



**I** I invoked the FSEDIT procedure on my data set and received a dialog box that told me the requested data set was damaged and asked if I wanted it repaired. How could the data set have become damaged?

**O** A data set could become damaged when changes are made to the observations in the data set and the SAS® System is not able to go through its normal closing routines on the data set due to some form of abnormal termination. In this case, changes written to memory are lost because SAS never has the opportunity to write them to disk. A typical abnormal termination for TSO users would be if your SAS session is timed out while you are making changes to your data set.

SAS maintains a damaged flag for data sets and their related index files. The damaged flag is set on the disk just before the first update occurs. The flag is reset when updates made to data in memory are written to disk. The purpose of the flag is to detect when updates made to data in memory were not written to disk. Keeping data in memory improves performance but opens a window for loss of data when the computing environment terminates abnormally. If the damaged flag is still set when a data set is opened, you receive the error that your data set is damaged.

The fact that changes made to the data set may not have been saved is an important piece of information that the developers felt SAS users needed to know. In Version 5 of the SAS System, an abnormal termination could lose changes, but because you were not alerted to this you may not have noticed for quite some time. The damaged data set message was added in Version 6 to alert you to possible problems the next time you access the data set, whether it be with the FSEDIT procedure or any other SAS statement that opens the data set.

The FSEDIT procedure edits a SAS data set in place. When you make a change to an observation and leave that observation by scrolling to another observation or adding a new observation, the damaged flag is set to indicate the data set has been changed and the changes are stored in memory. If you were to scroll back to the observation you changed, you would see the changes you just made.

Your changes are actually written from memory to disk when you either issue the SAVE command from the FSEDIT window, reach the AUTOSAVE value set for your FSEDIT session, or end your FSEDIT session. Buffers are also flushed at the discretion of the system when buffer space is needed for I/O. The damaged flag that indicates the data set has been changed is also reset at this time. If your SAS session abnormally terminates before one of these conditions is met, the SAS System is not able to perform its normal closing routines and changes

to the data set are not written from memory to disk. The damaged flag also remains set to indicate the data set has been changed.

If you attempt to read this data set again while the damaged flag is still set, the SAS System displays the message that your data set is damaged. This message indicates that changes were made to your data set that may not have been saved to disk.

If you respond to the message that you want the data set repaired, every page of the data set is read, looking for inconsistent structural data. The index is also examined for structural damage, such as if the index file is missing or has a number of update cycles that differs from the data set. If there are problems found with the index, the SAS System re-creates the index when possible. While simple indexes may be re-created, composite indexes cannot always be rebuilt.

Even though the data set has been repaired, changes made to the data set that were saved to memory and not written to disk are no longer in your data set. The message that your data set is damaged is to alert you that changes may have been lost. If your SAS session has abnormally terminated, then changes that were written to memory but not to disk cannot be recovered.

One way to reduce the chances of a damaged data set is to set the AUTOSAVE value in your FSEDIT session to 1. This value is stored with your modified FSEDIT screen in the General Parameters window. With AUTOSAVE set to 1, as soon as you make a change to an observation and leave the observation, the damaged flag is set, the update is made in memory, the update is written to disk, and the damaged flag is cleared. You should check to see if this option is going to be feasible with your data because saving a data set so frequently will degrade performance.

There may be more severe situations that cause a data set to be damaged other than the situation described here. If the entire operating system is abnormally terminated instead of just the SAS session, the data set may actually be corrupted, and attempts to repair the data set may not be successful. In this case, you would need to go to a backup of your data set.

*Answer provided by Ann Carpenter. Ann is a senior technical support analyst at SAS Institute where she specializes in the SAS/FSP® and SAS/ENGLISH® products. Copyright © 1994 by SAS Institute Inc.*

SAS, SAS/ENGLISH, and SAS/FSP are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.