



Best Practices for Forward Recovery of a User Catalog

Preparing for Recovery

Elliot Hamilton and Ron Ferguson

A problem with the recovery of a catalog could mean the difference between an unplanned outage of a few minutes and an outage of several hours. By planning ahead and understanding all of the elements necessary to achieve a successful forward recovery, you can minimize downtime.

Mainstar
Software
Corporation

P.O. Box 4132
Bellevue, WA 98009

www.mainstar.com

info@mainstar.com

While the theory is simple, and all of the technical hard work is performed by the utility program you are using, success depends on the details that you're responsible for.

Plan ahead for a successful catalog recovery.

The true test of a backup strategy is whether you can recover accurately and quickly. Because ICF user catalogs are one of the very few aspects of z/OS that present a single point of failure risk, a speedy and successful recovery is crucial for maintaining data access.

- A single user catalog could provide access to hundreds of thousands of data sets, or even millions of data sets, potentially supporting dozens or hundreds of applications.
- One user catalog might have one purpose, for example supporting your report management distribution system, such as RMDS from IBM. This catalog is critical, since you cannot access any business reports managed by RMDS if this catalog isn't available.
- If your DB2 user catalog becomes unavailable, the DB2 subsystem(s) cannot function.

By planning in advance for a successful recovery, you can minimize the downtime from hours to just a few minutes.

Preparing for Catalog Forward Recovery

The concept of catalog forward recovery is straightforward: SMF records that document changes to the catalog are applied to a point-in-time (PIT) backup copy of the catalog, creating an up-to-date catalog that can replace the damaged or missing catalog.

While the theory is simple, and all of the technical hard work is performed by the utility program you are using, success depends on the details that you're responsible for. If you don't set up good processes and make sure these processes are kept up to date, practiced, and tested on a regular basis by well-trained technical personnel, you could be in for any number of surprises – at the worst possible time.

To achieve successful recoveries, periodically review your processes by asking yourself the following questions.

Are Your Catalogs Backed Up?

Make absolutely certain that all catalogs are backed up. This can't be stated too often – you can't perform forward recovery without a backup. Periodically, list all user catalogs connected to the master catalog for each system, then compare that list to the catalogs actually being backed up.

Are Your Backups Accurate?

Double-check your return code monitoring process, to ensure that return codes indicating structural and logical errors with any catalog are analyzed, and the errors corrected.

Remember: the only reason you take a backup is to provide a recovery capability. If you can't recover, the backup is useless.

Can You Find the Most Recent Backup?

During a recovery, you can't afford any wasted time. You need to be able to quickly and accurately locate the most recent backup for any catalog in your environment. Periodically review your catalog backup procedures so that you are familiar with backup naming conventions, and where the backups are cataloged.

If the catalog requiring the recovery contains the catalog records for the catalog backups, identify the secondary user catalog that contains the duplexed backup copy of the catalog.

If you aren't taking a duplex copy of the catalog backups, modify your catalog procedures to create a second copy, assigning it to a different user catalog (by choosing a different high-level qualifier for the data set name of the backup copy).

Do You Know Your Aliases?

Know the aliases assigned to the catalog requiring recovery, and make certain the recovery process will recover them to the master catalog.

Double-check your backup process to ensure that aliases are backed up with the catalog backup.

Make certain that your recovery process will restore the aliases deleted from the master catalog during recovery on the system where the DELETE is run. These aliases are deleted either by your explicitly-issued IDCAMS DELETE command or by the one issued under the covers by IMPORT.

Any aliases defined since the last backup (the one you're recovering from) will not be on the backup, but will be reflected in the SMF records. ICFRU and other forward recovery utilities will re-apply these to the updated EXPORT-format backup copy of the user catalog that it's creating, and they will subsequently be added to the master catalog during IMPORT processing.

Ensure that all master catalogs that share access to this user catalog have the proper aliases for the user catalog being

recovered. IMPORT does not perform any alias updates to any master catalog other than the one for the system on which IMPORT is executing, leaving the responsibility to you to ensure that aliases are consistent across all master catalogs following forward recovery of a user catalog.

