



CHAPTER 37

SAS File Management

<i>Improving Performance</i>	519
<i>Moving SAS Files Between Operating Environments</i>	520
<i>Converting SAS Files</i>	520
<i>Repairing Damaged Files</i>	520
<i>Recovering SAS Data Files</i>	521
<i>Recovering Indexes</i>	522
<i>Recovering Catalogs</i>	522

Improving Performance

The SAS System offers tools to control the use of memory and other computer resources. Most SAS applications will run efficiently in your operating environment without using these features. However, if you develop applications under the following circumstances, you may want to experiment with tuning performance:

- You work with large data sets.
- You create production jobs that run repeatedly.
- You are responsible for establishing performance guidelines for a data center.
- You do interactive queries on large SAS data sets using SAS/FSP software.

The following table summarizes tools available to affect performance, and specifies where you can find documentation on the tools:

Table 37.1 Performance Tools Summary

For information about ...	See ...
the time required to run your application	STIMER or FULLSTIMER system options in the SAS documentation for your operating environment.
data set characteristics	CONTENTS statement for the DATASETS procedure in <i>SAS Procedures Guide</i> , the ATTRC and ATTRN functions in <i>SAS Language Reference: Dictionary</i> , and the DICTIONARY tables component for the SQL procedure in <i>SAS Procedures Guide</i> .
setting buffer size (page size)	BUFSIZE= data set option or system option in <i>SAS Language Reference: Dictionary</i> .
setting the number of page buffers	BUFNO= data set option or system option in <i>SAS Language Reference: Dictionary</i> .

For information about ...	See ...
compressing SAS data sets	COMPRESS= data set option or system option and the REUSE= data set option in <i>SAS Language Reference: Dictionary</i> .
indexing SAS data sets	“SAS Indexes” on page 433
programming more efficiently	<i>SAS Programming Tips: A Guide to Efficient SAS Processing</i>
programming with views	Chapter 29, “SAS Data Views,” on page 455

In addition, see the SAS documentation for your operating environment.

Moving SAS Files Between Operating Environments

The procedures for moving SAS files from one operating environment to another vary according to your operating environment, the member type and version of the SAS files you want to move, and the methods you have available for moving the files.

For details on this subject, see *Moving and Accessing SAS Files across Operating Environments*.

Converting SAS Files

Version 8 provides access to Version 8 and Version 7 SAS files, and to most Version 6 SAS files, without converting them. That is, when you migrate to Version 8, you can continue accessing your existing data as well as operating with both Version 6 and Version 8 simultaneously.

Accessing a SAS data library and its members is essentially the same in Version 8 as it is in Version 6. Depending on the type of SAS file and the SAS version being used, compatibility issues are generally handled automatically by the SAS System, by specifying an engine name in a LIBNAME statement or with the ENGINE= system option, or by converting a file.

For details, see Chapter 34, “Compatibility of Version 8 with Earlier Releases,” on page 493.

Repairing Damaged Files

The base engine detects possible damage to SAS data files (including indexes, integrity constraints, and the audit file) and SAS catalogs and provides a means for repairing some of the damage. If one of the following events occurs while you are updating a SAS file, SAS can recover the file and repair some of the damage:

- A system failure occurs while the data file or catalog is being updated.
- Damage occurs to the storage device where a data file resides. In this case, you can restore the damaged data file, the index, and the audit file from a backup device.
- The disk where the data file (including the index file and audit file) or catalog is stored becomes full before the file is completely written to it.
- An input/output error occurs while writing to the data file, index file, audit file, or catalog.

When the failure occurs, the observations or records that were not written to the data file or catalog are lost and some of the information about where values are stored is inconsistent. The next time SAS reads the file, it recognizes that the file's contents are damaged and repairs it to the extent possible in accordance with the setting for the DLDMGACTION= data set option or system option, which is available starting with Version 7.

Note: SAS is unable to repair or recover a view (a DATA step view, an SQL view, or a SAS/ACCESS view) or a stored compiled DATA step program. If a SAS file of type VIEW or PROGRAM is damaged, you must recreate it. Δ

Note: If the audit file for a SAS data file becomes damaged, you will not be able to process the data file until you terminate the audit trail. Then, you can initiate a new audit file or process the data file without one. Δ

Recovering SAS Data Files

To determine the type of action SAS will take when it tries to open a SAS data file that is damaged, set the data set option or system option DLDMGACTION=. That is, when a data file is detected as damaged, SAS will automatically respond based on your specification as follows:

DLDMGACTION=FAIL

tells SAS to stop the step without a prompt and issue an error message to the log indicating that the requested file is damaged. This specification gives the application control over the repair decision and provides awareness that a problem occurred.

To recover the damaged data file, you can issue the REPAIR statement in PROC DATASETS, which is documented in the *SAS Procedures Guide*.

DLDMGACTION=ABORT

tells SAS to terminate the step, issue an error message to the log indicating that the request file is damaged, and abort the SAS session.

