI.BBRC................................*
 19   11            104    >    65  *P970.V970.XREF3.TEMP.INDEX                   *
```

Figure 1: Catalog With Uneven Record Distribution, Causing Beneficial CA Splits

```
                          C A T A L O G   L I S T I N G
USERCATALOG---CATALOG.VSYS003.ICF
 ATTRIBUTES
  KEYLEN-------------45   AVGLRECL--------4,086   BUFSPACE--------9,728   CISIZE----------4,096
  RKP-----------------9   MAXLRECL-------32,400   CA SIZE-------737,280   CI/CA-------------180
STATISTICS
  REC TOTAL-----------0   CI SPLITS------55,499   EXCPS---------------0
  REC DELETED---129,263   CA SPLITS----------10   EXTENTS-------------1
  REC INSERTED--------0   FREESPACE CI%-------0
  REC UPDATED---------0   FREESPACE CA%-------0
  REC RETRIEVED-------0   FREE CIS-------3,103
INDEX--------CATALOG.VSYS003.ICF.CATINDEX
 ATTRIBUTES
  KEYLEN-------------45   RECORD SIZE-----1,529   CA SIZE--------39,936   CISIZE----------1,536
  AIX RKP-------------9                                                   CI/CA--------------26
  SHR(3,4)        IMBED   NOREPLICATE


                        Control Area Analysis Report
                                           ---------Control Area High Key Value----------
CA Number -Notes-  Free  Dead Cmpresd 0         1         2         3         4
Log  Phy           CIS   CIs  Key Avg 1---------0---------0---------0---------0------
------------------------------------------+---------+---------+---------+---------+------