Can You Find the SMF Dumps?

The SMF records keep track of the changes that have been made since your last backup. For an accurate recovery, you need to be able to quickly and accurately locate the SMF dump data sets from all systems in your environment. It's likely that you'll use a GDG for your SMF dump data sets.

By definition, SMF data sets are cataloged data sets. Therefore, they are cataloged in one of your user catalogs. If the catalog you're attempting to recover has these data sets cataloged within it, you have a major challenge on your hands to locate the necessary files. If you follow the recommendation below for pre-extracting the necessary type 6x records for catalog recovery, you should also consider duplexing this file and cataloging the second copy in a different catalog (by assigning it to a different alias with a different high-level qualifier on the data set name). If your plan relies on the entire SMF dump data set being 'fed' into your catalog forward recovery process, this situation will likely result in a very long recovery time.

For an accurate recovery, you need to be able to quickly and accurately locate the SMF dump data sets from all systems in your environment.

Are Your SMF Records Complete?

Ensure you are capturing and retaining the required SMF records *for all systems*. For every update to the ICF catalog environment, the catalog management system routines create an SMF record. The SMFPRMxx member of SYS1.PARMLIB identifies the SMF record types that your systems are recording.

- Type 61 (x'3D') represents the allocation and cataloging of a new ICF catalog record – typically a DEFINE of a new data set (VSAM, non-VSAM, GDG, etc.).
- Type 65 (x'41') represents the deletion and uncataloging of an existing ICF catalog record – typically a DELETE of an existing data set (VSAM, non-VSAM, GDG, etc.).
- Type 66 (x'42') represents an update of any kind to an existing ICF catalog record – such as adding or deleting a volume to the list of available volumes that a data set can occupy, or at the time of a first allocation to a new volume for an existing data set, or ALTERing any data set attributes that are contained in BCS records (such as the data set's name, SMS attributes, addition of new generation data sets, etc.).

- Type 60 (x'3C') represents any update to the VVDS. (This SMF record type is not used for a catalog forward recovery, so you only need to collect it if you are backing up your VVDS with a utility product such as Catalog RecoveryPlus (CR+) that supports forward recovery of the VVDS.)

Are Your SMF Records Being Thrown Out?

Identify the job that is (usually) initiated through an automated process to dump the SMF data when the current SYS1.MANx data set fills and switches.

IFASMFDP is the standard z/OS utility program that is executed in this job to dump and clear the SYS1.MANx data set.

Review the control statements for the IFASMFDP data set to ensure that the required SMF records are being retained.

Check for other possible post-processing jobs within your environment that access and use the dumped SMF files, ensuring that the SMF records you require for catalog forward recovery are not somehow being discarded.

When Are Your SMF Data Sets Dumped?

Know when and how frequently your systems typically switch and dump the SMF data sets. It's not uncommon for some systems to have light SMF recording activity, and the SYS1.MANx data set is force-switched when the automatically-scheduled job is kicked off. Other systems might have heavy SMF recording activity and the SYS1.MANx data set switches on its own, one or more times throughout the day, and kicks off the job to dump the SMF data.

When a catalog forward recovery is required, the more you know about this process, the easier and less error-prone it will be to gather up *all* necessary SMF data, from every system sharing access to the catalog. From some systems, it could require two to three SMF dump data sets; from others, only a single SMF dump data set.

The point is, you can't afford to miss even a single SMF dump data set during a forward recovery, as it will result in an incorrectly recovered catalog. This is a part of the process where there's *no margin for error*.

If you are using ICFRU as your catalog forward recovery utility, a message is issued during the SMF record selection and validation process (program CRURRSV).

Could Your SMF Data Management Be Improved?

Consider these recommendations and ideas for optimizing the forward recovery process through better SMF data management.

