

Ensuring Correct GDS Recovery

By Janet Sun

IBM DFSMSHsm Level 2 Support contacted Mainstar regarding a customer question about Generation Data Set (GDS) recovery. The customer reported that they had recovered generations of a Generation Data Group (GDG) and that the results were not as they expected. Had they uncovered a bug in DFSMSHsm or was something unusual happening at their recovery site? Upon our review it was clear that what happened at this recovery site could baffle even the best of technicians. We concluded that this scenario should be shared with all of our readers.

Consider the following scenario:

At the customer's home site a GDG with a limit of 3 currently has the following generations cataloged:

- G0001V00 relative generation -2
- G0002V00 relative generation -1
- G0003V00 relative generation 0

If an ABACKUP of all of the generations of this GDG is run at the home site and recovered with ARECOVER at the disaster recovery site, then the generations at the recovery site will be:

- G0001V00 relative generation -2
- G0002V00 relative generation -1
- G0003V00 relative generation 0

Back at the home site, new generations are created and cataloged as follows:

- G0004V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

Another ABACKUP is run at the home site and recovered at the disaster recovery site. The generations reflected at the recovery site are now:

- G0004V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

At the home site, the entire GDG is deleted and reallocated and the following generations are created:

- G0001V00 relative generation -2
- G0002V00 relative generation -1
- G0003V00 relative generation 0

Ensuring Correct GDS Recovery

When these data sets are backed up with ABACKUP at the home site and then recovered at the disaster recovery site, instead of having generations G0001V00, G0002V00 and G0003V00 cataloged, this is the result:

- G0003V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

This was not the desired result! How did this happen? The answer has to do with the way in which the relativity of the generations in a GDG is maintained. Even though a particular generation may have been created chronologically later, it is the generation number itself (the *nnnn* part of *GnnnnV00*) which will govern the relativity of the generations. Let's look at what happened.

At the disaster recovery site, all data sets from previous recoveries still existed. Before the last set of recoveries, here is what the catalog indicated exists in the GDG:

- G0004V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

The customer proceeded to perform the ARECOVER for G0001V00. Because they are recovering an earlier numbered generation than is currently cataloged, G0004V00 is rolled off to maintain the limit of 3 and replaced by G0001V00. This results in the following set of data:

- G0001V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

Then the customer recovered G0002V00. Because G0001V00 is earlier or lower by generation number, it is rolled off and replaced with G0002V00. Here is what the data sets look like at this point in time:

- G0002V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

Finally the customer recovered G0003V00. Once again, G0002V00 is earlier by generation number, so it is rolled off and replaced with G0003V00. Here are the results:

- G0003V00 relative generation -2
- G0005V00 relative generation -1
- G0006V00 relative generation 0

It was clear to us at Mainstar that the existence of generations 4 through 6 in the catalog prior to the last recovery caused the undesirable results.

How can you prevent this problem from occurring? One option is to use Incremental ABARS, a selectable feature of Backup & Recovery Manager Suite: ABARS Manager. Incremental ABARS captures the GDG base information taken at the time of backup. Incremental ABARS keeps track of when each generation is allocated and ensures that the correct generations exist after recovery by automatically deleting the out of date generations.

Here is how the same scenario would work using Incremental ABARS:

Ensuring Correct GDS Recovery

When the three generation data sets G0001V00, G0002V00, and G0003V00 are backed up with ABACKUP, they are kept as a collective set. Each time the GDG is backed up with ABACKUP, the generation data sets will be grouped in a collective set and each collective set is assigned a relative version. The most recently backed up copy of the GDG collective set will be relative version 0. At the disaster recovery site, one aggregate will be recovered and Incremental ABARS would then recover all of the data in the collective set.

The result would be that the GDG base would be restored into the catalog with a limit of three, and the most current generation, G0003V00, would be recovered, followed by G0002V00 and G0001V00 (relative generations 0, -1, -2). If there had been an earlier aggregate backup which contained G0004V00, G0005V00, and G0006V00, it would also be in a collective set, but would represent the -1 version of that collective set. Incremental ABARS will recover the most current, or 0 version, of the collective set unless told to do otherwise.

The following **ABARS Manager** screen shows how the incremental ABACKUP events are managed by **ABARS Manager**. Application EDGINCR1 has three generations of aggregate backups with the relative generations as follows:

- Backup date 30 September 2004 relative version 0
- Backup date 20 September 2004 relative version -1
- Backup date 10 September 2004 relative version -2

All of the data sets included in each of the three generations are part of one incremental set. Only the most current version of any data set and only the active generations of any GDG base, as of the most current incremental backup, are automatically recovered.

Menu Diagnostics Preferences												
ABM - ABARS Events										Row 57 of 213		
Command ==>										Scroll ==> PAGE		
Primary Cmds: REFresh,RESet,H,M,SD,SE,SN,SS										Date => 2004/11/09		
Line Cmds: N,S,SP,SM,VA,VC,VE,VI,VR,VS,+, -,DE										Time => 15:28		
B,C,I,M,O,R,RD,BI,EI,BL,EL,BS,ES,BR,ER,BC,EC										Sort => N D,E,N,S		
C	Name	+-	Gen	Incr	T	PGM	RCA	RCB	Date	Time	Total DSNs	Total Space
—	*	—	—	—	—	—	—	—	—	—	—	—
	EDGINCR1	-	0	0	I	AB	0	4	2004/09/30	15:34	1	55.33 Kb
	- EDGINCR1		0	-1	I	AB	0	0	2004/09/20	12:47	1	55.33 Kb
	- EDGINCR1		0	-2	BI	AB	0	0	2004/09/10	12:36	3	184.81 Kb
	EDG1	+	0		B	AB	-	3	2004/11/05	11:20	0	.00 b
	EM01		0		B	AB	0	0	2002/01/09	13:01	15	1.91 Mb
	IMX		0		B	AB	0	0	2003/03/24	09:46	6	436.00 Kb
***** Bottom of data *****												

Figure 1. Managing the incremental ABACKUP events in ABARS Manager

Another way to prevent this problem is to delete the catalog entries for the GDG with generations 4 through 6 in it prior to running the ARECOVER for the GDG with generations 1 through 3.

Removing the catalog entries at the disaster recovery site can be accomplished in a variety of ways. One option is to use CATSCRUB, which will soon be a selectable feature of **ABARS Manager**. If the generation data sets do not physically exist on DASD volumes at the recovery site, the CATSCRUB command can be used to remove the catalog entries.

Another option is to use **Catalog RecoveryPlus (CR+)** to selectively recover catalog entries that will not otherwise be recovered with ABARS. Using the selective recovery feature ensures that the catalogs will only contain entries for non-ABARS recovered data.

Ensuring Correct GDS Recovery

For more information on GDGs, see the Mainstar Information Document entitled, "Misconceptions and Mysteries About Generation Data Groups".

Janet Sun – Senior Systems Engineer for Mainstar Software Corporation.

Janet Sun has over 20 years of mainframe systems programming experience in the areas of operating system management, performance management, storage management, data security, web server management, and disaster recovery planning. She frequently presents at user conferences such as SHARE and the IBM Storage Symposium.

Mainstar is a registered trademark and Backup & Recovery Manager, ABARS Manager, Catalog RecoveryPlus, CATSCRUB, and Incremental ABARS are trademarks of Mainstar Software Corporation.

IBM is a registered trademark and DFSMSHsm is a trademark of the International Business Machines Corporation.