ICF Catalog
White Paper
From Mainstar Software Corporation

MAINSTAR

Recovering ICF Catalogs at the Disaster Recovery Site

By Colleen Gordon

Recovery from any type of disaster becomes more complicated as newer technologies are developed and implemented in the data center. In a z/OS environment, a focal point of a successful system recovery is recovering the ICF catalogs and the records they contain. The recovery methodologies available for ICF catalogs are: empty catalog recovery, full catalog recovery and selective catalog recovery. This paper describes these three methodologies and provides information to help the reader determine the best one for their data center.

Empty Catalog Recovery

If your installation is utilizing the DFSMShsm data mover, Aggregate Backup and Recovery Support (ABARS), empty catalog recovery is possible. When ABARS backs up a data set on tape, on a direct access storage device (DASD), or in DFSMShsm migration, the catalog entry is included in the backup. When a generation data set (GDS) is backed up, the generation data group (GDG) entry and the catalog entry is also included in the backup. In addition, ABARS supports placing ICF catalog data set names in the ALLOCATE list of an ABARS aggregate. This action results in ABARS backing up the allocation parameters of the ICF catalogs to tape for subsequent recovery at the recovery site.

At the recovery site, ARECOVER commands are executed for all aggregate backups taken at the backup site. ICF catalogs placed in the ALLOCATE list in the aggregate's selection data set are allocated empty at the recovery site using their original allocation parameters. The records within the catalogs are not recovered.

As ARECOVER events are executed and data is recovered, ABARS catalogs the restored data in the appropriate ICF catalog, including the GDG entries. This methodology can be used for all data recovered using ABARS. However, keep in mind that in an ABARS environment, some data must be recovered outside of ABARS. DFSMShsm must be active before ARECOVER commands can be issued, requiring that the system be IPL'd and active prior to using ABARS. Most often, system packs are recovered from full volume dumps using DFSMSdss or FDR. The catalogs that support system data are recovered full, meaning they are recovered with all of their catalog entries intact.

Tape Entries and ABARS

ABARS is the only data mover that supports tape data sets, except IEBGENER, a data mover not commonly used for mass movement of data other than tape. ABARS also supports data on DFSMShsm migration level two (ML2) tape, without recall, which is not supported by IEBGENER tape data sets, including virtual tape. When tape data sets are in the INCLUDE list of an ABARS aggregate, the tapes are mounted and the data are copied to some kind of removable media.

At the recovery site, tape data is written to the customer's specified output device: native tape, virtual tape, or DASD. Generally, in an ABARS environment, not all tape data is included in an ABARS aggregate. There are several tape data sets that are vaulted that are needed at the recovery site. The catalog entries for these tape data sets can be copied and restored into the empty catalog environment using ABARS. The tape data set names or data set name masks that are easily obtainable in the tape management system's vault list can be copied and pasted into the allocate list of an ABARS aggregate.

When the ABACKUP command is executed, all of the catalog entries that support the vaulted tape data sets are copied to the output device and restored into the allocated (empty) catalogs at the recovery site with the execution of the ARECOVER command. Other tape data sets' catalog entries can be supported in a similar fashion: creating a list of the data set names, cutting and pasting them into an ABARS allocate list, and executing ABACKUP.

Full (Populated) Catalog Recovery

The next methodology to discuss is full, or populated, catalog recovery. This is simply the act of backing up an ICF catalog with all of its entries using a data mover such as IDCAMS, DFSMSdss logical backup, or full volume dump. At the recovery site, the ICF catalog is restored exactly as it was at the backup site.

The difficulty with this type of catalog backup and restore methodology is that it is very difficult to get a synchronized backup of the catalog. Essentially, all access to the ICF catalogs must be paused and a backup taken. An un-synchronized backup results in discrepancies between the records in the ICF catalog and the data on the physical volumes (and in DFSMShsm migration) at the recovery site.

An un-synchronized recovery of a catalog can result in confusion and delays in processing at the recovery site. As an example, a batch job executing at the recovery site requires a data set that, according to the ICF catalog record at the time the catalog was backed up, is migrated. However, the data set was actually recalled some time after the backup of the catalog was taken. At the recovery site, the data set exists on a DASD volume. The batch job ABENDS and execution stops. The situation must be reported and researched to determine what the discrepancy is and how to resolve it. On average, a situation like this would take 30 minutes to resolve. At the recovery site, when time is critical, a delay of 30 minutes for every catalog discrepancy would have a severe impact on recovery time objectives.

What is needed to resolve this situation is to compare the records in all ICF catalogs at the recovery site to the physical data on the volumes and to correct the discrepancies before processing begins. In an DFSMShsm environment, a comparison of the entries in the ICF catalogs to the entries in the Migration Control Data Set (MCDS) is also required. Mainstar provides this functionality in **Catalog RecoveryPlus (CR+)** and the ***Fast*Audit/390** selectable product, **HSM *Fast*Audit**.

