A Customer Success Story
White Paper
From Mainstar Software Corporation

MAINSTAR

Restoring Catalogs at the Recovery Site

By Ken Rodger

One of the essential elements of a successful disaster recovery (DR) test is the creation of healthy ICF catalogs that accurately reflect the actual environment of the recovery site. Catalog mismatches can result in data sets not being found, or in the inability to restore (and catalog) data sets which already have catalog entries. The latter condition is fairly common when restoring full catalogs, and can be both annoying and time-consuming to resolve.

Mainstar's ICF catalog management tool, **Catalog RecoveryPlus (CR+)**, offers two methods to create healthy ICF catalogs at the recovery site: the selective recovery mechanisms of the RECOVER command, and the CATSCRUB command.

This paper discusses how one **CR+** customer, a financial services company, switched from an internally developed methodology using IDCAMS commands and Catalog Solution (from EMC) to **CR+** CATSCRUB. By doing so, they saved time and money, and also expanded their testing capabilities.

The Original Process

The customer's original process involved fully recovering some user catalogs, partially recovering others, and creating some empty user catalogs with just the aliases related to them.

After using IDCAMS to define all of the user catalogs, the full catalog recoveries were performed using EMC Catalog Solution. IDCAMS was then used to define aliases, GDG bases, and tape data sets in the catalogs to be partially recovered, and those catalogs that would remain empty until populated by application activity. ADRDSSU restores of selected DASD volumes were used for the partial recoveries (and this part of the process was especially susceptible to delays due to duplicate or incorrect catalog entries). Finally, both partially recovered and empty catalogs were populated as logical application recoveries took place.

Additionally, the IDCAMS DEFINE statements that built the GDG bases and tape data set entries were generated daily at the home site via the Catalog Search Interface (CSI) using locally developed programs.

Concerns with the Original Process

In overall elapsed time, the process of restoring nine full user catalogs, nine partially recovered catalogs, and fourteen empty catalogs required 130 minutes to complete.

An additional problem with this method was that one particular user catalog, containing more than 800,000 GDG bases, was never recovered. Using IDCAMS, the time required to define this huge number of GDG bases and the tape data set entries prohibited this catalog from being built. As a result, recovering the applications using this catalog had never been attempted during a DR test, so there was no way to know for certain if recovery was even possible.

Moving to CR+ CATSCRUB

The conversion to **CR+** CATSCRUB was easy, consisting of a straightforward two-step process: simply executing the **CR+** RECOVER then CATSCRUB commands.

The RECOVER syntax is quite simple:

```
RECOVER BCS(icf.ucatnamemask)
        INFILE(INDD) NOLOCK          -
        EXCLUDE-DSN(**.SPFTEMP*.**)
```

The INFILE keyword specifies a DD statement containing **CR+** catalog backups from the home site. RECOVER BCS(**) would have recovered all catalogs, so the customer used a catalog name mask to identify the user catalog groups in each execution of the RECOVER command for fully recovered, partially recovered, and empty catalogs. For those catalogs to be recovered with only GDG bases, an additional keyword needs to be added. The keyword INCLUDE-TYPE(EMPTY-GDG) results in an empty catalog

with GDG base definitions. If a catalog is recovered in this manner, then there will be no need to perform further processing with CATSCRUB.

The catalog recoveries are followed by the **CR+** CATSCRUB processing. As you can see in Figure 1, CATSCRUB is highly flexible, allowing a great deal of process tailoring with global filtering keywords to identify categories of catalog records to be deleted or kept as they are encountered. In this example, catalog records for migrated data sets and data set records in various error conditions are deleted. The result of this 'scrubbing' process leaves all tape data set catalog entries intact. Also, the DASD data set entries now precisely match the data sets actually present when CATSCRUB is run.

At the DR site, within 35 minutes of the **CR+** process start, the nine fully recovered, nine partially recovered, and fourteen empty catalogs are ready, saving 95 minutes over the old process. This translates to a 73% reduction in the time required.

Upon moving to **CR+** CATSCRUB, the daily process of generating the IDCAMS DEFINE statements became obsolete, as all of these definitions are captured by the **CR+** catalog backup, and are automatically applied during a **CR+** catalog recovery. Due to the automatic application of GDG definitions in **CR+**, the recovery of the user catalog containing more than 800,000 GDG bases could at last be tested. This meant that switching to CATSCRUB resolved a major exposure.

Because these catalogs reflect what's actually at the recovery site, the ADRDSSU logical restores are now run without catalog mismatch errors, also greatly improving recovery time, and allowing time for more testing. Pre-test preparation and maintenance have also been reduced. Finally, the applications using the catalog that couldn't be recovered can now test their recovery plans.

```
CATSCRUB                                                          -
      EXECUTE                                                     -
      BCS(icf.ucatnamemask)                                       -
      NONVSAM-MULTI-VOL-ERR(DELETE,RC(05) CONTINUE)               -
      NONVSAM-MIGRATED(DELETE,RC(05) CONTINUE)                    -
      NONVSAM-ALIAS-NO-REALNAME(DELETE,RC(05),CONTINUE)           -
      VSAMSPHERE-MULTI-VOL-ERR(DELETE,RC(01),CONTINUE)            -
      VSAMSPHERE-MIGRATED(DELETE,RC(01),CONTINUE)                 -
      GDS-MULTI-VOL-ERR(DELETE,RC(03),CONTINUE)                   -
      GDS-NOT-FOUND(DELETE,RC(03),CONTINUE)                       -
      GDS-MIGRATED(DELETE,RC(03),CONTINUE)                        -
      MESSAGETEXT(FULL)                                           -
      MATCH-VOLSER(**)
/*
```

Figure 1: The CATSCRUB Command

Conclusion

The ICF catalog facility is the key control structure that provides access to your z/OS data. Protecting that access to your data by preparing your ICF catalogs at the recovery site is therefore crucial to a successful recovery.

CR+ offers simple, flexible, and effective mechanisms to help you succeed at that task. With the **CR+** CATSCRUB command, this financial services company expedites their backup and recovery, and ensures that all of their ICF catalogs are protected. To find out how **CR+** can help your data center meet its unique ICF catalog management challenges, contact us at product_info@mainstar.com.

Ken Rodger – Systems Engineer

Ken Rodger joined Mainstar Software Corporation as a Systems Engineer in 2001 with over twenty years of experience in mainframe performance tuning, capacity planning and storage management. Ken has worked in industries as diverse as oil, healthcare, publishing, human resources consulting, and manufacturing. Ken now works with users to provide comprehensive solutions.