DLDMGACTION=REPAIR

tells SAS to automatically repair the file and rebuild indexes, integrity constraints, and the audit file as well. If the repair is successful, a message is issued to the log indicating that the open and repair were successful. If the repair is unsuccessful, processing stops without a prompt and an error message is issued to the log indicating the requested file is damaged.

Note: If the data file is large, the time needed to repair it can be long. Δ

DLDMGACTION=PROMPT

tells SAS to provide the same behavior that exists in Version 6 for both interactive mode and batch mode. For interactive mode, SAS displays a requestor window that asks you to select the FAIL, ABORT, or REPAIR action. For batch mode, the files fail to open.

For a data file, the date and time of the last repair and a count of the total number of repairs is automatically maintained. To display the damage log, use PROC CONTENTS as shown below:

```
proc contents data=sasuser.census;
run;
```

Output 37.1 Output of CONTENTS Procedure

```

                                The CONTENTS Procedure

Data Set Name: SASUSER.CENSUS      Observations:      27
Member Type:   DATA              Variables:         4
Engine:        V8                 Indexes:           0
Created:       12:39 Monday, January 4, 1999  Observation Length: 32
Last Modified: 11:30 Tuesday, January 5, 1999 Deleted Observations: 0
Protection:                               Compressed:        NO
Data Set Type:                               Sorted:           NO
Label:

-----Engine/Host Dependent Information-----

Data Set Page Size:      8192
Number of Data Set Pages: 1
First Data Page:        1
Max Obs per Page:       254
Obs in First Data Page: 27
Number of Data Set Repairs: 1
Last Repair:            12:46 Tuesday, January 5, 1999

```

Recovering Indexes

In addition to the failures listed earlier, you can damage the indexes for SAS data files by using an operating environment command to delete, copy, or rename a SAS data file, but not its associated index file. The index is repaired similarly to the `DLDMGACTION=` option as described for SAS data files, or you can use the `REPAIR` statement in `PROC DATASETS` to rebuild composite and simple indexes that were damaged.

You cannot use the `REPAIR` statement to recover indexes that were deleted by one of the following actions:

- copying a SAS data file by some means other than `PROC COPY` or `PROC DATASETS`, for example, using a `DATA` step
- using the `FORCE` option in the `SORT` procedure to write over the original data file.

In the above cases, the index must be rebuilt explicitly using the `PROC DATASETS INDEX CREATE` statement.

Recovering Catalogs

To determine the type of action that SAS will take when it tries to open a SAS catalog that is damaged, set the system option `DLDMGACTION=`. Then when a catalog is detected as damaged, SAS will automatically respond based on your specification.

Note: There are two types of catalog damage:

- localized damage* is caused by a disk condition, which results in some data in memory not being flushed to disk. The catalog entries that are currently open for update are marked as damaged. Each damaged entry is checked to determine if all the records can be read without error.
- severe damage* is caused by a severe I/O error. The entire catalog is marked as damaged.

Δ **DLDMGACTION=FAIL**

tells SAS to stop the step without a prompt and issue an error message to the log indicating that the requested file is damaged. This specification gives the application control over the repair decision and provides awareness that a problem occurred.

To recover the damaged catalog, you can issue the REPAIR statement in PROC DATASETS, which is documented in the *SAS Procedures Guide*. Note that when you use the REPAIR statement to restore a catalog, you receive a warning for entries that have possible damage. Entries that have been restored may not include updates that were not written to disk before the damage occurred.

DLDMGACTION=ABORT

tells SAS to terminate the step, issue an error message to the log indicating that the requested file is damaged, and abort the SAS session.

DLDMGACTION=REPAIR

for localized damage, tells SAS to automatically check the catalog to see which entries are damaged. If there is an error reading an entry, the entry is copied. If an error occurs during the copy process, then the entry is automatically deleted. For severe damage, the entire catalog is copied to a new catalog.

DLDMGACTION=PROMPT

for localized damage, tells SAS to provide the same behavior that exists in Version 6 for both interactive mode and batch mode. For interactive mode, SAS displays a requestor window that asks you to select the FAIL, ABORT, or REPAIR action. For batch mode, the files fail to open. For severe damage, the entire catalog is copied to a new catalog.

Unlike data files, a damaged log is not maintained for a catalog.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *SAS Language Reference: Concepts*, Cary, NC: SAS Institute Inc., 1999. 554 pages.

SAS Language Reference: Concepts

Copyright © 1999 SAS Institute Inc., Cary, NC, USA.

ISBN 1-58025-441-1

All rights reserved. Printed in the United States of America. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, by any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute, Inc.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of the software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, November 1999

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries.® indicates USA registration.

IBM, ACF/VTAM, AIX, APPN, MVS/ESA, OS/2, OS/390, VM/ESA, and VTAM are registered trademarks or trademarks of International Business Machines Corporation. ® indicates USA registration.

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

The Institute is a private company devoted to the support and further development of its software and related services.