The CATSCRUB Command in CR+

The CATSCRUB command in **CR+** compares entries in the specified ICF catalogs to the data sets on the storage volumes at the recovery site. It reads through each Basic Catalog Structure or BCS for each catalog record, and follows the volume cell pointer out to the physical volume to determine if the physical data set for it exists. If it does not exist, then the catalog record is deleted. If the volume cell pointer identifies a VOLSER that is not physically at the recovery site or is not varied online, the record is deleted.

The CATSCRUB command in **CR+** provides user control for anomaly detection and processing. As an example, you can choose to keep all tape data set records and/or all GDG base records.

MCDS Synchronization Using HSM *Fast*Audit/390

For those using DFSMShsm, migrated entries must also be synchronized. **HSM *Fast*Audit** provides an audit facility to quickly compare all records in the ICF catalogs to the records in the MCDS and build FIX statements for subsequent execution.

In a recovery situation, the ICF catalog and the MCDS may not have been backed up at exactly the same time. Depending on your situation, you may select either the ICF catalog or the MCDS as the basis for the audit..

As an example, if you select the ICF catalog as the basis, then all records in the MCDS are compared to the ICF catalogs and if a record in the MCDS is not found in the ICF catalogs, the record is deleted out of the MCDS. Conversely, if you select the MCDS as the basis for the audit, records in the

MCDS are compared to the ICF catalogs, and if a record for a data set that is in the MCDS does not exist in the ICF catalog, a new ICF catalog record is built and the migrated data set is re-cataloged. In addition, a data set's catalog entry may indicate that it's migrated to ML1 when in fact it's migrated to ML2 or to a primary storage volume. The audits in **HSM *Fast*Audit** find these discrepancies as well and automatically build the FIXCDS or IDCAMS commands to resolve them.

Selective Catalog Recovery

In today's data center, several backup and recovery methodologies may be in use to meet the varying needs of the data. As an example, some data may reside on DASD devices that are remotely mirrored to an offsite facility while other data in the same data center are backed up and recovered using ABARS or full volume dump. The ICF catalogs that support the data center may all be on mirrored devices even though not all of the data in the data center is mirrored. The scenario varies from site to site as does the method of providing for ICF catalog recovery.

In a mixed environment where some catalogs need to be recovered full while others need to be empty or scrubbed, software products are needed to support these functions. **CR+** provides the CATSCRUB facility, described previously in this paper, which scrubs entries for data that do not exist on the physical volumes out of the catalogs. However, in some cases a selective recovery of catalog entries may be a better solution. **CR+** provides this functionality in the RECOVER command.

Let's look at an example in which catalog data is recovered into catalogs that are allocated empty at the recovery site using ABARS.

In this scenario, some data resides on volumes that are remotely mirrored and some data are recovered using ABARS. The recovery site is a fully functional data center with existing catalogs for the data located at that site. ABARS is used to restore some catalogs empty while other catalogs reside on volumes that are mirrored.

Using the **CR+** RECOVER command, you can restore into existing catalogs by utilizing the INTO-EMPTY keyword. With this keyword specified, RECOVER makes no attempt to delete and re-define a new BCS, but rather locates an existing BCS and restores the backup records into it. The RECOVER command ensures that the catalog is empty, except for the catalog's self-describing records. If the catalog is not empty, processing fails.

You can restore certain catalog records from the backup into the empty catalog by data set name or by type. As an example, if some of the data in the catalog represents data not recovered by ABARS or data on the mirrored volumes, the catalog records can be recovered by including a data set name mask, such as PAY.**, in the INCLUDE-DSN keyword. Only records from the backup file whose key matches the specified data set name mask are recovered.

In addition, catalog records can be recovered by type. As an example, you can select to recover only nonVSAM DASD entries or only nonVSAM TAPE catalog entries. The INCLUDE-TYPE and EXCLUDE-TYPE keywords also support GDG, VSAM, and VOLCAT (L and W record types).

Conclusion

There isn't a one-size-fits-all solution to catalog recovery at a disaster or alternate site. Data centers need options and flexibility to suit their individual needs. Whether you're an ABARS user or not, Mainstar provides a variety of solutions to help you recover ICF catalogs at the recovery site. In some cases, data centers need multiple solutions because they are using multiple methodologies. We understand this at Mainstar and we provide a variety of product solutions that can be customized to suit your needs.

For more information, visit our web site at www.mainstar.com or e-mail us at product_info@mainstar.com.

Colleen Gordon – Professional Services Manager and SE Manager for Mainstar Software Corporation.
Colleen has an extensive background in Business Continuation and Storage Management and has worked with customers all over the world. Colleen is the co-author of the IBM Redbook, ABARS and Mainstar Solutions, *and has written articles for* Disaster Recovery Journal, Enterprise Systems Journal, *and other industry magazines.*