

A VSAM/ICF Catalog White Paper

From Mainstar Software Corporation



Catalogs Do Break! — Part 1 – *With More Frequency and Regularity Than Most People Realize*

By Ronald K. Ferguson

Preface: "I've never had a catalog failure", or "catalogs don't break these days" are perhaps two of the most often-heard comments about ICF catalogs. Yet, in installations that truly believe this, there are probably 'broken' catalogs right now, and no one even knows about it. The truth is, catalogs break in myriad ways, frequently without anyone even knowing about it, and in ways that are not debilitating to overall successful processing through that catalog.

As more installations run Mainstar's Catalog RecoveryPlus, and begin to run the diagnostic routines that will identify if you have catalog problems, the more feedback we're getting that indicates catalog failures are much more commonplace, whether subtle types of failures that only affect a few records, or serious errors that wreak havoc across wide parts of the catalog.

Also, many people are only referring to the BCS when thinking about 'catalog failures', without realizing that the VVDS is the other half of an ICF catalog. VVDSs also break, and in ways that are not understood. VVDS failures occur on RAID volumes, where there's a blithe attitude that a mirrored copy will protect you from this sort of hazard, but in fact, only presents you with two copies of the failure.

This White Paper will attempt to illustrate that catalogs do break, and with catalogs possibly being one of the weakest links in your entire OS/390 environment, ignoring this potential for business disruption can be a disaster far greater than the other disasters you are already planning for.

ICF Catalogs – Possibly Your Weakest Link

Here's something to try when you have a little spare time on your hands:

1. Run an IDCAMS LISTCAT summary of your ICF catalog environment to determine how many datasets you have cataloged in each catalog, then tally that up across all catalogs to determine the total number of cataloged datasets in your environment.
2. Next, consider that virtually all jobs accessing those datasets rely on the catalog entry to identify the location of the dataset, and without that information, production job streams would have to be radically modified to identify that location.
3. Then, check around to determine how often (and how regularly) you run thorough diagnostics against your ICF catalog environment. Do you run an IDCAMS EXAMINE INDEXTEST against each BCS every time you back it up? Do you run an EXAMINE DATATEST at some regularity? Do you regularly run IDCAMS DIAGNOSE ICFCATALOG and DIAGNOSE VVDS commands?
4. Finally, how often do you run a serious catalog recovery exercise, to identify flaws in your system environment, or flaws in your technical expertise, that will make a catalog recovery difficult or much more time consuming than can be afforded. Consider how often your installation tests its DR plan (and why), versus how often you test your catalog recovery process. My theory is that your CIO or Data Center Manager is aware of the overall risks you face against natural and man-made disasters that can close

down your entire data center, but that many people don't realize that the odds are probably much greater for a catastrophic catalog failure to debilitate your processing capability. Is your CIO or Data Center Manager aware of this?

Catalogs Do Break These Days

On the 4th of July, which happens to be America's great fireworks holiday, fireworks of another kind were encountered at a major auto manufacturer, when two broken ICF catalogs were found (the affected systems were on Release 2.9, and they are a JES3 shop). Everything that led up to the circumstance seemed innocent enough, but in fact, it created even more serious ICF catalog recovery problems. Here's a short synopsis of the problem. The details as I know them are a bit sketchy, but the point is to illustrate that catalogs do break, and very unexpectedly.

During the night, this large DB2 shop (and who isn't?) experienced some DB2 problems, which resulted in many log files being written. Each log file is cataloged, and with the number being generated, this resulted in the catalog suddenly 'getting full', with an explosion out to the maximum 123 extents.

When the next scheduled backup of the catalog was attempted (using Catalog RecoveryPlus BACKUP), the internally-generated IDCAMS EXAMINE INDEXTEST produced a message indicating many duplicate records in the index. Then, when it tried to read the data component of the BCS as an ESDS (which is the standard way that Catalog RecoveryPlus processing works), the subsequent backup process itself failed, getting error message:

```
IEC161I 032(008,042,IGG0CLEG)-002
```

The return code and reason code of the message indicate a no-record-found condition, or an end-of-data occurred while reading the catalog.