- A significant amount of time can be saved during forward recovery processing if a *copy* of the SMF records required for the forward recovery is written to its own unique data set.
- The type 6x SMF records constitute only a small portion of your entire SMF record collection. This way, *only* those records need to be processed at the time of forward recovery. Sorting the SMF records into ascending time sequence is required unless your copy process sorts the records from the various systems sharing access to the catalog into a single file. (Keeping the data sorted ahead of time isn't recommended, as the complexity and variability of the data to be used will increase the error factor.)
- A separate data set containing only the type 6x SMF records can be split out and created at the time the SYS1.MANx data set is dumped with IFASMFDP.
- Create a GDG for the SMF type 6x records that are split out. Set the limit to allow for enough dump cycles of the SYS1.MANx data sets that may occur between the catalog backup jobs. For example, if the catalog backups are taken daily, you'll want to determine the maximum number of times the SYS1.MANx data sets may be dumped in a one-day period, ensuring there are at least 24 hours of type 6x SMF data available. If the catalog backups are taken more frequently than once a day, you can retain less than 24 hours' worth of data.
- Catalog recovery will be faster if the SMF data that you have to process resides in a disk data set, rather than tape. For many installations, however, the sheer volume of SMF data precludes storage of more than one or two days' worth of SMF data.
- Here's an example of what you can do. Create two GDGs, one defined as a disk data set, and one with a generation limit that matches the total amount of disk storage space you can afford for your most recent SMF data. Our suggestion would be to ensure this is at least equal to the amount of SMF data generated between catalog backup jobs, in which case, all of the SMF data you'll need for a catalog recovery will now reside on disk. The other GDG can be defined as a tape data set, and it's where the SMF dump data sets on disk can 'roll off' when their limit is reached on disk.
- An alternative suggestion, if you've chosen the pre-extraction of 6x records as your methodology, is to keep the extracted

6x records in a disk data set, while the large volume of all SMF records are sent straight to tape.

Your installation undoubtedly slices and dices SMF data in several ways already, for the various installation uses that you have for SMF, so there's no single methodology for handling the SMF dump data that works for everyone. You have to decide on the solution that works best for your particular situation. The key, though, is to have the process well thought out and tested, so that it works when you need it most – at the critical time when a catalog failure occurs.

Conclusion

An outage of just one catalog can seriously affect your organization. By implementing solid preparation processes, you can minimize the time needed to recover and build a foundation for an accurate, successful recovery. These processes need to ensure the following:

- Your catalogs backed up
- The backups are accurate
- You can quickly find the most recent backup
- You know your aliases and that the recovery process will restore them
- You can find the SMF dumps
- The SMF records are complete
- You know how when and how frequently SMF data sets are dumped
- Your SMF data management is optimized

As the requirement for 24x7 data access grows, safeguarding your catalogs is more crucial than ever. CR+ provides powerful features for ICF catalog management, backup, and recovery. With effective tools for secure recovery and better day-to-day management, you'll be able to access your data, even if a simple "finger check" goes wrong or a disaster occurs.

For more information visit our web site, www.mainstar.com, or contact us at experts@mainstar.com to arrange a personal briefing.

Elliot Hamilton

Elliot Hamilton has been involved with IBM's MVT, VS1, SVS, XA, and numerous variations of MVS since the mid-1970s. He gained VSAM and catalog expertise through experience as a systems Programmer and as a System Engineer and Instructor for Amdahl Corporation. As one of the early instructors for Mainstar Software Corporation (formerly Software Information Services, Inc.), Elliot presented VSAM and ICF catalog classes around the world. When Mainstar's customers' needs changed, Elliot worked with the Mainstar Software Development staff on numerous projects, including Catalog RecoveryPlus (CR+). Currently, Elliot provides Level 3 Support for CR+, using his expertise in identifying tricky and difficult problems within ICF catalogs to search out brilliant CR+ solutions.

Ron Ferguson

Mainstar's founder, 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.

Mainstar is a registered trademark of Mainstar Software Corporation. Catalog RecoveryPlus is a trademark of Mainstar Software Corporation. Mainstar Software Corporation is a wholly owned subsidiary of Rocket Software, Inc. All other products or company names are used for identification purposes only and may be trademarks of their respective owners.

010-0114-01 (01/10/07)