  0    0 CA-SPLIT  121   40    12    *AB21.FREGND.PRCATDB.POA02Y2.I0001.A042.......*
  1   13     1     101   19    11    *AC34.AGREGND.X101A...........................*
  2    1 CA-SPLIT  112   16    10    *AC34.VR848...................................*
  3   12     1     115   21    11    *AC34.XA10103E.TSONDERD ......................*
  4    2 CA-SPLIT  113   22    11    *AD73.R130004E.TRTINDA........................*
  5   17     1     106   16    10    *AD73.R733000E.D391301........................*
  6    3 CA-SPLIT  114   21    11    *AG14.R788003E.SRTSG101 ......................*
  7   14     1     105   15    10    *AG14.T504000K.XLAFLGSC.......................*
  8    4 CA-SPLIT  119   29    11    *AG14.P844000E.D5681361.......................*
  9   18     1     109   23    11    *AG14.Y103523E.ES91DRDA ......................*
 10    5 CA-SPLIT  101   15    10    *AG14.N262....................................*
 11   11     1     105   19    11    *AG14.GRDA000E.E130601........................*
 12    6 CA-SPLIT  101   12     9    *AG14.PRIMAR1E.DEFCNTND ......................*
 13   15     1     102   12     9    *AQ68.POLACCT1.VSINOUT1.......................*
 14    7 CA-SPLIT  113   24    11    *AQ68.ARMDA05E.CICSTIND ......................*
 15   19     1     108   22    11    *AQ68.F955112B.VSPT0015.......................*
 16    8 CA-SPLIT  100   10     9    *AQ68.PRODON12................................*
 17   16     1      99   11     9    *AR10.SYSARTD1.VSRT1201.......................*
 18    9           128   48    12    *AR10.XRRTANI.PPRC............................*
 19   10           131   49    12    *SYS1.VVDS.VONLG41                            *
```

Figure 2: A Catalog With Dead CIs, Causing Unnecessary CA Splits

Reorg Recommendations

Consider that a CA split requires hundreds of I/Os and upwards of an elapsed second to complete. During ICF catalog management processing with this catalog, a total of 10 seconds of elapsed time has been spent performing these 10 total CA splits. That isn't a long time in the overall scheme of things, but it also isn't something you should cause just for the fun of it. When someone reorgs this catalog, they will effectively remove all of the free space that's now imbedded within the CAs, and will put this catalog back to FREESPACE(0 0) initial load status. What that means is, new CA splits, at the very locations where they've already occurred, will then have to be done once again just to create the free space where it's needed.

Oftentimes, that is the problem with CA splits that cause people to reorg their VSAM file or catalog - they do a reorg with performance improvement in mind, little realizing that they are actually <u>decreasing</u> the performance. In fact, there's nothing worse than doing a CA split all over again - one that you wouldn't have had to do if the first place if the reorg had not been.

The other reason many people do reorgs is to reclaim DASD space - space in a file that they consider is wasted because of the dreaded CA splits throughout it. As you can hopefully see from our analysis, if we reorg this file, it will definitely shrink back to about 11 CAs, but almost immediately after the reorg, when new dataset defines are done, new CA splits are likely to occur at the same locations as before (i.e., at the same high activity dataset name locations), and the file will increase to 20 CAs. What was intended as a DASD reclamation is, in fact, just a temporary 'feel good' effort, to the detriment of the catalog's performance.

Somebody else might think that a reorg 'cleans up' the key direct or sequential access performance. That also isn't true, for long-winded reasons that are best left to another White Paper. Suffice it to say that access performance on a catalog is just about equal before a reorg as after, regardless of how many splits have occurred within it.

DEAD CIs

Many catalogs suffer serious DASD space and performance problems due to the long-standing problem related to an incorrect index component CI size assigned to the catalog. The VSAM Manager report in ***Figure 2*** illustrates the problem, and will be used for this discussion.

Note in the abbreviated section of the Catalog Listing report, that this catalog has the following characteristics:

- A data CISZ of 4096
- A CI/CA count of 180
- An index CISZ of 1536, resulting in all index records for this catalog being 1529 bytes (the CISZ minus 7 bytes for an RDF and CIDF)

In the Control Area Analysis report, the DEAD CIs and the Compressed Key Average columns have been included (they were removed from the report in ***Figure 1*** for sake of clarity, as well as to fit the report on the page).

Here's the background behind this problem. Each control area (CA) of a VSAM KSDS is indexed by a single Level 1 index record (also called the Sequence Set record). This record must be large enough to index all of the CIs in the associated CA, with the high key of each CI (in compressed format), and the CI number pointer value. The dead CI problem occurs when the index record length isn't large enough to index all CIs, and whenever that happens, the remaining CIs within the CA are effectively 'dead', or more politely stated, unusable. The effect on the dataset (catalog, in this case) is wasted space, but worse, premature and unnecessary CA splits.

All documentation on VSAM and ICF catalogs (such as Managing Catalogs, and the new IBM Redbook, *VSAM Demystified*) continues to recommend that you let the index CI size default, and herein lies the problem. Currently, the assumption inside the IDCAMS code for calculating the default index CISZ value assumes better key compression than catalogs (and datasets) actually get (the code assumes between 5-8 bytes after compression).

Consider this calculation: With an index record size of 1529 (see the third bullet on the previous page), divided by 180 CI/CA, this will only work properly when the average compressed key length for the record is less than 8½ bytes (e.g., 1529 / 180 = 8.5 bytes).

Looking at the VSAM Manager report in ***Figure 2***, it's easy to see by scanning the Compressed Key Average column in the CA Analysis that this catalog was found to *never* have an average compressed key less than 9 bytes.

Therefore, VSAM Manager's analysis has found that every CA within the catalog has a substantial number of dead CIs. Worse, even if this catalog had free CIs allocated for it with the FREESPACE parameter, these free CIs would be down at the bottom of the CA (where the dead CIs also occur), resulting in an immediate CA split on the very first record insertion.

The remedy for this problem is to ensure that the index CISZ value is never left to default, but rather is explicitly specified in the DEFINE USERCATALOG command. To determine the correct CISZ for an existing catalog, it is best to run VSAM Manager's MAPBCS on it, and the resulting Recommendations Report will provide for you the optimum CISZ to specify. For a new catalog, explicitly specify an index CISZ of 2048, then run a VSAM Manager MAPBCS after the catalog is fairly well populated, so that the final optimum size can be specified for the next catalog reorganization.

For our catalog in ***Figure 2***, where the largest number in the Compressed Key Average column is 12, we would multiply that by 180 CIs/CA, to get 2160. This tells us that our index CISZ would have to be larger than this to correct for the dead CI problem – I'd recommend 4096 for this BCS. Dead CIs are a problem with both application datasets and catalogs.

IBM is currently working with Mainstar to determine the optimum default index CISZ for catalog definitions, based on a widespread field analysis of key compression. Until that analysis is complete, and the code within IDCAMS has been changed to reflect a more accurate default index CISZ, it is strongly recommended that you never let this parameter default, and instead, explicitly specify an index CISZ. This problem can be so damaging to a catalog, that it is recommended to run VSAM Manager against every catalog in your installation, and for any that are found with this problem, you should schedule a reorg at your earliest opportunity to correct it.

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.