A RECOVER from the last backup was then run, using as input, SMF data from the time of backup to the current time, giving them a forward recovery of the catalog.

When the problem was researched, APAR OW45637 turned up, which provided the following information (paraphrased):

One of the BCS' index records contains duplicate index pointers (and other possible errors) – the possibility exists that index records may be damaged in catalogs that are shared between systems, where another system making updates to the BCS is using incorrect information, causing damaged index records.

A PTF was supplied for this APAR, and the user has had no further indication of this problem.

Another 'Broken' BCS Incident

Another installation, this time a very large bank, experienced a somewhat similar catalog problem. They too are a Catalog RecoveryPlus user, and reported to us that BACKUP for the catalog was executing successfully, but the internally-issued IDCAMS EXAMINE command, with the INDEXTEST keyword, was turning up "some 6000-odd duplicate entries" in the BCS. This installation was ignoring these errors, waiting for some scheduled downtime when they could have a closer look at this BCS. One night this past May, though, the EXAMINE process within BACKUP suddenly issued the following message, and then abended with a S0C4-11:

```
CI ON CURRENT LEVEL NOT POINTED TO  
BY NEXT LEVEL
```

This user had two 2 options:

1) Recover from the last good backup and apply any SMF updates for a forward recovery, but due to the amount of SMF data involved (six systems sharing access to this catalog), they felt this was not a good option.

2) Try to get a valid backup, using Catalog RecoveryPlus BACKUP, as it processes just the data component of the BCS, bypassing the index (which the message above is indicating has errors). This option would not have been valid if they were using any other catalog backup utility (such as EXPORT), which relies on the index to access the data inside the catalog.

They ran BACKUP with NOEXAMINE specified, and then performed a RECOVER from the backup just taken – everything worked well, and in just a few minutes they had a nice clean catalog, with no further duplicate index entries reported, and every BACKUP since that time has run with RC=0.

Still Another 'Broken' BCS

Here's another example of a broken catalog, where the only indication is that datasets can't be located just after they are created. The catalog appeared to be working OK, but several members of the technical staff were saying there was some sort of problem with it. The catalog support person decided to "do some DIAGNOSE and EXAMINE runs" against the catalog, and found serious errors that indicated out-of-sequence and/or duplicate keyed records inside the data component of the catalog.

Here's where it gets exciting – a check was made to locate the latest backup for the catalog, only to find that somehow *this catalog was not being backed up!* Their catalog backup tool is currently HSM, and now when they tried a backup of this catalog, it failed, HSM uses EXPORT 'under the covers', and when it encounters structural errors, it fails.

The catalog support person had just learned of Mainstar's Catalog RecoveryPlus, and when he heard that it backs up using just the data component of the catalog, then sorts the records as part of the recovery (and drops out duplicate records that it finds), he brought in Catalog RecoveryPlus to determine if it would help him with his recovery. Mainstar is now helping him to get this catalog cleaned up, and resolve all of the problems.

This type of problem is not uncommon, and the situation that led up to its identification is very interesting. Because many people don't feel that catalogs break these days, there sometimes isn't sufficient emphasis on making sure that all catalogs in the shop are in the backup list.

Worse, with the hodge-podge naming convention that many installations have for their catalogs, it's very difficult to keep a close eye on catalog backups. With Catalog RecoveryPlus' unique extended mask filters for identifying catalogs to be backed up, you have a much greater chance of selecting all catalogs in your BACKUP command.

Perhaps you have been merging and consolidating data centers, with the result that catalog naming conventions are a mess – you can use the Catalog RecoveryPlus NEWNAME facility with the RECOVER command, enabling you to easily and safely convert your catalog

names over to a good, solid naming convention, ensuring that all catalogs are backed up. Finally, with the ability to run EXAMINE and DIAGNOSE, at your discretion, on every catalog at the time it is backed up, you are assured of finding errors when they occur, not weeks or months later when they've already produced lots of problems for you.

Catalogs Don't Break Only in Catastrophic Ways

Oftentimes, catalogs can have internal problems that aren't structural in nature, but rather, just some 'bad information' within certain records.

Take for example, the experience of a large installation that was trying to split a very large catalog into several smaller catalogs (the large catalog came to them, sight unseen, when they merged with another company). Attempts to split out parts of the catalog using MERGECAT failed whenever it encountered cataloged dataset records containing a completely erroneous volser entry – i.e., a volser that contained all unprintable values, and which didn't represent any known volser that had ever been in the installation. How this volser value got into the catalog records is a mystery, but nevertheless, it was there. The reason MERGECAT failed is that one of its important functions is to update the catalog 'back-pointer' in the VVDS record for each catalog record moved to the target catalog, and to do this, a valid volser has to be in the catalog record in order to update the VVDS on that volume.

On the first alias group of catalog records to be moved, there were only a handful that contained these erroneous volser values, and it was fairly easy to manually delete these entries out of the catalog before running MERGECAT. The second alias group, though, encountered thousands upon thousands of catalog records with the erroneous volser – far too many to correct before running MERGECAT. As a result, getting this catalog sufficiently cleaned up before attempting MERGECAT is the requirement, and therefore, this is actually a case of a 'broken' catalog, not because it isn't functioning as it should, but because it has garbage in it that prohibits execution of necessary maintenance commands.

VVDS Corruption

Very little attention is given to the potential for VVDS failures. The VVDS is absolutely crucial in locating and opening VSAM datasets, and if one is corrupted, you're faced with the loss of access to every VSAM dataset on the volume (not to mention proper HSM support for non-VSAM datasets on the volume too).

Consider this situation: A very large bank lost two VVDSs in the space of a two month period. While the volumes were 'development' volumes, meaning they were volumes that were being used by development application programmers, they contained enough datasets that this installation estimates it would have set back 300 application programmers upwards of two weeks of development effort, while the volumes were being recovered from the most recent volume-level backups. In this instance, the customer used the VVDS backup and forward recovery capability, and had the volumes back up and running in a very short time.

The point of this is, identifying critical volumes, such as complex DB2 or IMS database volumes, where large databases might be scattered across dozens of volumes, or as in the example above, volumes that would be hard to recreate from an entire volume-level backup, are good candidates for VVDS backup. With the Catalog RecoveryPlus BACKUP VVDS command, the elapsed time to back up a typical-sized VVDS is between 1-2 seconds, making it quite feasible to back up just the most critical ones. As the person stated who was in the 'hot seat' for the two VVDS repairs described above, "VVDS failures in my book are much nastier and insidious than catalog failures". This is definitely a case where an ounce of prevention is worth a pound of cure.

Testing Catalog Recovery

As mentioned in the opening paragraphs of this White Paper, it appears there is not enough emphasis given to catalog recovery testing – and testing with catalogs that are true to your actual production system, with SMF data collection also tested, and all of the other factors that come into play when a real catalog failure occurs. That's the reason your installation actually performs its disaster recovery at an off-site location, so that

everything related to an actual recovery can be tested. If you aren't performing true-to-life catalog recovery testing on a regular and scheduled basis, you're probably going to find it's a real disaster on your hands when the time comes.

Conclusion

Catalogs do break, and in fact, much more frequently than most people realize. The seriousness of the breakage ranges from:

- CAS abends that might require you to close and re-open a catalog each time the failure spot is 'touched',
- to a catastrophic failure that requires an immediate recovery (typically a forward recovery) before the catalog can be used again,
- to troublesome structural or data content errors inside the catalog, resulting in unpredictable or intermittent catalog problems that only surface when that particular location in the catalog is accessed.

Nevertheless, all of these should be classified in the broken catalog category, but with varying degrees of seriousness (and therefore, urgency to get it repaired).

One last suggestion – there is an IBM Redbook, titled *Integrated Catalog Facility Backup and Recovery*, SG24-5644. Starting in September 2001, I will be part of a 4-person, 5-week residency, to update this Redbook to incorporate substantial new information on how to use Mainstar's Catalog RecoveryPlus in your ICF catalog backup/recovery plan, and your day-to-day ICF catalog management. I expect it will be available by late 2001, and you can bet that Mainstar will encourage everyone to get a copy of this Redbook.

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.